

DEFINABILITY AND CATEGORICITY IN CONTINUOUS LOGIC

By

James E. Hanson

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

(MATHEMATICS)

at the

UNIVERSITY OF WISCONSIN – MADISON

2020

Date of final oral examination: April 21, 2020

The dissertation is approved by the following members of the Final Oral Committee:

Uri Andrews, Associate Professor, Mathematics

Omer Mermelstein, Assistant Professor, Mathematics

Joseph S. Miller, Professor, Mathematics

Andreas Seeger, Professor, Mathematics

Abstract

After a self-contained development of continuous first-order logic, we study the phenomena of definability and categoricity in continuous logic.

The classical Baldwin-Lachlan characterization of uncountably categorical theories is known to fail in continuous logic in that not every inseparably categorical theory has a strongly minimal set. We investigate these issues by developing the theory of strongly minimal sets in continuous logic and by examining inseparably categorical expansions of Banach space.

To this end, we introduce and characterize ‘dictionary theories,’ theories in which definable sets are prevalent enough that many constructions familiar in discrete logic can be carried out, and we show that ω -stable theories and randomizations of arbitrary continuous theories are dictionary. We also introduce, in the context of Banach theories, ‘indiscernible subspaces,’ which we use to improve a result of Shelah and Usvyatsov [SU19]. Both of these notions are applicable outside of the context of inseparably categorical theories.

We construct or present a slew of counterexamples, including an ω -stable theory with no Vaughtian pairs which fails to be inseparably categorical and an inseparably categorical theory with strongly minimal sets in its home sort only over models of sufficiently high dimension.

In order to investigate notions of approximate categoricity, we give a formalism for approximate isomorphism in continuous logic simultaneously generalizing those of Ben Yaacov [BY08b] and Ben Yaacov, Doucha, Nies, and Tsankov [BDNT17], which are

largely incompatible. We introduce distortion systems, which are a mild generalization of perturbation systems [BY08b]. With this we explicitly exhibit Scott sentences for perturbation systems, such as the Banach-Mazur distance and the Lipschitz distance between metric spaces. Our formalism is simultaneously characterized syntactically, by distortion systems, and semantically, by certain elementary classes of two-sorted structures that witness approximate isomorphism. We also make progress towards an analog of Morley's theorem for inseparable approximate categoricity, showing that if there is some uncountable cardinal κ such that every model of size κ is 'approximately saturated,' in the appropriate sense, then the same is true for all uncountable cardinalities. Finally, we present some non-trivial examples of these phenomena and highlight an apparent interaction between ordinary separable categoricity and inseparable approximate categoricity.

Acknowledgments

I would like to thank my advisor, Uri Andrews for his mentorship and assistance and for teaching the class on continuous logic which originally spurred my interest in the topic.

I would also like to thank Professors Mariya Soskova, Joseph Miller, and Steffen Lempp for their help and advice and for making the UW–Madison Logic Group so welcoming to students.

For their feedback on my thesis or some of its constituent papers, I would like to thank Professors John Baldwin, C. Ward Henson, and Tomás Ibarlucía. I would also like to thank Professors Isaac Goldbring, Bradd Hart, and H. Jerome Keisler for stimulating discussions involving research, and I would like to thank Professors Kyle Gannon and Omer Mermelstein for stimulating discussions *not* involving research.

I'd like to thank Patrick and Dominic and Stephen for remotely helping to keep me sane since roughly January 2020. And Phil, for locally keeping me sane since approximately 2015.

And deepest thanks

To my brother Andrew for his regular encouragement,

To my father Bradley for his constant support,

And to my mother Diane for her continuous love

and infinite patience.

Contents

Abstract	i
Acknowledgments	iii
List of Figures	x
List of Tables	xi
Introduction and Preliminaries	1
0.1 Outline	3
0.2 Comparison of Discrete Logic and Continuous Logic	5
0.3 Metric Space Preliminaries	9
1 Syntax and Semantics	17
1.1 Signatures	17
1.2 (Pre-)structures	18
1.3 Terms and Formulas	23
1.4 The Logical Norm and the Density of Restricted Formulas	38
1.5 Compactness I: Type Spaces	45
1.6 Compactness II: Ultraproducts	51
1.7 Logical Operations on Topological Formulas	56
1.8 Saturated, Special, and Monster Models	66

2	Topometry and Definability	76
2.1	Topometric Spaces	77
2.2	Example: The Halo	85
2.3	Locatable and Definable Sets	87
2.3.1	Locatable Sets	88
2.3.2	Definable Sets	98
2.3.3	Examples: The Discrete Interval and the Polarized Square	113
2.3.4	Finite Counting Quantifiers and Algebraic Sets	128
2.3.5	Definable Functions	135
2.4	Dictionary Type Spaces and Theories	145
2.5	Small Type Spaces and ω -stability	166
3	Imaginaries and Interpretations	177
3.1	Many-Sorted Signatures	177
3.2	Imaginary Sorts and T^{eq}	183
3.3	T^{eq} for Dictionary T	200
3.4	Additional Imaginary Constructions	201
3.5	Compact Parameter Sorts and Almost Uniformity	214
3.6	Interpretations	220
4	Minimality and Categoricity	223
4.1	(Strongly) Minimal Sets in Continuous Logic	226
4.1.1	Some Characterizations	228
4.1.2	The Pregeometry of a Strongly Minimal Set	233
4.1.3	Strongly Minimal Types	236

4.1.4	Approximately (Strongly) Minimal Pairs	241
4.2	Categoricity	248
4.3	Vaughtian Pairs	249
4.4	Theories with Strongly Minimal Sets over the Prime Model	252
4.4.1	Main Theorem	252
4.4.2	Strongly Minimal Sets over the Prime Model	254
4.4.3	The Number of Separable Models	256
4.5	Generalizations of Strongly Minimal Sets	261
4.5.1	Orthonormalizable Sets	261
4.5.2	Peripheral Types	264
5	Structures and Theories	272
5.1	Ultrametric Structures	272
5.2	Locally Compact Structures	278
5.3	Strongly Minimal Theories	280
5.3.1	Strongly Minimal Theories That Do Not Interpret Infinite Discrete Theories	280
5.3.2	Characterization of Strongly Minimal Groups	291
5.4	Banach and Hilbert Structures	305
5.4.1	Banach Theory Background	306
5.4.2	Asymptotically Hilbertian Spaces Do Not Interpret a Strongly Minimal Set	310
5.4.3	Indiscernible Subspaces	311
5.4.4	Minimal Wide Types	315

5.5	Randomizations	320
6	Approximate Isomorphism and Approximate Categoricity	328
6.1	Approximate Isomorphism	331
6.1.1	Induced Metrics on Type Space	344
6.1.2	Theories of Approximate Isomorphism	348
6.2	Special Cases	354
6.2.1	Elementary and Finitary Gromov-Hausdorff-Kadets Distances . .	354
6.2.2	Banach-Mazur Distance	358
6.2.3	Approximate Isomorphism in Discrete Logic	364
6.3	Scott Sentences for Functional Approximation Fragments	365
6.4	Parameters in Distortion Systems and d_Δ	374
6.5	Approximately Atomic Types	377
6.6	Approximate Categoricity	383
6.6.1	Separable Categoricity	384
6.6.2	Inseparable Approximate Categoricity	393
6.6.3	Some Examples and the Relationship between Different Notions of Categoricity	406
A	A Compendium of Relevant Topological Facts	413
A.1	(Pseudo-)metric Spaces	413
A.2	Compact Hausdorff Spaces	418
B	Continuous Analogs of Classical Discrete Results	423
B.1	Some Results that Translate Directly	423

B.2	Two Proofs of Morley’s Theorem	430
B.2.1	Common Material	431
B.2.2	Ben Yaacov’s Proof after Morley	436
B.2.3	A Proof after Chang and Keisler	446
B.3	Encoding Structures as Metric Spaces	452
B.3.1	Preliminaries	455
B.3.2	Expansions	456
B.3.3	Countable Disjoint Unions of Sorts	460
B.3.4	Making Everything Lipschitz	467
B.3.5	Encoding in Metric Spaces	472
B.4	On ‘Finite Axiomatizability’ in Continuous Logic	482
B.5	On the Definition of NSOP in Continuous Logic	485
C	Counterexamples in Continuous Logic	496
C.1	Topometry and Definability	496
C.2	Minimality and Categoricity	506
C.2.1	A Theory with Strongly Minimal Sets, but Only over Models of Dimension $\geq n$	513
C.2.2	A Counterexample to a Direct Translation of the Baldwin-Lachlan Characterization	522
C.3	Banach and Hilbert Structures	525
C.3.1	No Infinitary Ramsey-Dvoretzky-Milman Phenomena in General .	525
C.3.2	(Type-)Definability of Indiscernible Subspaces and Complex Banach Structures	527

C.3.3	Non-minimal Wide Types	531
C.3.4	Minimal Wide Types That are Not Strongly Minimal Wide	532
C.4	Approximate Isomorphism and Approximate Categoricity	542
C.4.1	An Irregular Distortion System	544
D	More on Definable Sets	548
D.1	Almost Clopen Sets	548
D.2	Dictionary Type Spaces and Generic Separators	551
D.3	Definable Sets in Type Spaces Homeomorphic to $[0, 1]$	566
	List of Symbols	570
	Metric Spaces	570
	Signatures, Terms, and Formulas (Single-Sorted)	571
	(Pre-)structures	575
	Theories	577
	Type Space and Other Topometric Spaces	578
	Many-Sorted Signatures	579
	Bibliography	582

List of Figures

1	\mathfrak{H} and a, b, c	85
2	Minimizing $d(ac, b'c')$ for $b'c' \equiv bc$	85
3	The type spaces $S_{xy}(T)$ and $S_x(T_c)$	87
4	$S_1(\text{DI})$ with a cartoon of its ‘microscopic’ metric topology.	114
5	$S_1(\text{PS})$ with a cartoon of its ‘microscopic’ metric topology.	116
6	A locatable subset L of $S_1(\text{PS})$	117
7	Separating $p_{a,b}$ and $p_{c,e}$ with definable neighborhoods.	118
8	D from Proposition 2.3.49	124
9	L from Proposition 2.3.54	127
10	$d(p, o) \leq \sqrt{2(1 - d_{\text{inf}}(p, \mathfrak{H}))}$ in $S_1(\mathfrak{H})$	265
11	$\alpha(\varepsilon)$ in Proposition A.1.8 part (i)	416
12	$\gamma(\delta)$ and $\beta(\delta)$ in Proposition A.1.8 part (ii)	416
13	$\llbracket Q = r \rrbracket$ is not definable for any $r \in [0, 1]$	500
14	Connected \mathbb{R}^2 -embeddable dictionary type space with no d -atomic points	501
15	Dictionary type space with a closed subspace that fails to be dictionary	503
16	Infinite type space with precisely 3 definable sets.	504
17	Two sided definable endpoint that is not d -atomic	505

List of Tables

1	General Translation	5
2	Types and Type Spaces	7
3	Imaginaries	7
4	acl/dcl and Definable (Partial) Functions	8
5	ω -stability	8
6	Models	9
7	Combinations of approximate categoricity conditions.	406

Introduction and Preliminaries

Continuous logic is a generalization of discrete first-order logic designed to study classes of structures with complete metrics and uniformly continuous functions, referred to as metric structures [BYBHU08, BYU10]. With it one can apply methodology common in model theory to structures such as Banach spaces and Banach spaces with ‘accessories,’ e.g. C^* -algebras and other operator algebras, systems, or spaces; probability algebras, possibly with a designated group action such as those studied in ergodic theory; Berkovich spaces and other \mathbb{R} -trees; and *randomizations* [BYJK09], structures encoding a family of random variables taking values in another structure.

For the last forty or so years, model theory has been dominated by stability theory and its generalizations. These approaches focus on dividing lines, subtle combinatorial tameness properties of theories that give strong structural information about definable sets in their models. Identifying logically tame structures in everyday mathematics can be extremely fruitful. For instance, Hrushovski’s lauded proof of the geometric Mordell-Lang conjecture in arbitrary characteristic relies heavily on stability theory. Continuous logic expands the applicability of stability theory by taming more structures of interest. One of the most drastic known examples of this is the class of Hilbert spaces, which when treated as discrete structures are not as tame as they ‘ought’ to be.

A discrete first-order theory T is called κ -categorical, for a cardinality κ , if any two models of T of cardinality κ are isomorphic. Morley showed in the 1960s that a countable theory that is categorical in some uncountable cardinality is categorical in every uncountable cardinality. Familiar examples of this are infinite dimensional vector

spaces over a fixed field and algebraically closed fields of a fixed characteristic. In the 1970s, Shelah's work on generalizing Morley's theorem to uncountable languages led to the development of stability theory.

Also in the 1970s, Baldwin and Lachlan developed a better structural understanding of uncountably categorical theories as part of their characterization of the countable models of uncountably categorical theories. Central in this structural understanding is the concept of a strongly minimal set, a definable set in which every definable subset is finite or co-finite. Strongly minimal sets have good notions of dimension, analogous to vector space dimension or transcendence degree in fields. The structural picture that Baldwin and Lachlan gave is that a theory is uncountably categorical if and only if it is ω -stable (which ensures the presence of strongly minimal sets) and has no 'Vaughtian pairs,' a condition that ensures that a model is entirely determined by any infinite definable subset of it, such as in particular any strongly minimal set. Since strongly minimal sets are entirely determined by dimension, the class of models of an uncountably categorical theory is very simple. An uncountable model can only have a dimension equal to its cardinality and a countable model can only have countable or finite dimension and dimension is closed upwards, implying that the theory can only have either 1 or \aleph_0 many countable models.

In the 2000s, Ben Yaacov was able to generalize the original proof of Morley's theorem to continuous logic [BY05] (with the understanding that 'size' is measured by density character and non-compactness, rather than cardinality), and Usvyatsov independently gave a proof using different methods [SU11]. Just as in discrete logic, a countable theory that is categorical for some uncountable cardinality is categorical for all uncountable cardinalities.

The structural picture of Baldwin and Lachlan does not work in continuous logic, despite the fact that there is a natural generalization of strong minimality to continuous logic. Hilbert spaces are uncountably categorical but have no strongly minimal sets. This is particularly tantalizing because Hilbert spaces have an exceedingly clear notion of dimension.

0.1 Outline

This thesis is split into roughly two themes, although these themes are not neatly reflected in the organization of its chapters and sections. We start with some introductory and preliminary material followed by a largely self-contained formalization and exposition of continuous logic in Chapter 1. Chapter 2 explores various issues surrounding the notion of definable sets in continuous logic, the first of our two themes. In that chapter, we introduce and characterize a novel class of continuous first-order theories, *dictionary theories*, in which definable sets are in some sense prevalent enough that many constructions common in discrete logic can be carried out. While we use this for our main results, this machinery ought to be of general use outside of the context of the particular theorems in this thesis. Chapter 3 finishes some expository material involving many-sorted signatures and imaginary sorts but also contains some further material regarding definable sets. Furthermore, Section 5.5 contains some results regarding definable sets in randomizations, Appendix C contains many counterexamples, and Appendix D contains more results both regarding definable sets in general.

Chapters 4, 5, and 6 contain the second of our two themes, categoricity and approximate categoricity of continuous first-order theories. In Chapter 4, we develop the

machinery of strongly minimal sets in continuous logic and prove our main result, a partial Baldwin-Lachlan characterization of inseparably categorical theories that do contain strongly minimal sets. In Chapter 5, we present results relevant to specific classes of structures, including refinements of the Baldwin-Lachlan characterizations in the special cases of ultrametric and locally compact theories, characterizations of strongly minimal theories that do not interpret infinite discrete structures and continuous strongly minimal groups, and some general results regarding inseparably categorical expansions of Banach spaces. In the section on Banach spaces, Section 5.4, we introduce the concept of an *indiscernible subspace* and use it to improve a result of Shelah and Usvyatsov in [SU19] and to show that every expansion of a Banach space admits an infinite indiscernible set. This concept ought to be generally applicable to the model theory of expansions of Banach spaces. In Chapter 6, we present a formalism for approximate isomorphism of metric structures simultaneously generalizing those of [BY08b] and [BDNT17], which are largely incompatible; we extend some results to this broader context, specifically the approximate Ryll-Nardzewski theorem from [BY08b] and the existence of Scott sentences characterizing approximate isomorphism from [BDNT17]; and we prove one direction of an approximate version of Morley’s theorem for inseparably approximately categorical theories. Appendix C contains counterexamples relevant to this second theme as well.

Appendices A and B are references containing, respectively, topological facts that are frequently used but which may be unfamiliar to model theorists and proofs of generalizations of familiar model theoretic facts to the context of continuous logic, including two proofs of Morley’s theorem.

0.2 Comparison of Discrete Logic and Continuous Logic

In this informal and skippable section we will present an overview of the similarities and differences between discrete and continuous model theory. Our focus is on the topics relevant to this thesis as a whole.¹ While attempting to communicate intuition directly is often less than useless, it is inevitable that discrete model theorists are going to try to port their intuition to continuous logic, at least implicitly, and our hope is that this section will help with that. Despite appearances, continuous first-order logic is closely related to discrete first-order logic, arguably more closely related than $\mathcal{L}_{\omega_1\omega}$. Compactness and the Löwenheim-Skolem theorem still hold—with some small modifications—for instance. Many of these modifications are roughly predictable with enough experience. The most important, but also the vaguest, intuitions are presented in Table 1.

Table 1: General Translation

Discrete Logic	Continuous Logic
Equality (=)	Metric (d)
	Equality ($d(x, y) = 0$)
Cardinality	Density character (where compact \sim finite)
	ε -covering/packing number, ε -metric entropy
Negation, Disconnectedness	<i>(Potential) connectedness</i>
Finite	Finite at every $\varepsilon > 0$
	(Metrically) compact
Countable	Finite or countable/separable
Uncountable	Uncountable/inseparable

Each of these translations shows up in more than one way. Obviously structures

¹The most notable omissions are stability theory beyond ω -stability and neo-stability in general.

now have metrics, but there is also now a natural metric on type space, separate from the compact logical topology. The sizes of structures and type spaces are both now correctly measured with their metric density character, rather than their cardinality. Type spaces can fail to be disconnected, but so can structures themselves. In the context of approximate isomorphism, these translations even apply to the class of models of a theory of a given size: it has a metric, its density character is meaningful, and it can be connected. The blurring of the line between finite and countably infinite is felt constantly. It is necessary to admit ω -product sorts into the definition of imaginary, and, relatedly, finite tuples of parameters can behave like ω -tuples of parameters. Because of this, pronounced changes tend to happen to the countable case:² It is necessary to introduce the notions of approximate ω -saturation and approximate ω -homogeneity, but no analogous concepts are needed for larger cardinalities.

Key

Italics	<i>Represents failure of analogous fact</i>
Grayed	Represents object or property that corresponds closely but is no longer guaranteed to occur in analogous situations

There is, of course, some subjectivity with what constitutes the ‘correct’ analogous fact. For example, Tables 6 and 5 are written with the tacit assumption that approximate ω -saturation is the ‘correct’ continuous analog of ω -saturation, and as such the fact that an ω -stable theory may fail to have an ω -saturated separable model isn’t regarded as an italicized failure.

²Contradicting this trend, however, the Ryll-Nardzewski theorem is left remarkably unscathed, but the characterization of inseparably categorical theories is one of the main topics of this thesis.

Table 2: Types and Type Spaces

Discrete Logic	Continuous Logic
Formula	Real formula
	Clopen formula
	Open formula
Ctbl. partial type (ctbly. type-dfbl.)	Closed formula (= zeroset)
Uncountable partial type	Uncountable partial type
Type is a set of formulas	Type is a set of closed formulas
Complete type is a maximal consistent set of formulas	Complete type is a maximal consistent set of closed formulas
n -type, n finite	(Uniformly) d -finite type [BYU07]
ω -type	α -type, $\alpha \leq \omega$
α -type, $\alpha > \omega$	α -type, $\alpha > \omega$
Topology on type space	(Logic) topology on type space
Equality of types (discrete topology on type space)	d -metric between types (metric topology on type space)
Cardinality of type space	Metric density character of type space
Definable set	Definable set (zeroset that admits relative quantification)
	$\{0, 1\}$ -definable set (type definable and co-type-definable)
Isolated/principal/atomic type	d -atomic type
	Topologically isolated type
Type determined by definable sets	<i>Non-trivial type space can fail to contain any non-trivial definable sets</i>

Table 3: Imaginaries

Discrete Logic	Continuous Logic
Finitary product sort	Finitary or countable product sort
Definable equivalence relation	Definable $\{0, 1\}$ -valued eq. relation
	Definable pseudo-metric (= countably type-definable equivalence relation)
Countably type-definable eq. relation	
Passing to definable set not taken as primitive imaginary-forming operation	<i>Passing to definable set not 'derivable' from passing to quotients</i>

Table 4: acl/dcl and Definable (Partial) Functions

Discrete Logic	Continuous Logic
Type is algebraic iff has finitely many realizations in any model	Type is algebraic iff has compactly many realizations in any model
Realizations of algebraic partial type are definable set	Realizations of algebraic partial type are definable set
acl is same as bounded closure	acl is same as bounded closure
acl/dcl are abstract closure operators (increasing, idempotent)	acl/dcl are abstract closure operators (increasing, idempotent)
acl/dcl have finite character	acl/dcl have approximate finite character at every $\varepsilon > 0$
	acl/dcl have countable character
$a \in \text{dcl}(B)$ iff global type of a is isolated by $\text{tp}(a/B)$	$a \in \text{dcl}(B)$ iff global type of a is isolated by $\text{tp}(a/B)$
$a \in \text{dcl}(B)$ iff there is definable function f such that $a = f(\bar{b})$ for some $\bar{b} \in B$	$a \in \text{dcl}(B)$ iff there is definable partial fn. f such that $a = f(\bar{b})$ for some $\bar{b} \in B$
The arity of a definable function is finite	The arity of a definable (partial) function is finite or countable
Every formula $\varphi(\bar{x}, y)$ defines a partial fn. on some (possibly empty) dfbl. set	Every real fmla. $\varphi(\bar{x}, y)$ dfns. a partial fn. on some (possibly empty) zeroset
Every definable partial function on a definable set extends to a definable total function	<i>Not every dfbl. partial fn. on a zeroset extends to a dfbl. partial fn. on a dfbl. set, and not every dfbl. partial fn. on a dfbl. set extends to a dfbl. total fn.</i>
A function is definable iff there is a formula that defines its graph	A function is dfbl. iff there is a closed formula that dfns. its graph and dfns. the graph of a function in every model

Table 5: ω -stability

Discrete Logic	Continuous Logic
ω -stable iff every type space over countably many parameters is countable	ω -stable iff every type space over countably many parameters is metrically separable
ω -stable iff every type has ordinal Morley rank	ω -stable iff every type has ordinal ε -Morley rank for every $\varepsilon > 0$
ω -stable theory has saturated model of every infinite cardinality	ω -stable theory has approximately ω -saturated separable model
	ω -stable theory has saturated model of every uncountable cardinality

Table 6: Models

Discrete Logic	Continuous Logic
Cardinality of structure	Density character of structure
Countable prime model iff isolated types dense	Separable prime model iff d -atomic types dense
Countable prime model is unique	Separable prime model is unique
ω -saturated	Approximately ω -saturated ω -saturated
Countable ω -saturated model is unique	Separable apprx. ω -sat. model is unique
κ -saturated, $\kappa > \omega$	κ -saturated, $\kappa > \omega$
Saturated model of card. κ is unique	Sat. model of den. char. κ is unique
A sequence of partial types in a countable theory can be omitted iff they are all topologically nowhere dense.	<i>Omitting partial types in models is very complicated</i> [FM18]
	A sequence of partial types in a countable theory can be omitted in a pre-model iff they are all nowhere dense.
	A seq. of complete types can be omitted in a model iff they are all non- d -atomic.
ω -categorical iff every type space finite	ω -categorical iff every type space metrically compact
ω -categorical theory is ω -categorical after adding finitely many constants	ω -cat. theory is ‘apprx. ω -categorical’ after adding finitely many constants
	<i>ω-cat. theory can fail to be ω-cat. after adding finitely many constants</i>
Unique countable model of ω -categorical theory is prime, ω -saturated, and ω -homogeneous	Unique separable model of ω -categorical theory is prime, apprx. ω -saturated, and apprx. ω -homogeneous
A countable complete theory cannot have precisely 2 countable models	<i>There is a countable complete theory with precisely 2 separable models</i>

0.3 Metric Space Preliminaries

(Pseudo-)metric spaces are totally ubiquitous in continuous logic. Each of the following objects has at least one natural (pseudo-)metric on it:

- A given structure \mathfrak{M} .³
- The class of models of a given theory T . (Proposition 6.2.2)
- The collection of \mathcal{L} -terms for a given signature \mathcal{L} . (Definition 1.4.1)
- The collection of real \mathcal{L} -formulas for a given signature \mathcal{L} . (Definition 1.4.1)
- The set of n -types over a given theory T . (Definition 2.0.1)
- The collection of definable sets in a given theory T . (Definition 2.3.27)
- The collection of definable functions in a given theory T . (Definition 2.3.88)
- The collection of imaginary sorts over a given theory T . (Definition 3.4.10)

As such, it seems prudent to establish basic facts regarding and notation for metric spaces at the beginning of this thesis, rather than in an appendix. Everything in this section applies equally well to pseudo-metric spaces, but to avoid writing ‘pseudo-’ over and over again, we have written everything in terms of metric spaces.

Definition 0.3.1. For any metric space (X, d) , we define the following.

- For any $x \in X$, the *open (resp. closed) ball of radius ε with center x* , written $B_{<\varepsilon}^d(x)$ (resp. $B_{\leq\varepsilon}^d(x)$), is the set $\{y \in X : d(x, y) < \varepsilon\}$ (resp. $\{y \in X : d(x, y) \leq \varepsilon\}$).
- The *diameter of X* , written $\text{diam } X$, is $\sup\{d(x, y) : x, y \in X\}$.
- For any set $A \subseteq X$, the *point-set distance from x to A* , written $d_{\text{inf}}(x, A)$ is $\inf\{d(x, a) : a \in A\}$, where $\inf \emptyset = \text{diam } X$.

³Even in continuous logic without a metric. (Remark 1.2.3)

- The *open (resp. closed) ε -fattening* of A , written $A^{d < \varepsilon}$ (resp. $A^{d \leq \varepsilon}$), is the set $\{y \in X : d_{\text{inf}}(x, A) < \varepsilon\}$ (resp. $\{y \in X : d_{\text{inf}}(x, A) \leq \varepsilon\}$). \triangleleft

We will drop the superscript d when it is clear from context. Note that $A^{< \varepsilon} = \bigcup_{y \in A} B_{< \varepsilon}(y)$, but an analogous statement is not true for $A^{\leq \varepsilon}$ except in special circumstances. Fortunately these special circumstances occur frequently in the contexts that we care about.

Definition 0.3.2. For two metric spaces X and Y and any $r \geq \frac{1}{2} \min\{\text{diam } X, \text{diam } Y\}$, the *disjoint union of X and Y with separation r* , written $X \sqcup_r Y$, is a metric space with underlying set $X \sqcup Y$ and a metric ρ satisfying $\rho \upharpoonright X = d^X$, $\rho \upharpoonright Y = d^Y$, and for any $x \in X$ and $y \in Y$, $\rho(x, y) = r$. \triangleleft

Definition 0.3.3. For any metric space (X, d) , the *density character of X* , written $\#^{\text{dc}} X$ or $\#^{\text{dc}}(X, d)$, is the smallest cardinality of a dense subset of X . X is called *separable* if $\#^{\text{dc}} X \leq \aleph_0$ and *inseparable* if $\#^{\text{dc}} X \geq \aleph_1$. \triangleleft

Definition 0.3.4 (Counting Notions). For any metric space (X, d) , $Y \subseteq X$, $\sqsupset \in \{\geq, >\}$, $\sqsubset \in \{<, \leq\}$, and $\varepsilon > 0$, we define the following.

- A set $Z \subseteq X$ is $(\sqsupset \varepsilon)$ -*separated* if for any distinct $z_0, z_1 \in Z$, $d(z_0, z_1) \sqsupset \varepsilon$.
- The $(\sqsupset \varepsilon)$ -*metric entropy of Y* , written $\#_{\sqsupset \varepsilon}^{\text{ent}} Y$, is $\sup\{|Z| : Z \subseteq Y, Z \text{ } (\sqsupset \varepsilon)\text{-separated}\}$.
- A set $Z \subseteq X$ $(\sqsubset \varepsilon)$ -*covers* Y if for any $y \in Y$ there exists a $z \in Z$ such that $d(y, z) \sqsubset \varepsilon$.⁴

⁴Again, note that $(\leq \varepsilon)$ -covering is not quite the same thing as $Y \subseteq Z^{\leq \varepsilon}$. If one was so inclined, one could define a notion of *weak $(\leq \varepsilon)$ -covering number* in terms of this weaker notion of covering, but it doesn't seem to be very natural in that it doesn't fit into a statement like Proposition 0.3.5 neatly.

- The *external* $(\sqsubset \varepsilon)$ -covering number of Y in X , written $\#_{\sqsubset \varepsilon}^{\text{ecov}, X} Y$, is $\inf\{|Z| : Z \subseteq X, Z \text{ } (\sqsubset \varepsilon)\text{-covers } Y\}$.
- The *internal* $(\sqsubset \varepsilon)$ -covering number of Y , written $\#_{\sqsubset \varepsilon}^{\text{icov}} Y$, is $\#_{\sqsubset \varepsilon}^{\text{ecov}, Y} Y$.
- The *external* $(\sqsubset \varepsilon)$ -packing number of Y in X , written $\#_{\sqsubset \varepsilon}^{\text{epac}, X} Y$, is $\sup\{|Z| : Z \subseteq Y, (\forall z_0, z_1 \in Z) z_0 \neq z_1 \Rightarrow B_{\sqsubset \varepsilon}(z_0) \cap B_{\sqsubset \varepsilon}(z_1) = \emptyset\}$.
- The *internal* $(\sqsubset \varepsilon)$ -packing number of Y , written $\#_{\sqsubset \varepsilon}^{\text{ipac}} Y$, is $\#_{\sqsubset \varepsilon}^{\text{epac}, Y} Y$.
- The $(\sqsubset \varepsilon)$ -partition number of Y , written $\#_{\sqsubset \varepsilon}^{\text{par}} Y$, is $\inf\{|P| : P \text{ a partition of } Y, (\forall Z \in P) \text{diam}(Z) \sqsubset \varepsilon\}$.⁵

We may drop the superscript X if it is clear from context. If we need to refer to an unspecified or partially specified counting notion, we may use notation such as $\#_{\sqsubset \varepsilon}^*$, $\#_{\sqsubset \varepsilon}^*$, or $\#_{\sqsubset \varepsilon}^*$, with $*$ the symbolic name of some counting notion and \sqsubset , \sqsubset , or \sqsubset representing some appropriate element of $\{<, \leq, >, \geq\}$. \triangleleft

Proposition 0.3.5. *Let (X, d) be a metric space, and let $Y \subseteq X$ be some subsets. For any $\varepsilon > 0$,*

$$\begin{array}{ccccccc}
 & & & \#_{\leq \varepsilon}^{\text{ipac}} Y & & & \\
 & & & \sqsubset & & \sqsupset & \\
 \#_{> 2\varepsilon}^{\text{ent}} Y & \leq & \#_{\leq \varepsilon}^{\text{epac}, X} Y & & \#_{\leq \varepsilon}^{\text{icov}} Y & \leq & \#_{> \varepsilon}^{\text{ent}} Y & \leq & \#_{\leq \varepsilon}^{\text{par}} Y. \\
 & & \sqsupset & & \sqsupset & & \sqsupset & & \\
 & & \#_{\leq 2\varepsilon}^{\text{par}} Y & \leq & \#_{\leq \varepsilon}^{\text{ecov}, X} Y & & & &
 \end{array}$$

For any $\varepsilon > 0$ and any $\delta > 0$ or $\delta = 0$, if Y is compact,

⁵The $(< \varepsilon)$ -partition number also doesn't seem to be very natural outside of the context of compact metric spaces. Specifically, note that while $\text{diam } B_{\leq \varepsilon}(x) \leq 2\varepsilon$ is always true, $\text{diam } B_{< \varepsilon}(x) < 2\varepsilon$ is not. Perhaps the 'correct' notion is a partition into sets Z satisfying $d(z_0, z_1) < 2\varepsilon$ for all $z_0, z_1 \in Z$.

$$\begin{array}{ccccccc}
& & & \#_{<\varepsilon}^{\text{ipac}} Y & & & \\
& & & \swarrow & & \searrow & \\
\#_{\geq 2\varepsilon}^{\text{ent}} Y & \leq & \#_{<\varepsilon}^{\text{epac}, X} Y & & \#_{<\varepsilon}^{\text{icov}} Y & \leq & \#_{\geq \varepsilon}^{\text{ent}} Y \leq \#_{<\varepsilon}^{\text{par}} Y. \\
& & \searrow & & \swarrow & & \\
& & \#_{< 2\varepsilon + \delta}^{\text{par}} Y & \leq & \#_{<\varepsilon}^{\text{ecov}, X} Y & &
\end{array}$$

If $0 < \varepsilon < \delta$ then

- $\#_{\geq \delta}^{\text{ent}} Y \leq \#_{>\varepsilon}^{\text{ent}} Y \leq \#_{\geq \varepsilon}^{\text{ent}} Y$,
- $\#_{<\delta}^* Y \leq \#_{\leq \varepsilon}^* Y \leq \#_{<\varepsilon}^* Y$, for $* \in \{\text{icov}, \text{ipac}, \text{par}\}$, and
- $\#_{<\delta}^{*, X} Y \leq \#_{\leq \varepsilon}^{*, X} Y \leq \#_{<\varepsilon}^{*, X} Y$, for $* \in \{\text{ecov}, \text{epac}\}$.

If $Z \subseteq Y \subseteq X \subseteq W$, then

- $\#_{\sqsupset \varepsilon}^{\text{ent}} Z \leq \#_{\sqsupset \varepsilon}^{\text{ent}} Y$,
- $\#_{\sqsupset \varepsilon}^{\text{par}} Z \leq \#_{\sqsupset \varepsilon}^{\text{par}} Y$, and
- $\#_{\sqsupset \varepsilon}^{*, W} Z \leq \#_{\sqsupset \varepsilon}^{*, X} Y$ for $* \in \{\text{ecov}, \text{epac}\}$.

Proof. Let F be a set of minimal cardinality such that $Y \subseteq F^{\leq \varepsilon}$. For each $f \in F$, let Y_f be a subset of $Y \cap B_{\leq \varepsilon}(f)$ chosen so that $\{Y_f\}_{f \in F}$ is a partition of Y . By construction, $\text{diam } Y_f \leq 2\varepsilon$ and $\bigcup_{f \in F} Y_f = Y$, so we have $\#_{\leq 2\varepsilon}^{\text{par}} Y \leq |F| = \#_{\leq \varepsilon}^{\text{ecov}, X} Y$. A similar argument shows that $\#_{\leq 2\varepsilon + \delta}^{\text{par}} Y \leq |F| = \#_{\leq \varepsilon}^{\text{ecov}, X} Y$ for every $\delta > 0$.

If Y is compact and F is a set of minimal cardinality such that $Y \subseteq F^{<\varepsilon}$, then there is some $\gamma < \varepsilon$ such that $Y \subseteq F^{<\gamma}$, and by the same argument we get $\#_{\leq 2\varepsilon}^{\text{par}} Y \leq |F| = \#_{\leq \varepsilon}^{\text{ecov}, X} Y$.

We will write the rest of the proofs for counting notions corresponding to $\leq \varepsilon$ or $> \varepsilon$. The proofs for $< \varepsilon$ and $\geq \varepsilon$ are identical.

Fix $\varepsilon > 0$. Let A be a maximal ($> \varepsilon$)-separated set. By definition this implies that $|A| \leq \#_{>\varepsilon}^{\text{ent}} Y$. Every $y \in Y$ must satisfy $d_{\text{inf}}(y, A) \leq \varepsilon$, otherwise A wouldn't be maximal, so we have $\#_{\leq\varepsilon}^{\text{icov}} Y \leq |A| \leq \#_{>\varepsilon}^{\text{ent}} Y$.

$\#_{\leq\varepsilon}^{\text{ecov},X} Y \leq \#_{\leq\varepsilon}^{\text{icov}} Y$ is obvious (any witness to an upper bound on $\#_{\leq\varepsilon}^{\text{icov}} Y$ is also a witness to an upper bound on $\#_{\leq\varepsilon}^{\text{ecov},X} Y$).

Let B be a subset of Y of minimal cardinality such that $Y \subseteq B^{\leq\varepsilon}$. Let C be any subset of Y such that for any distinct $c_0, c_1 \in C$, $B_{\leq\varepsilon}(c_0) \cap B_{\leq\varepsilon}(c_1) \cap Y = \emptyset$. For every b in B there must be at most one $c \in C$ such that $d(b, c) \leq \varepsilon$, otherwise this would contradict the packing condition (since $B \subseteq Y$). Therefore $|C| \leq |B| = \#_{\leq\varepsilon}^{\text{icov}}$. Since we can do this for any such C we have that $\#_{\leq\varepsilon}^{\text{ipac}} Y \leq \#_{\leq\varepsilon}^{\text{icov}} Y$. The proof that $\#_{\leq\varepsilon}^{\text{epac},X} Y \leq \#_{\leq\varepsilon}^{\text{ecov},X} Y$ is essentially the same.

$\#_{\leq\varepsilon}^{\text{epac},X} Y \leq \#_{\leq\varepsilon}^{\text{ipac}} Y$ is obvious (any witness to a lower bound on $\#_{\leq\varepsilon}^{\text{epac},X} Y$ is also a witness to a lower bound on $\#_{\leq\varepsilon}^{\text{ipac}} Y$).

Let E be a ($> 2\varepsilon$)-separated subset of Y . By construction, for any distinct $e_0, e_1 \in E$, $B_{\leq\varepsilon}(e_0) \cap B_{\leq\varepsilon}(e_1) = \emptyset$, so we have that $|E| \leq \#_{\leq\varepsilon}^{\text{epac},X} Y$. Since we can do this for any such E we have that $\#_{>2\varepsilon}^{\text{ent}} Y \leq \#_{\leq\varepsilon}^{\text{epac},X} Y$.

Let P be a partition of Y of minimum cardinality such that for every $Z \in P$, $\text{diam } Z \leq \varepsilon$. For a ($> \varepsilon$)-separated set A , for every $Z \in P$, $|A \cap Z| \leq 1$, so we have $\#_{>\varepsilon}^{\text{ent}} Y \leq \#_{\leq\varepsilon}^{\text{par}} Y$.

The inequalities after the two diagrams are all straightforward. \square

Note that we do not have any results of the form $Z \subseteq Y \Rightarrow \#_{<\varepsilon}^{\text{icov}} Z \leq \#_{<\varepsilon}^{\text{icov}} Y$. Neither of the internal counting notions are monotonic in the set (consider $\varepsilon = \frac{2}{3}$ and the sets $\{0, 1\} \subseteq \{0, \frac{1}{2}, 1\}$ with the standard metric from \mathbb{R}).

One of the advantages of having many different counting notions is that it gives us many different ways to prove that a subset of a metric space is compact.

Corollary 0.3.6. *For any counting notion $\#_{\square\varepsilon}^*$, a set $Y \subseteq X$ is pre-compact (i.e. has compact metric completion) if and only if $\#_{\square\varepsilon}^* Y$ is finite for every $\varepsilon > 0$.*

Proof. By Proposition 0.3.5, we only need to show this for one particular counting notion. The condition that $\#_{<2\varepsilon}^{\text{icov}} \bar{Y} \leq \#_{<\varepsilon}^{\text{icov}} Y$ is finite for every $\varepsilon > 0$ is equivalent to total boundedness, which is equivalent to pre-compactness. \square

Corollary 0.3.7. *For any metric space X and any counting notion $\#_{\square\varepsilon}^*$, $\lim_{\varepsilon \rightarrow 0} \#_{\square\varepsilon}^* X = \#^{\text{dc}} X$. In particular, if $\text{cf}(\#^{\text{dc}} X) > \aleph_0$, then for any counting notion $\#_{\square\varepsilon}^*$ for any sufficiently small $\varepsilon > 0$, $\#_{\square\varepsilon}^* X = \#^{\text{dc}} X$.*

Proof. By Proposition 0.3.5, we only need to show this for one particular counting notion. If $\#^{\text{dc}} X$ is finite, then this is obvious, so assume that $\#^{\text{dc}} X$ is infinite. It is obvious that $\#_{<\varepsilon}^{\text{icov}} X \leq \#^{\text{dc}} X$ for every $\varepsilon > 0$. Conversely, if for each $k < \omega$ we let $A_k \subseteq X$ have minimal cardinality such that $A_k^{<1/k} = X$, then $\bigcup_{k < \omega} A_k$ is a dense subset of X , and we get the required limit. \square

It is easy to construct examples where $(\forall \varepsilon > 0) \#_{\square\varepsilon}^* X < \#^{\text{dc}} X$ when $\text{cf}(\#^{\text{dc}} X) = \aleph_0$ (see Counterexample C.0.1)).

The following is not completely obvious, as $(\geq \varepsilon)$ -separated subsets of maximal cardinality do not exist in general.

Proposition 0.3.8. *For any metric space X , family $\{Y_i\}_{i \in I}$ of subsets of X , $\sqsupset \in \{\geq, >\}$, and $\varepsilon > 0$, $\#_{\sqsupset\varepsilon}^{\text{ent}} \bigcup_{i \in I} Y_i \leq \sum_{i \in I} \#_{\sqsupset\varepsilon}^{\text{ent}} Y_i$.*

Proof. Let A be a $(\sqsupset \varepsilon)$ -separated subset of $\bigcup_{i \in I} Y_i$. By partitioning A into $A_i \subseteq Y_i$ we see that $|A| = \sum_{i \in I} |A_i| \leq \sum_{i \in I} \#_{\sqsupset \varepsilon}^{\text{ent}} Y_i$. Since this holds for any $(\sqsupset \varepsilon)$ -separated subset of $\bigcup_{i \in I} Y_i$ we have $\#_{\sqsupset \varepsilon}^{\text{ent}} \bigcup_{i \in I} Y_i \leq \sum_{i \in I} \#_{\sqsupset \varepsilon}^{\text{ent}} Y_i$. \square

Definition 0.3.9. A modulus, ω , is a continuous, non-decreasing function $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with $\omega(0) = 0$.

Given metric spaces (X, d^X) and (Y, d^Y) and a modulus ω , a function $f : X \rightarrow Y$ is *uniformly continuous with regards to modulus ω* or *ω -uniformly continuous* if for all $x_0, x_1 \in X$,

$$d^Y(f(x_0), f(x_1)) \leq \omega(d^X(x_0, x_1)).$$

A function is *uniformly continuous* if it is ω -uniformly continuous for some modulus ω .

\triangleleft

Note that this definition is actually slightly stronger than the standard definition of uniformly continuous, although when Y is bounded, they are the same (see Corollary A.1.9).

Chapter 1

Syntax and Semantics

1.1 Signatures

Definition 1.1.1. A (*single-sorted*) *metric signature* \mathcal{L} (also called a language or a vocabulary) consists of a collection of predicate symbols \mathcal{P} and a collection of function symbols \mathcal{F} . When it is necessary to specify the signature in question we will write expressions such as $\mathcal{P}(\mathcal{L})$.

- For each predicate or function symbol $s \in \mathcal{P} \cup \mathcal{F}$, there is an *arity*, $a(s) \leq \omega$.
- For each predicate or function symbol $s \in \mathcal{P} \cup \mathcal{F}$ with non-zero arity, there is a modulus ω_s , called the *modulus of uniform continuity of s* .
- For each predicate symbol $P \in \mathcal{P}$, there is a *codomain interval*, a compact interval in \mathbb{R} , $I(P)$.

In addition there is a positive real number, the *diameter bound (of \mathcal{L})*, written $\text{db}(\mathcal{L})$.

There is a special predicate symbol d with $a(d) = 2$, $\omega_d(x) = 2x$, and $I(d) = [0, \text{db}(\mathcal{L})]$.

Functions with zero arity are referred to as *constants*. The collection of constants is denoted \mathcal{C} .

The *cardinality* of a metric signature, written $|\mathcal{L}|$, is $|\mathcal{P}| + |\mathcal{F}|$.

If \mathcal{L}' is a metric signature that contains a subset of the symbols of \mathcal{L} , we say that \mathcal{L}' is a *reduct* of \mathcal{L} and conversely we say that \mathcal{L} is an *expansion* of \mathcal{L}' .

The *empty signature* (with diameter bound r) is the signature \mathcal{L} with $\text{db}(\mathcal{L}) = r$, $\mathcal{P} = \{d\}$, and $\mathcal{F} = \emptyset$. \triangleleft

In the use of all following defined terms containing ‘ \mathcal{L} -,’ we will omit it if it is unimportant or clear from context.

Remark 1.1.2. Given that formulas with finitely many free variables are logically complete and that every theory is interdefinable with one in which every predicate and function has finite arity, one might be tempted to restrict our formalism to these cases. The problem with this is that definable sets may depend on countably many parameters non-trivially, and in this case we do not have the luxury to restrict attention to those with only finitely many parameters. Likewise, dcl is witnessed by definable partial functions which may in general have arity ω in a non-trivial way, which supports the contention that function symbols of arity ω are natural.

1.2 (Pre-)structures

Part of Definition 1.2.2 depends on Definition 1.2.1, which depends on a different part of Definition 1.2.2, but—be assured—there is no circularity.

Definition 1.2.1. For $\bar{a}, \bar{b} \in \mathfrak{M}^n$, if n is finite, we will write $d^{\mathfrak{M}}(\bar{a}, \bar{b})$ to mean $\max_{i < n} d^{\mathfrak{M}}(a_i, b_i)$, and if $n = \omega$ we will write $d^{\mathfrak{M}}(\bar{a}, \bar{b})$ to mean $\sup_{i < \omega} 2^{-i} d^{\mathfrak{M}}(a_i, b_i)$.

We will also write $d(\bar{x}, \bar{y})$ for the corresponding formulas, i.e. $\max\{dx_0y_0, \dots, dx_{n-1}y_{n-1}\}$ and $\sup_{i < \omega} 2^{-i} da_i b_i$.

If we need to write out the elements of the tuple explicitly, we will write either $d(x_0x_1 \dots x_{n-1}, y_0y_1 \dots y_{n-1})$ or $d(x_0, x_1, \dots, x_{n-1}; y_0, y_1, \dots, y_{n-1})$. \triangleleft

Uniform continuity of functions on product spaces will always mean uniform continuity with regards to one of these product metrics, unless otherwise stated.

Note that we have defined the metrics on tuples so that they always satisfy $d^{\mathfrak{M}}(\bar{a}, \bar{b}) \leq d^{\mathfrak{M}}(a_i, b_i)$, for every $i < |\bar{a}|$. Also note that while the max metric on finite products of metric spaces is in some precise sense a canonical choice,¹ the specific metric we have for ω -tuples is an arbitrary choice. What is important is that it metrizes the product topology and, in particular, that formulas on ω -many variables are uniformly continuous with regards to this metric. There is no canonical choice accomplishing these things, but any such choice will be uniformly equivalent to the one we have here, and typically the precise choice of product metric will be unimportant.

Definition 1.2.2. Given a metric signature \mathcal{L} , an \mathcal{L} -pre-structure \mathfrak{M} is a (possibly empty) pseudo-metric space $(M, d^{\mathfrak{M}})$ of diameter $\leq \text{db}(\mathcal{L})$, together with the following data:

- For each predicate symbol $P \in \mathcal{P} \setminus \{d\}$, $P^{\mathfrak{M}} : M^{a(P)} \rightarrow I(P)$ is a function.
- For each function symbol $f \in \mathcal{F}$, $f^{\mathfrak{M}} : M^{a(f)} \rightarrow M$ is a function.

If the arity of $s \in \mathcal{P} \cup \mathcal{F}$ is non-zero, then $s^{\mathfrak{M}}$ is required to be ω_s -uniformly continuous.

¹The max metric corresponds to the product in the category of metric spaces with 1-Lipschitz maps.

A *reduced \mathcal{L} -pre-structure* is an \mathcal{L} -pre-structure \mathfrak{M} such that $(M, d^{\mathfrak{M}})$ is a metric space (i.e. for any x, y if $d(x, y) = 0$, then $x = y$), and an *\mathcal{L} -structure* is an \mathcal{L} -pre-structure such that $(M, d^{\mathfrak{M}})$ is a complete metric space.

By an abuse of notation we will typically write the underlying set of a (pre-)structure as \mathfrak{M} , to emphasize that it is the underlying set of some (pre-)structure in question.

If \mathfrak{M} and \mathfrak{N} are (pre-)structures and $M \subseteq N$, if $P^{\mathfrak{M}}(\bar{a}) = P^{\mathfrak{N}}(\bar{a})$ for every predicate symbol $P \in \mathcal{P}$ and tuple $\bar{a} \in \mathfrak{M}^{a(P)}$, and if $f^{\mathfrak{M}}(\bar{a}) = f^{\mathfrak{N}}(\bar{a})$ for every function symbol $f \in \mathcal{F}$ and tuple $\bar{a} \in \mathfrak{M}^{a(f)}$, then we say that \mathfrak{M} is a *sub- $(\mathcal{L}-)$ (pre-)structure of \mathfrak{N}* , written $\mathfrak{M} \subseteq \mathfrak{N}$. Note that if $A \subseteq N$ is a set that is closed under all functions $f^{\mathfrak{N}}$, then there is a canonical way of regarding A as the universe of a sub- \mathcal{L} -pre-structure of \mathfrak{N} . We refer to this as the *induced structure* on A .

Given a pre-structure, \mathfrak{M} , the *reduction of \mathfrak{M}* , written ${}^r\mathfrak{M}$, is the pre-structure resulting from modding $(M, d^{\mathfrak{M}})$ by the equivalence relation $x \sim y \Leftrightarrow d^{\mathfrak{M}}(x, y) = 0$. By the uniform continuity requirements, this is still an \mathcal{L} -pre-structure. Moreover it is clearly a reduced \mathcal{L} -pre-structure. Given a reduced pre-structure, \mathfrak{N} , the *completion of \mathfrak{N}* , written $\overline{\mathfrak{N}}$, is the \mathcal{L} -structure such that \overline{N} is the metric completion of N . All predicate and function symbols are extended by uniform continuity. We will regard \mathfrak{N} as a sub-structure of $\overline{\mathfrak{N}}$ in the canonical way. If \mathfrak{M} is a pre-structure, we will write $\overline{\mathfrak{M}}$ for $\overline{{}^r\mathfrak{M}}$, and we may refer to this as the completion of \mathfrak{M} .

If \mathcal{L}' is a reduct of \mathcal{L} and \mathfrak{M} is an \mathcal{L} -(pre-)structure, then the *\mathcal{L}' -reduct of \mathfrak{M}* , written $\mathfrak{M} \upharpoonright \mathcal{L}'$ is the \mathcal{L}' structure with the same underlying set such that the interpretations of the remaining predicate, function, and constant symbols are the same. \triangleleft

Remark 1.2.3. It is possible to formalize continuous logic without a metric, analogously to first-order logic without equality [GK19]. From this point of view the requirements that d be interpreted as a pseudo-metric are the analog of the axioms asserting that $=$ is an equivalence relation and the requirements that the interpretations of symbols be uniformly continuous are equivalent to the predicate substitution axioms (e.g. $\forall xy[x = y \rightarrow (P(x) \leftrightarrow P(y))]$), which are sometimes called the *equality axioms*.

This approach has the advantage that signatures do not need to specify moduli of uniform continuity and that the definition of ultraproducts becomes simpler. In fact, in the proof of Łoś's theorem (Proposition 1.6.2), we need to implicitly use the notion of continuous logic without a metric.

The difference, however, between continuous logic and continuous logic without a metric is smaller than the difference between discrete logic and discrete logic without equality. In a countable signature (or more generally a signature with countably many non-constant symbols) without a metric and in which every predicate and function symbol has finite arity,² there is a (uniformly) definable pseudo-metric with regards to which all predicates are automatically uniformly continuous. Namely, let $\{\varphi_i(x, \bar{z}_i)\}_{i < \omega}$ be an enumeration of all restricted atomic $\mathcal{L}(x\bar{z})$ -formulas containing no constant symbols, with \bar{z} an ω -tuple of variables, and let $r_i = \max\{\sup I(\varphi_i) - \inf I(\varphi_i), 1\}$ and

$$d(x, y) = \sup_{i < \omega} \frac{1}{2^i r_i} \sup_{\bar{z}_i} |\varphi_i(x, \bar{z}_i) - \varphi_i(y, \bar{z}_i)|.$$

With regards to this metric all predicate and function symbols are automatically

uniformly continuous (Lipschitz continuous, even). So for such signatures the two approaches are in some sense equivalent,³ although the natural notion of ‘structure’ is different in that requiring structures to be complete with regards to that metric is less well motivated than it is in continuous logic with a metric.

In this way it is also possible to regard discrete logic without equality as a special case of continuous logic (although, again, the intended semantics are typically different, discrete model theory without equality typically deals with what we are calling pre-structures or sometimes reduced pre-structures [KM99]).

Definition 1.2.4. If \mathcal{L} is a signature and C is a set of constant symbols (some of which may be in \mathcal{L}), then \mathcal{L}_C is the signature obtained by adding C to \mathcal{L} as constant symbols.

If \mathfrak{M} is an \mathcal{L} -pre-structure and A is a set of elements of \mathfrak{M} , then we define \mathcal{L}_A to be a signature extending \mathcal{L} with each element a of A as a new constant symbol, the *name of a* , written $n(a)$. We also define \mathfrak{M}_A , also written (\mathfrak{M}, A) , to be an \mathcal{L}_A -pre-structure expanding \mathfrak{M} such that for each $a \in A$, $n(a)^{\mathfrak{M}} = a$. By an abuse of notation we will typically write a for $n(a)$. ◁

²Predicates and functions with arity ω are slightly more difficult to formalize in continuous logic without a metric, but every signature is interdefinable (Definition B.3.7) with a signature in which every predicate and function symbol has finite arity.

³Even in uncountable signatures there is a subtler equivalence between the approaches given by converting non-metric signatures to many-sorted metric signatures with a sort for each countable reduct of the original signature with appropriate transition maps. See the discussion after Theorem 2.22 in [BY05].

1.3 Terms and Formulas

We allow arbitrary collections of variable symbols (in particular so that we can formalize κ -types for arbitrarily large κ), but we have a special class of variable symbols called the *variable symbols for binding*, written $\mathcal{V}_b = \{\dot{v}_i\}_{i < \omega_1}$, indexed by ω_1 .

Definition 1.3.1. Given a metric signature \mathcal{L} and a collection of variable symbols V (which may or may not be disjoint from \mathcal{V}_b), the collection of $\mathcal{L}(V)$ -terms is defined inductively. We also simultaneously define the collection of *restricted* $\mathcal{L}(V)$ -terms.

- If $v \in V$ is a variable symbol, then v is an $\mathcal{L}(V)$ -term and is restricted.
- If $f \in \mathcal{F}$ is a function symbol and \bar{t} is an $a(f)$ -tuple of $\mathcal{L}(V)$ -terms, then $f\bar{t}$ is an $\mathcal{L}(V)$ -term. If \bar{t} is a tuple of restricted terms and is either finite or eventually constant, then $f\bar{t}$ is restricted.

When we want to emphasize that we are thinking of t as an $\mathcal{L}(V)$ -term, we will write it as $t(V)$. V will frequently be a particular tuple of variables (e.g. $\mathcal{L}(\bar{x})$ -terms). An \mathcal{L} -term is an $\mathcal{L}(V)$ -term for some set of variable symbols V . \triangleleft

Definition 1.3.2. Given a metric signature \mathcal{L} and a set of variable symbols V , *real* $\mathcal{L}(V)$ -formulas are defined inductively. Also we define the *codomain interval* of φ , written $I(\varphi)$, (which is always a compact subinterval of \mathbb{R}) and the *free variables* of φ , written $\text{fv}(\varphi)$.

- If $P \in \mathcal{P}$ is a predicate symbol and \bar{t} is an $a(P)$ -tuple of $\mathcal{L}(V)$ -terms, then $P\bar{t}$ is a formula. $I(P\bar{t}) = I(P)$. $\text{fv}(P\bar{t})$ is the set of all variable symbols occurring in \bar{t} .
- For any $n \leq \omega$, if $F : \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$ is a partially defined function and $\bar{\varphi}$ is an n -tuple of $\mathcal{L}(V)$ -formulas such that F is defined and continuous (in the product topology)

on the set $X = \prod_{i < n} I(\varphi_i)$, then $G\bar{\varphi}$, with $G = F \upharpoonright X$,⁴ is an $\mathcal{L}(V)$ -formula. We will write this as $F\bar{\varphi}$. $I(F\bar{\varphi})$ is the image of X under F . (Note that this is always a compact interval, possibly of zero length.) $\text{fv}(F\bar{\varphi}) = \bigcup_{i < n} \text{fv}(\varphi_i)$.

- If W is a finite or countable set or tuple of variable symbols for binding and φ is an $\mathcal{L}(VW)$ -formula, then $\text{inf}_W \varphi$ and $\text{sup}_W \varphi$ are $\mathcal{L}(V)$ -formulas. $I(\text{inf}_W \varphi) = I(\text{sup}_W \varphi) = I(\varphi)$, and $\text{fv}(\text{inf}_W \varphi) = \text{fv}(\text{sup}_W \varphi) = \text{fv}(\varphi) \setminus W$.

We will typically write quantification over a singleton set of variables as $\text{inf}_v \varphi$ or $\text{sup}_v \varphi$, rather than $\text{inf}_{\{v\}} \varphi$ or $\text{sup}_{\{v\}} \varphi$.

When we need to write an expression that could be either inf or sup we will use the metasyntactic variable Q , possibly with a subscript. In such cases we will not put the variable in the subscript, so we will write expressions such as $Q_0 x Q_1 y \varphi$.⁵

If φ is a formula, \bar{x} is a tuple of distinct variable or constant symbols, and \bar{t} is a tuple of terms of the same length, then we write $\varphi[\frac{\bar{t}}{\bar{x}}]$ for φ with all free instances of each x_i replaced with t_i simultaneously.

If φ is a formula and \bar{x} is a tuple of distinct variable or constant symbols, then we will write $Q\bar{x}\varphi$ for $Q\bar{v}(\varphi[\frac{\bar{v}}{\bar{x}}])$, where $\bar{v} = v_i v_{i+1} \dots$ and i is the smallest index greater than all indices of variable symbols for binding appearing in φ .

When we want to emphasize that we are thinking of φ as an $\mathcal{L}(V)$ -formula, we will write it as $\varphi(V)$ (again, typically with a tuple of variables). An \mathcal{L} -formula is an $\mathcal{L}(V)$ -formula for some V . A *real sentence* is a real formula φ such that $\text{fv}(\varphi) = \emptyset$. A *quantifier-free* formula is a formula containing no instances of inf or sup. An *atomic*

⁴The only reason for this convention of restricting to a specific domain on which the function is continuous is that we would like to only have 2^{\aleph_0} connectives, rather than $2^{2^{\aleph_0}}$ connectives.

⁵We have not adopted it here, but a similarly compact and legible notation could be used for inf and sup as well, with an inverted I and S, respectively, yielding expressions such as $IxSy\varphi$.

formula is a formula of the form $P\bar{x}$. An *affine atomic* formula is a formula of the form $c + \sum_{i < N} r_i \cdot \varphi_i$, with c and r_i real numbers and φ_i atomic for every $i < N$. A *finitary* formula is a formula containing no infinitary connectives. An *unnested* atomic formula is one that is either of the form $P\bar{x}$ or $dyf\bar{x}$, with \bar{x} a tuple of variables. An *unnested* formula is one in which all atomic sub-formulas are unnested.

By an abuse of notation we will denote the collection of $\mathcal{L}(V)$ -formulas by $\mathcal{L}(V)$. \triangleleft

The permissiveness of the second bullet point in this definition is largely for the sake of three particular connectives, namely $\bar{\varphi} \mapsto \sum_{i < \omega} \varphi_i$, $\bar{\varphi} \mapsto \sup_{i < \omega} \varphi_i$, and $\bar{\varphi} \mapsto \inf_{i < \omega} \varphi_i$, which we would like to be able to use freely without equivocating. (In fact we have already used one of these in Definition 1.2.1.) Note, however, that $\bar{\varphi} \mapsto \sup_{i < \omega} \varphi_i$ and $\bar{\varphi} \mapsto \inf_{i < \omega} \varphi_i$ are not continuous on $[0, 1]^\omega$, even though they are well defined on that set. Allowing sup and inf on $[0, 1]^\omega$ would take us to a continuous analog of $\mathcal{L}_{\omega_1\omega}$. We will typically use sup as a connective on sets of the form $\prod_{i < \omega} [0, r_i]$ with r_i a sequence of positive numbers limiting to 0 (as in Definition 1.2.1), and inf will be similarly limited. Likewise, infinite sums will always be restricted to sets of the form $\prod_{i < \omega} [s_i, r_i]$ with $s_i \leq 0 \leq r_i$ and $\sum_{i < \omega} s_i$ and $\sum_{i < \omega} r_i$ convergent.

We have chosen to restrict bound variables to a particular special class of variables to compensate for allowing arbitrarily large collections of variables. We want to keep the collection of $\mathcal{L}(V)$ -formulas a set, rather than an unbounded proper class.

Quantification over ω -tuples of variables may seem hopelessly non-first-order, but consider the following: In discrete first-order logic, if φ is a formula, then $\exists x_0 \exists x_1 \exists x_2 \dots \varphi$ is logically equivalent to a first-order formula, as φ only has finitely many free variables. This is of course a fairly trivial observation, but an analogous thing is happening here. To any $\varepsilon > 0$, the value of a formula with an infinite string of quantifiers only depends

on finitely many of them. We won't need it here, but we could have even allowed infinite alternating quantifier strings and non-well-founded (but still linearly ordered) infinite quantifier strings.

Formulas involving functions with standard non-Polish notation will be freely written using the standard notation. So for instance, if φ and ψ are formulas then $\varphi + \psi$ is also a formula. We will also freely write expressions like $\varphi + \psi + \chi$, instead of $(\varphi + \psi) + \chi$.

Notation 1.3.3 (Common Connectives). The *monus function*, written $x \dot{-} y$, is defined as $\max\{x - y, 0\}$. This notation is common in continuous logic literature, but we won't need it very often here.

If we need an infix notation for max and min, we will use \uparrow and \downarrow , respectively. We will avoid using \vee and \wedge for these operations to avoid confusion with the logical operations,⁶ which we will also be using.

For real numbers r, s with $r \leq s$, the $[r, s]$ -clamping function,⁷ written $[x]_r^s$, is $\min\{\max\{x, r\}, s\}$. Note that $[x]_r^s = r$ if $x \leq r$, $[x]_r^s = x$ if $r \leq x \leq s$, and $[x]_r^s = s$ if $x \geq s$, and so in particular for any formula φ , $I([\varphi]_r^s) \subseteq [r, s]$. \triangleleft

A common use of the clamping function is to force sums to converge: If $\{r_i\}_{i < \omega}$ is a sequence of positive real numbers such that $\sum_{i < \omega} r_i < \infty$, then for any sequence of formulas $\{\varphi_i\}_{i < \omega}$, $\sum_{i < \omega} [\varphi_i]_{-r_i}^{r_i}$ is a formula. This is closely related to the forced limit function presented in [BYU10].

Definition 1.3.4. A *restricted atomic formula* is an atomic formula $P\bar{t}$ such that \bar{t} is a tuple of restricted terms that is either finite or eventually constant.

⁶And frankly, \vee and \wedge are already confusing as lattice theoretic notation.

⁷The term *clamp* or *clamping*, while not established in mathematics, is common in computer graphics.

Restricted formulas are the smallest class of formulas containing the restricted atomic formulas and the formula 1 (i.e. the zero-ary connective 1) and closed under quantification over 1-tuples of variables of the form \dot{v}_i for $i < \omega$ and the connectives $x + y$, $\max\{x, y\}$, $\min\{x, y\}$, and $x \mapsto r \cdot x$ for rational r . The collection of restricted $\mathcal{L}(V)$ -formulas is written $R\mathcal{L}(V)$. \triangleleft

We will, of course, typically write r instead of $r \cdot 1$ in restricted formulas and use other common notational conveniences such as $\max\{x, y, z\}$.

Remark 1.3.5. The notion of formula used in [BYBHU08] is equivalent to what we are calling a finitary formula. I have elected to formalize continuous logic in terms of infinitary connectives, which [BYBHU08] alludes to as a possibility (specifically around Proposition 9.3). Instead of using infinitary connectives they use the notion of a *definable predicate* for a fixed structure \mathfrak{M} . A definable predicate on \mathfrak{M} is a function that is the uniform limit of the evaluations of formulas. [BYBHU08] also has a slightly more restrictive definition of restricted formulas, which is a less significant difference.

Occasionally it will be convenient to allow formulas to take on values in topological spaces other than \mathbb{R} or to allow a formula to have some extra parameters in some topological space.

Definition 1.3.6. For any topological space X , signature \mathcal{L} , and set of variable symbols V , an X -valued $\mathcal{L}(V)$ -formula, φ , is an expression of the form $F\bar{\psi}$, with $\bar{\psi}$ a tuple of real $\mathcal{L}(V)$ -formulas and a partially defined function $F : \subseteq \mathbb{R}^{|\bar{\psi}|} \rightarrow X$ that is defined and continuous on the set $\prod_{i < |\bar{\psi}|} I(\psi_i)$. We set $\text{fv}(\varphi) = \bigcup_{i < |\bar{\psi}|} \text{fv}(\psi_i)$. If Y is some subspace

of X and we are in a context in which we know $F\bar{\psi}$ always evaluates to an element of Y (such as in the context of a specific theory), we may refer to $F\bar{\psi}$ as a Y -valued formula. \triangleleft

Obviously \mathbb{R} -valued formulas are themselves real formulas. Most commonly X will be \mathbb{C} or some compact metric space (such as the unit circle in \mathbb{C}). It is possible to give an account of (first-order) quantifiers for X -valued formulas, but we will not.

Definition 1.3.7. Given a pre-structure \mathfrak{M} and a (possibly empty) set of variable symbols V , a *variable assignment (for V on \mathfrak{M})* or a *V -assignment (on \mathfrak{M})*, ι , is a function $\iota : V \rightarrow \mathfrak{M}$.

If ι is a variable assignment, \bar{x} is a tuple of distinct variable symbols (possibly not in the domain of ι), and \bar{a} is a tuple of elements of \mathfrak{M} of the same length, then $\iota[\frac{\bar{a}}{\bar{x}}]$ is the variable assignment where $\iota[\frac{\bar{a}}{\bar{x}}](v) = \iota(v)$ if $v \notin \bar{x}$ and $\iota[\frac{\bar{a}}{\bar{x}}](v) = a_i$ if $v = x_i \in \bar{x}$. \triangleleft

Definition 1.3.8. Let \mathfrak{M} be an \mathcal{L} -pre-structure, let \bar{x} be a tuple of variable symbols, and let ι be a variable assignment for \bar{x} in \mathfrak{M} .

- Given an $\mathcal{L}(V)$ -term t , the *evaluation of t (on ι in \mathfrak{M})*, written $t^{\mathfrak{M}}(\iota)$, is an element of \mathfrak{M} defined inductively:

- If $t = v \in \text{dom}(\iota)$, then $t^{\mathfrak{M}}(\iota) = \iota(v)$.
- If $t = f\bar{s}$, then $t^{\mathfrak{M}}(\iota) = f^{\mathfrak{M}}(s_0^{\mathfrak{M}}(\iota), s_1^{\mathfrak{M}}(\iota), \dots)$.

- Given an $\mathcal{L}(V)$ -formula φ , the *evaluation of φ (on ι in \mathfrak{M})*, written $\varphi^{\mathfrak{M}}(\iota)$, is defined inductively:

- If $\varphi = P\bar{t}$, then $\varphi^{\mathfrak{M}}(\iota) = P^{\mathfrak{M}}(t_0^{\mathfrak{M}}(\iota), t_1^{\mathfrak{M}}(\iota), \dots)$.

- If $\varphi = F\bar{\psi}$, then $\varphi^{\mathfrak{M}}(\iota) = F(\psi_0^{\mathfrak{M}}(\iota), \psi_1^{\mathfrak{M}}(\iota), \dots)$.
- If $\varphi = \left\{ \begin{array}{c} \inf_W \\ \sup_W \end{array} \right\} \psi$, where W is some set of variable symbols, and \bar{v} is an enumeration of W , then

$$\varphi^{\mathfrak{M}}(\iota) = \left\{ \begin{array}{c} \inf \\ \sup \end{array} \right\} \left\{ \psi^{\mathfrak{M}} \left(\iota \left[\begin{array}{c} \bar{a} \\ \bar{v} \end{array} \right] \right) : \bar{a} \in M \right\},$$

if M is non-empty, otherwise $\varphi^{\mathfrak{M}}(\iota) = \left\{ \begin{array}{c} \max \\ \min \end{array} \right\} I(\psi)$.

We will typically drop \mathfrak{M} when it is clear from context, most commonly in expressions involving d . In general for an $\mathcal{L}(V)$ -term, t , we will consider $t^{\mathfrak{M}}$ as a function $M^V \rightarrow M$ given by $\iota \mapsto t^{\mathfrak{M}}(\iota)$. We will do likewise for formulas. \triangleleft

Definition 1.3.9. For real $\mathcal{L}(V)$ -formulas φ, ψ , expressions of the form $\varphi \square \psi$ with $\square \in \{\leq, \geq, =\}$ are called *closed $\mathcal{L}(V)$ -formulas* and with $\square \in \{<, >, \neq\}$ are called *open $\mathcal{L}(V)$ -formulas*. Open and closed $\mathcal{L}(V)$ -formulas are referred to collectively as *topological $\mathcal{L}(V)$ -formulas*. Topological \mathcal{L} -formulas are topological $\mathcal{L}(V)$ -formulas for any V . The set of all closed $\mathcal{L}(V)$ -formulas is written $C\mathcal{L}(V)$. The set of all open $\mathcal{L}(V)$ -formulas is written $O\mathcal{L}(V)$.

For a topological formula $Q = (\varphi \square \psi)$, the *negation of Q* , written $\neg Q$, is the topological formula $\varphi \square' \psi$ where \square' is the logical negation of \square (i.e. $\neg(\varphi \leq \psi) = (\varphi > \psi)$). Note that the negation of a closed formula is an open formula and vice versa.

The *free variables* of a topological formula are the free variables of its constituent real formulas. We will typically represent closed formulas with capital letters suggestive of closed sets, such as F and G , and we will represent open formulas with capital letters suggestive of open sets, such as U and V . When we wish to emphasize that these are

$\mathcal{L}(V)$ -formulas we will write expressions such as $F(V)$. *Quantifier-free, atomic, affine atomic, finitary, unnested, or restricted* topological formulas are topological formulas in which both real formulas are quantifier-free, atomic, affine atomic, finitary, unnested, or restricted, respectively. A *topological sentence* is a topological formula in which both real formulas are real sentences.

If φ is a real $\mathcal{L}(V)$ -formula, ψ is a real $\mathcal{L}(V')$ -formula (which in particular may be fixed real numbers), and $\square \in \{<, >, \leq, \geq, =, \neq\}$, then we write $\mathfrak{M} \models \varphi(\iota) \square \psi(\iota')$ to mean $\varphi^{\mathfrak{M}}(\iota) \square \psi^{\mathfrak{M}}(\iota')$. We may also write expressions such as $\mathfrak{M} \models \varphi \leq \psi < \chi$, with the obvious meaning.

If $X(\bar{x}, \bar{y})$ is a topological formula, \mathfrak{M} is a pre-structure, and $\bar{a} \in \mathfrak{M}$ is a tuple assigned to \bar{y} , then $X(\mathfrak{M}, \bar{a})$ is the set of all tuples \bar{b} such that $\mathfrak{M} \models X(\bar{b}, \bar{a})$.

If Σ is a set of topological $\mathcal{L}(V)$ -formulas, we write $\mathfrak{M} \models \Sigma(\iota)$ to mean $\mathfrak{M} \models Q(\iota)$ for each $Q \in \Sigma$. If R is a topological $\mathcal{L}(V)$ -formula, then we write $\Sigma \models R$ to mean $\mathfrak{M} \models R(\iota)$ for every \mathfrak{M} and ι such that $\mathfrak{M} \models \Sigma(\iota)$.

If T is a theory, \mathfrak{M} is a pre-structure, and $\mathfrak{M} \models T$, we say that \mathfrak{M} is a *pre-model of T* . If \mathfrak{M} is a reduced pre-structure, we say that it is a *reduced pre-model of T* . If \mathfrak{M} is a structure, we say that it is a *model of T* . We may refer to a (pre-)structure as a (pre-)model when we are implicitly thinking about a specific theory.

If we wish to emphasize that Σ is a set of $\mathcal{L}(V)$ -formulas, we will write $\Sigma(V)$ (again, typically with a tuple of variables).

If an open $\mathcal{L}(V)$ -formula U is known to be logically equivalent (modulo a partial type Σ) to a closed $\mathcal{L}(V)$ -formula F (i.e. $\mathfrak{M} \models \Sigma(\iota), U(\iota) \Leftrightarrow \mathfrak{M} \models \Sigma(\iota), F(\iota)$), then we may refer to either U or F as a *clopen $\mathcal{L}(V)$ -formula (over Σ)*. \triangleleft

Notation 1.3.10. Instead of writing $\varphi^{\mathfrak{M}}(\iota)$ we will typically have a particular tuple

\bar{x} of variable symbols in mind, and we will write $\varphi^{\mathfrak{M}}(\bar{a})$ to mean $\varphi^{\mathfrak{M}}(\iota)$ where ι is the variable assignment such that $\iota(x_i) = a_i$. We will do likewise for terms, formulas, and sets of formulas.

In all of these cases if we wish to emphasize the relevance of certain constants as parameters, we will also write them explicitly in expressions such as $\varphi(\bar{x}, \bar{a})$ or $\varphi(\bar{x}; \bar{a})$.

◁

Definition 1.3.11. Given a (pre-)structure \mathfrak{M} and a variable assignment ι for V , the *type of ι (in \mathfrak{M})*, written $\text{tp}_{\mathfrak{M}}(\iota)$, is the set of all closed $\mathcal{L}(V)$ -formulas F such that $\mathfrak{M} \models F(\iota)$. If $A \subseteq \mathfrak{M}$ is a set of elements, we will write $\text{tp}_{\mathfrak{M}}(\iota/A)$ for $\text{tp}_{\mathfrak{M}_A}(\iota)$. When \mathfrak{M} is clear from context we will typically write $\text{tp}(\iota)$ or $\text{tp}(\iota/A)$.

If ι and ι' are variable assignments for the same set of variables, and A is some set of parameters, we write $\iota \equiv_A \iota'$ to mean $\text{tp}(\iota/A) = \text{tp}(\iota'/A)$. We will omit A when it is empty.

The *theory of \mathfrak{M}* , written $\text{Th}(\mathfrak{M})$, is $\text{tp}_{\mathfrak{M}}(\emptyset)$. The *elementary diagram of \mathfrak{M}* , written $\text{eldiag}(\mathfrak{M})$, is the theory of $\mathfrak{M}_{\mathfrak{M}}$ as an $\mathcal{L}_{\mathfrak{M}}$ -structure. Two \mathcal{L} -(pre-)structures are *elementarily equivalent*, written $\mathfrak{M} \equiv \mathfrak{N}$, if $\text{Th}(\mathfrak{M}) = \text{Th}(\mathfrak{N})$.

◁

Definition 1.3.12. Given (pre-)structures, \mathfrak{M} and \mathfrak{N} , and a function $f : M \rightarrow N$, we say that f is an *elementary map*, written $f : \mathfrak{M} \preceq \mathfrak{N}$, if for every tuple $\bar{a} \in \mathfrak{M}$ and formula $\varphi \in \mathcal{L}(\bar{x})$, we have $\varphi^{\mathfrak{M}}(\bar{a}) = \varphi^{\mathfrak{N}}(f(\bar{a}))$, where $f(\bar{a})$ is the tuple such that $(f(\bar{a}))_i = f(a_i)$.

If \mathfrak{M} is a sub-(pre-)structure of \mathfrak{N} , we say that \mathfrak{M} is an *elementary sub-(pre-)structure of \mathfrak{N}* and conversely that \mathfrak{N} is an *elementary extension of \mathfrak{M}* , written $\mathfrak{M} \preceq \mathfrak{N}$, if the inclusion map is an elementary map.

◁

Remark 1.3.13. [BYBHU08] uses the notational conventions of writing $\mathfrak{M} \models \varphi$ to mean $\mathfrak{M} \models \varphi = 0$. I will avoid using this notation as I feel that its primary functionality is superseded by the language of open and closed formulas developed in Section 1.7, which are more directly analogous to formulas in discrete logic. Relatedly, I will refrain from any declaration that 0 represents true and 1 represents false (or vice versa).

What I am calling open and closed formulas are typically referred to as open and closed conditions, in [BYBHU08] and elsewhere. I feel that this terminology would be psychologically difficult to maintain while employing the machinery of Subsection 1.7, which we will employ frequently. There are also two other notions of formula, which we will use less frequently, specifically, X -valued formulas for arbitrary topological spaces X and type-set formulas (Definition 1.7.7). Given this proliferous overloading of the word ‘formula,’ I feel that it would be good to give a general organizing mindset for these four concepts.

An $\mathcal{L}(V)$ -formula is an object that takes an \mathcal{L} -structure and a V -assignment and produces some kind of output, either an element of some topological space, such as \mathbb{R} , or one of the values true and false. The free variables of an \mathcal{L} -formula are the smallest V such that it is an $\mathcal{L}(V)$ -formula. (Such a set doesn’t automatically exist, but we will always define formulas in such a way that they do.) This perspective might be disturbingly broad, as it would tell us that the operation of finding the type of a tuple is *itself* a type-space-valued formula, but I think that it is inevitable in the context of continuous logic. After all, any discrete theory in a countable signature has a real formula φ such that $a \equiv b$ if and only if $\varphi(a) = \varphi(b)$ (since every topologically separable, totally disconnected compact Hausdorff space embeds into \mathbb{R}).

The formulas we will consider here are usually first-order, in the sense that the value of the formula on some input (\mathfrak{M}, ι) only depends on the type of ι .⁸ (First-order) formulas are very closely related to functions on or subsets of the relevant type space (Definition 1.5.5). The only distinction is that an $\mathcal{L}(V)$ -formula is defined in such a way that it is automatically a $\mathcal{L}'(V')$ -formula for any $\mathcal{L}' \supseteq \mathcal{L}$ and $V' \supseteq V$, and this is primarily a matter of bookkeeping.

It is possible to reduce true-false formulas to X -valued formulas, specifically with different topologies on the set $\{\top, \perp\}$. The discrete topology results in clopen formulas, the two Sierpiński topologies result in open and closed formulas (in countable signatures), and the indiscrete topology results in type-set formulas (arbitrary sets of types), but this approach, while conceptually simpler, doesn't seem very pragmatic.

We will typically use lowercase Greek letters (e.g. φ, ψ, χ , etc.) for formulas taking on values in a topological space, and we will typically use uppercase Roman letters (e.g. F, G, U, V , etc.) for formulas that evaluate to true or false. The only systematic exception will be definable sets, in which case we will use D for both the corresponding closed formula and the corresponding distance predicate. The meaning of an instance of the word 'formula' without other clarifying words will hopefully be clear from context. A phrase such as 'a formula φ ' without additional context will typically refer to a real formula, analogously to the implicit convention in topology that a phrase such as 'a continuous function f ' refers to a real valued continuous function by default.

⁸There is some circularity in this statement, as types are themselves defined in terms of closed formulas. A philosophically pure approach would be to formalize first-order types directly and then define first-order formulas in terms of types. This is possible (without ultraproducts, even), but seems somewhat tedious and unwise pedagogically speaking, given most logicians' familiarity with the standard formalization of discrete first-order logic.

The following facts will be used frequently and implicitly.

Fact 1.3.14. *Let \mathcal{L} be a signature, \mathfrak{M} and \mathfrak{N} be \mathcal{L} -pre-structures, V be a set of variables, and ι be a V -assignment.*

1. *The concepts of quotient pre-structure and completion of a reduced pre-structure, given in Definition 1.2.2, are well defined. (As in the quotient and completion are themselves pre-structures.)*
2. *The collection of restricted $\mathcal{L}(V)$ -terms has cardinality at most $\aleph_0 + |\mathcal{L}| + |V|$. The collection of $\mathcal{L}(V)$ -terms has cardinality at most $(|\mathcal{L}| + |V|)^{\aleph_0}$.*
3. *The collection of restricted real $\mathcal{L}(V)$ -formulas has cardinality $\aleph_0 + |\mathcal{L}| + |V|$. The collection of real $\mathcal{L}(V)$ -formulas has cardinality $(\aleph_0 + |\mathcal{L}| + |V|)^{\aleph_0}$.*
4. *Every restricted formula contains at most finitely many predicate and function symbols and has at most finitely many free variables. Every formula contains at most countably many predicate and function symbols and has at most countably many free variables.*
5. *An \mathcal{L} -formula φ is an $\mathcal{L}(V)$ -formula if and only if $\text{fv}(\varphi) \subseteq V$.*
6. *For any formula $\varphi \in \mathcal{L}(V)$, $\varphi^{\mathfrak{M}}(\iota) \in I(\varphi)$.*
7. *If t is an $\mathcal{L}(\bar{x})$ -term where \bar{x} is an at most countable tuple of variable symbols, then there is a modulus ω_t , depending only on t and (if \bar{x} is infinite) the particular enumeration \bar{x} , such that $t^{\mathfrak{M}}$ is ω_t -uniformly continuous (as a function on $M^{|\bar{x}|}$, relative to the metric in Definition 1.2.1).*

8. If φ is an $\mathcal{L}(\bar{x})$ -formula where \bar{x} is an at most countable tuple of variable symbols, then there is a modulus ω_φ , depending only on φ and (if \bar{x} is infinite) the particular enumeration \bar{x} , such that $\varphi^{\mathfrak{M}}$ is ω_t -uniformly continuous (as a function on $M^{|\bar{x}|}$, relative to the metric in Definition 1.2.1).
9. We can freely convert between variable assignments and fresh constant symbols: If V is a set of variable symbols, \bar{v} is a tuple of distinct variable symbols, \bar{c} is a tuple of distinct constant symbols (not in $\mathcal{C}(\mathcal{L})$) of the same length, and \bar{a} is a tuple of elements of \mathfrak{M} of the same length, then for any $\mathcal{L}(V\bar{v})$ -formula $\varphi(\bar{v})$, $\varphi^{\mathfrak{M}}(\iota) = \varphi[\frac{\bar{c}}{\bar{v}}]^{\mathfrak{M}[\frac{\bar{a}}{\bar{c}}]}(\iota \upharpoonright V)$, where ι is a $V\bar{v}$ -assignment mapping v_i to a_i for each i , $\varphi[\frac{\bar{c}}{\bar{v}}]$ is the $\mathcal{L}_{\bar{c}}(V)$ -formula resulting from replacing each free instance of each v_i with c_i , and $\mathfrak{M}[\frac{\bar{a}}{\bar{c}}]$ is the $\mathcal{L}_{\bar{c}}$ -structure in which $c_i^{\mathfrak{M}} = a_i$.
10. The natural quotient map of \mathfrak{M} to ${}^r\mathfrak{M}$ is an elementary map, and the natural inclusion map of ${}^r\mathfrak{M}$ to $\overline{\mathfrak{M}}$ is an elementary map (so in particular a reduced pre-structure is always an elementary sub-structure of its completion).
11. Elementary maps are isometries, so in particular an elementary map from a reduced pre-structure to a pre-structure is necessarily injective. (Although elementary maps between pre-structures in general may fail to be injective.)
12. If there is an elementary map from \mathfrak{M} to \mathfrak{N} , then $\mathfrak{M} \equiv \mathfrak{N}$.
13. $\mathfrak{M} \preceq \mathfrak{N}$ if and only if \mathfrak{M} is a sub-pre-structure of \mathfrak{N} and $\mathfrak{M}_{\mathfrak{M}} \equiv \mathfrak{N}_{\mathfrak{M}}$.

Given Fact 1.3.14 parts 7 and 8 we will like to have canonical names for these moduli of uniform continuity. It is possible to compute sufficient moduli more directly than we do in the following definition, but for our purposes this is more trouble than it's

worth. (In particular, dealing with combining or rearranging infinite tuples of variables is tedious.)

Definition 1.3.15. For an at most countable tuple of variable symbols \bar{x} and an $\mathcal{L}(\bar{x}\bar{y})$ -term $t(\bar{x}\bar{y})$, we define the (*syntactic*) *modulus of uniform continuity of t with regards to \bar{x}* as

$$\omega_{t,\bar{x}}(r) = \inf \left\{ m \cdot r + c : m, c \geq 0, \right. \\ \left. (\forall \mathfrak{M}, \text{ an } \mathcal{L}_0\text{-str.})(\forall \bar{a}, \bar{b}, \bar{c} \in \mathfrak{M}) d(t^{\mathfrak{M}}(\bar{a}\bar{c}), t^{\mathfrak{M}}(\bar{b}\bar{c})) \leq m \cdot d(\bar{a}, \bar{b}) + c \right\},$$

where $\mathcal{L}_0 \subseteq \mathcal{L}$ is the reduct consisting only of those predicate and function symbols occurring in t . For an $\mathcal{L}(\bar{x}\bar{y})$ -formula $\varphi(\bar{x}\bar{y})$, the (*syntactic*) *modulus of uniform continuity of φ with regards to \bar{x}* , written $\omega_{\varphi,\bar{x}}$, is defined similarly. If \bar{x} has finite length and $\bar{y} = \emptyset$, we will typically write ω_t or ω_φ . ◁

We are calling the moduli ‘syntactic’ to emphasize that they do not depend on the particular structure in question, only the term or formula and the particular enumeration of variable symbols. The complexity of this definition comes from the fact that we have required that moduli of uniform continuity be continuous. The benefit of this is that we can use them freely in formulas as connectives. It is a straightforward but somewhat tedious exercise to verify that for any term $t(\bar{x}\bar{y})$ and any structure \mathfrak{M} and tuple $\bar{b} \in \mathfrak{M}$ of the same length as \bar{y} , the function $(\bar{x} \mapsto t^{\mathfrak{M}}(\bar{x}\bar{b})) : \mathfrak{M}^{|\bar{x}|} \rightarrow \mathfrak{M}$ is indeed $\omega_{t,\bar{x}}$ -uniformly continuous, and likewise for formulas. It is also not hard to verify (using Lemma A.1.7) that $\omega_{t,\bar{x}}$ and $\omega_{\varphi,\bar{x}}$

- are moduli, in the sense of Definition 0.3.9;

- are convex downwards, and therefore sub-additive;
- are bounded above by $\text{db}(\mathcal{L})$ and $\sup I(\varphi) - \inf I(\varphi)$, respectively;
- do not depend on the particular enumeration \bar{x} , if $|\bar{x}|$ is finite;

and that $\omega_{t,\bar{x}}$ and $\omega_{\varphi,\bar{x}}$ only depend on t and φ and not on the ambient signature. Note that a useful property of sub-additivity is that if d is a pseudo-metric and $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is sub-additive and satisfies $\alpha(0) = 0$, then $\alpha \circ d$ is a pseudo-metric.

Now that we have syntactic moduli for formulas, we can define the following familiar concept from discrete logic.

Definition 1.3.16. For any signature \mathcal{L} , the *Morleyization of \mathcal{L}* , written \mathcal{L}^{Mor} , is an expansion consisting of \mathcal{L} together with for each real formula $\varphi(\bar{x})$ a predicate symbol $P_{\varphi(\bar{x})}$ with $I(P_{\varphi(\bar{x})}) = I(\varphi)$, $a(P_{\varphi(\bar{x})}) = |\bar{x}|$, and $\omega_{P_{\varphi(\bar{x})}} = \omega_{\varphi,\bar{x}}$.

For any signature \mathcal{L} , the *admissibility axioms for \mathcal{L}* are axioms for each real formula $\varphi(\bar{x})$ of the form $\sup_{\bar{x}} |\varphi(\bar{x}) - P_{\varphi(\bar{x})}(\bar{x})| = 0$.

For any \mathcal{L} -theory T , the *Morleyization of T* , written T^{Mor} , is T together with the admissibility axioms for \mathcal{L} .

For any \mathcal{L} -pre-structure \mathfrak{M} , the *Morleyization of \mathfrak{M}* , written \mathfrak{M}^{Mor} , is the expansion of \mathfrak{M} to an \mathcal{L}^{Mor} -pre-structure satisfying $P_{\varphi(\bar{x})}^{\mathfrak{M}^{Mor}}(\bar{a}) = \varphi^{\mathfrak{M}}(\bar{a})$ for every formula φ and $\bar{a} \in \mathfrak{M}$. ◁

It is not hard to check that \mathfrak{M}^{Mor} is always an \mathcal{L}^{Mor} -pre-structure and that if $\mathfrak{M} \models T$, then $\mathfrak{M}^{Mor} \models T^{Mor}$.

1.4 The Logical Norm and the Density of Restricted Formulas

Definition 1.4.1. Given an $\mathcal{L}(V)$ -formula, φ , the *logical norm of φ* , written $\|\varphi\|_{\equiv}$, is given by

$$\|\varphi\|_{\equiv} = \sup\{|\varphi^{\mathfrak{M}}(\iota)|; \iota : V \rightarrow \mathfrak{M}\}.$$

Given a set of $\mathcal{L}(V)$ -formulas, Σ , the *logical norm of φ modulo Σ* , written $\|\varphi\|_{\Sigma}$, is given by

$$\|\varphi\|_{\Sigma} = \sup\{|\varphi^{\mathfrak{M}}(\iota)|; \iota : V \rightarrow \mathfrak{M} \models \Sigma(\iota)\},$$

where $\sup \emptyset = 0$.

Given two $\mathcal{L}(V)$ -formulas, φ, ψ , the *logical distance between φ and ψ* is $\|\varphi - \psi\|_{\equiv}$. The *logical distance between φ and ψ modulo Σ* is defined similarly. φ and ψ are *logically equivalent*, written $\varphi \equiv \psi$, if $\|\varphi - \psi\|_{\equiv} = 0$. φ and ψ are *logically equivalent modulo Σ* , written $\varphi \equiv_{\Sigma} \psi$, if $\|\varphi - \psi\|_{\Sigma} = 0$.

Given two $\mathcal{L}(V)$ -terms, t, s , and $X \in \{\equiv, \Sigma\}$, we define

$$d_X(t, s) = \|d(x, t) - d(x, s)\|_X,$$

where x is a variable symbol not in V . ◁

Note that if $\Sigma \subseteq \Sigma'$, then $\|\varphi\|_{\equiv} \geq \|\varphi\|_{\Sigma} \geq \|\varphi\|_{\Sigma'}$, and likewise for the metric on terms.

Lemma 1.4.2. *For any formula $\varphi(\bar{x}, \bar{y}), \psi(\bar{x}, \bar{y})$, with \bar{x} at most countable, and partial*

type $\Sigma(\bar{y})$, $\|\inf_{\bar{x}} \varphi - \inf_{\bar{x}} \psi\|_{\Sigma} \leq \|\varphi - \psi\|_{\Sigma}$, and likewise for sup.

Proof. For any \mathfrak{M} and $\bar{a} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, for any $\bar{b} \in \mathfrak{M}$, by assumption we have that $\mathfrak{M} \models \varphi(\bar{b}, \bar{a}) \leq \psi(\bar{b}, \bar{a}) + \|\inf_{\bar{x}} \varphi - \inf_{\bar{x}} \psi\|_{\Sigma}$, therefore $\mathfrak{M} \models \inf_{\bar{x}} \varphi(\bar{x}, \bar{a}) \leq \inf_{\bar{x}} \psi(\bar{x}, \bar{a}) + \|\inf_{\bar{x}} \varphi - \inf_{\bar{x}} \psi\|_{\Sigma}$. Therefore, by symmetry the required inequality holds.

The sup case follows from the inf case by considering $-\varphi$ and $-\psi$. \square

Proposition 1.4.3. $d_{\equiv}(t, s) = \sup_{\mathfrak{M}, \iota} d^{\mathfrak{M}}(t^{\mathfrak{M}}(\iota), s^{\mathfrak{M}}(\iota))$, and a similar fact holds for d_{Σ} .

Proof. For any r , if we have \mathfrak{M} and ι such that $d^{\mathfrak{M}}(t^{\mathfrak{M}}(\iota), s^{\mathfrak{M}}(\iota)) > r$, then we have a witness that $d_{\equiv}(t, s) > r$ as well, so we have that $d_{\equiv}(t, s) \geq \sup_{\mathfrak{M}, \iota} d^{\mathfrak{M}}(t^{\mathfrak{M}}(\iota), s^{\mathfrak{M}}(\iota))$. Conversely, if we have \mathfrak{N} , ι , and a such that $\mathfrak{N} \models |d(t(\iota), a) - d(s(\iota), a)| > r$, then by the reverse triangle inequality we have that $\mathfrak{N} \models d(t(\iota), s(\iota)) > r$, as well, so we have that $d_{\equiv}(t, s) \leq \sup_{\mathfrak{M}, \iota} d^{\mathfrak{M}}(t^{\mathfrak{M}}(\iota), s^{\mathfrak{M}}(\iota))$, and the desired equality follows. The proof for d_{Σ} is the same. \square

For the following proposition, note that, through a hilarious accident of history, the term ‘semi-norm’ is the norm analog of the term ‘pseudo-metric,’ as in both lift the requirement that distance zero things are identical.

Proposition 1.4.4. For $X \in \{\equiv, \Sigma\}$, d_X is a pseudo-metric on the set of $\mathcal{L}(V)$ -terms and $\|\cdot\|_X$ is a semi-norm on the set of $\mathcal{L}(V)$ -formulas, where we treat the collection of $\mathcal{L}(V)$ -formulas as an \mathbb{R} -vector space in the obvious way.

Proof. d_X is a supremum of a family of pseudo-metrics and is therefore a pseudo-metric. $\|\cdot\|_X$ is a supremum of a family of semi-norms and is therefore a semi-norm. (It does not matter that the indexing family is a proper class rather than a set.) \square

We will use the following fact from topology (which follows, in particular, from the Stone–Weierstrass theorem (see Fact A.2.4) and some basic identities involving max, min, and + which are mentioned in the proof of Proposition 1.4.12).

Fact 1.4.5. *For any $k \leq \omega$, if $\{I_n\}_{n < k}$ is a sequence of compact intervals and $F : \prod_{n < k} I_n \rightarrow \mathbb{R}$ is a continuous function, then for any $\varepsilon > 0$ there is a function $G : \mathbb{R}^k \rightarrow \mathbb{R}$ of the form $G(\bar{x}) = \max_{n < N} \min_{m < M} A_{nm}(\bar{x})$, where each A_{nm} is an affine function with rational coefficients and constant term (in particular, depending on only finitely many variables in \bar{x}), such that for all $\bar{x} \in \prod_{n < k} I_n$, $|F(\bar{x}) - G(\bar{x})| \leq \varepsilon$.*

Proposition 1.4.6. *For any \mathcal{L} and V ,*

- (i) *the set of restricted $\mathcal{L}(V)$ -terms is dense in the set of $\mathcal{L}(V)$ -terms under d_{\equiv} , and*
- (ii) *the set of restricted $\mathcal{L}(V)$ -formulas is dense in the set of $\mathcal{L}(V)$ -formulas under $\|\cdot\|_{\equiv}$.*

Furthermore, an $\mathcal{L}(V)$ -term or formula X can be approximated arbitrarily well by restricted $\mathcal{L}(V)$ -term or formulas Y such that $\text{fv}(Y) \subseteq \text{fv}(X)$.

Proof. (i) For $v \in V$, $d_{\equiv}(v, v) = 0$, so v is in the closure of the collection of restricted $\mathcal{L}(V)$ -terms. Let $f\bar{t}$ be an $\mathcal{L}(V)$ -term, and assume that each t_i is known to be the limit under d_{\equiv} of restricted $\mathcal{L}(V)$ -term with no additional free variables. Fix $\varepsilon > 0$. Find $\delta > 0$ small enough that $\omega_f(\delta) < \varepsilon$. Assume that $a(f) < \omega$. Then for each $i < \omega$, find s_i , a restricted $\mathcal{L}(V)$ -term, with $\text{fv}(s_i) \subseteq \text{fv}(t_i)$. Then we have that $d_{\equiv}(f\bar{t}, f\bar{s}) < \varepsilon$. Now assume that $a(f) = \omega$. Find $n < \omega$ large enough that $2^{-n} < \delta$. For each $i \leq n$, find s_i such that $d_{\equiv}(t_i, s_i) < \delta$. Now we have that $d_{\equiv}(f\bar{t}, f s_0 s_1 \dots s_{n-1} s_n s_n s_n \dots) < \varepsilon$.

So by induction, every $\mathcal{L}(V)$ -term can be approximated arbitrarily well by restricted $\mathcal{L}(V)$ -terms containing no additional free variables.

(ii) Let $P\bar{t}$ be an atomic $\mathcal{L}(V)$ -formula. By the same argument as in part (i), we can find restricted atomic $\mathcal{L}(V)$ -formulas $P\bar{s}$ containing no additional free variables, such that $\|P\bar{t} - P\bar{s}\|_{\equiv} < \varepsilon$ for arbitrarily small $\varepsilon > 0$.

Now given a $\mathcal{L}(V)$ -formula (for some V) of the form $F\bar{\varphi}$ for some connective F and some k -tuple of formulas $\bar{\varphi}$, assume that for each φ_i we've shown that there are restricted $\mathcal{L}(V)$ -formulas arbitrarily close to φ_i under $\|\cdot\|_{\equiv}$.

Fix $\varepsilon > 0$. Using Fact 1.4.5, find a G (of the form described in the fact) such that for all $\bar{x} \in X$, $|F(\bar{x}) - G(\bar{x})| \leq \frac{\varepsilon}{2}$, and such that G only depends on the first ℓ variables for some finite $\ell \leq k$. Since G is piecewise affine, it is Lipschitz, and there is a $\delta > 0$ such that if $\|\bar{x} - \bar{y}\|_{\infty} \leq \delta$, then $|G(\bar{x}) - G(\bar{y})| \leq \frac{\varepsilon}{2}$.

So now for each φ_i with $i \leq \ell$, find a restricted $\mathcal{L}(V)$ -formula ψ_i such that $\|\varphi_i - \psi_i\|_{\equiv} \leq \delta$. Now consider the formula $G\bar{\psi} = G\psi_0\psi_1 \dots \psi_{\ell-1}$. By construction, for any \mathfrak{M} and V -assignment ι , we have that

$$\mathfrak{M} \models |G\bar{\psi}(\iota) - F\bar{\varphi}(\iota)| \leq |G\bar{\psi}(\iota) - G\bar{\varphi}(\iota)| + |G\bar{\varphi}(\iota) - F\bar{\varphi}(\iota)| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Therefore $\|F\bar{\varphi} - G\bar{\psi}\|_{\equiv} \leq \varepsilon$. Since we can do this for any $\varepsilon > 0$, we have that restricted $\mathcal{L}(V)$ -formulas are arbitrarily close to $F\bar{\varphi}$ under $\|\cdot\|_{\equiv}$.

Now given $\mathcal{L}(V)$ -formula of the form $Q\bar{x}\varphi$ (where $Q \in \{\inf, \sup\}$), assume that we've shown that φ is approximated arbitrarily well by restricted $\mathcal{L}(Vx)$ -formulas under $\|\cdot\|_{\equiv}$. Pick $\varepsilon > 0$, and let ψ be a restricted $\mathcal{L}(Vx)$ -formula such that $\|\varphi - \psi\|_{\equiv} \leq \varepsilon$. Let \bar{x}_0 be the initial segment of \bar{x} consisting of those variables that actually occur in ψ (this is

always finite). Direct computation shows that $\|Q\bar{x}\varphi - Q\bar{x}_0\psi\|_{\equiv} \leq \varepsilon$. Since we can do this for arbitrarily small $\varepsilon > 0$, we have that restricted $\mathcal{L}(V)$ -formulas are arbitrarily close to $Qx\varphi$ under $\|\cdot\|_{\equiv}$.

Therefore by induction (simultaneously for different choices of V), we have that for any V , any $\mathcal{L}(V)$ -formula φ can be approximated arbitrarily well by restricted $\mathcal{L}(V)$ -formulas under $\|\cdot\|_{\equiv}$.

All of the restricted formulas ψ constructed in this proof to approximate a formula φ have the property that $\text{fv}(\psi) \subseteq \text{fv}(\varphi)$, so the last statement holds as well. \square

Proposition 1.4.7. *For any \mathcal{L} and any V , the collection of $\mathcal{L}(V)$ -formulas modulo logical equivalence is a complete metric space under $\|\cdot\|_{\equiv}$. For any set of $\mathcal{L}(V)$ -formulas Σ , the same statement holds for the collection of $\mathcal{L}(V)$ -formulas modulo logical equivalence modulo Σ , with the norm $\|\cdot\|_{\Sigma}$.*

Proof. Let $\{\chi_i\}_{i < \omega}$ be a sequence of $\mathcal{L}(V)$ -formulas such that for each $i < \omega$, $\|\chi_i - \chi_{i+1}\| \leq 2^{-i}$. Let $\eta_0 = \chi_0$, and for each $i > 0$, let $\eta_i = \chi_i - \chi_{i-1}$. Note that by construction, for any $i > 0$, for any \mathfrak{M} , and for any V -assignment ι , we have that $\mathfrak{M} \models -2^{-i} \leq \eta_i(\iota) \leq 2^{-i}$.

Let $\zeta_0 = \eta_0$, and for $i > 0$ let $\zeta_i = [\eta_i]_{-2^{-i}}^{2^{-i}}$. Note that by construction $\eta_i \equiv \zeta_i$, but for $i > 0$, $I(\zeta_i) \subseteq [-2^{-i}, 2^{-i}]$. So now we have that $\psi = \sum_{i < \omega} \zeta_i$ is an $\mathcal{L}(V)$ -formula. A direct computation shows that $\lim_{i \rightarrow \infty} \chi_i^{\mathfrak{M}}(\iota) = \psi^{\mathfrak{M}}(\iota)$ for all \mathfrak{M} and ι .

The same proof works for formulas modulo Σ . \square

Corollary 1.4.8. *For any \mathcal{L} and any V , the collection of $\mathcal{L}(V)$ -formulas modulo logical equivalence is a unital real Banach algebra under $\|\cdot\|_{\equiv}$, with operations $\varphi \mapsto r \cdot \varphi$, $(\varphi, \psi) \mapsto \varphi + \psi$, and $(\varphi, \psi) \mapsto \varphi \cdot \psi$.*

For any set of $\mathcal{L}(V)$ -formulas Σ , the collection of $\mathcal{L}(V)$ -formulas modulo logical equivalence modulo Σ also forms a real Banach algebra. If Σ is satisfiable (i.e. there exists an \mathfrak{M} and an ι such that $\mathfrak{M} \models \Sigma(\iota)$), then this algebra is unital. If Σ is unsatisfiable, then this algebra is trivial.

Now we will see that, despite the formal similarity between the syntax of continuous first-order logic and $\mathcal{L}_{\omega_1\omega}$, continuous logic does not actually have meaningful formula rank above ω (either in the sense of syntactic rank or the sense of quantifier rank) up to logical equivalence.

Corollary 1.4.9. *Every formula φ is logically equivalent to a formula $\psi = F\bar{\chi}$, where $\bar{\chi} = \{\chi_i\}_{i < \omega}$ is a sequence of restricted formulas.*

Proof. Follow the proof of Proposition 1.4.7 with χ_i being a sequence of restricted formulas limiting to φ . Let ψ be the formula constructed from the χ_i 's. By construction ψ is logically equivalent to φ . \square

Corollary 1.4.10. *For any $\mathcal{L}(V)$ -formula φ and any $\varepsilon > 0$, there is a restricted $\mathcal{L}(V)$ -formula ψ such that $\text{fv}(\psi) \subseteq \text{fv}(\varphi)$ and such that $\|\varphi - \psi\|_{\equiv} \leq \varepsilon$ and $\mathfrak{M} \models \varphi(\iota) \geq \psi(\iota)$ for every \mathfrak{M} and V -assignment ι . (And likewise for \leq .)*

Proof. Assume without loss that ε is rational. Find a restricted $\mathcal{L}(V)$ -formula χ such that $\|\varphi - \chi\|_{\equiv} \leq \frac{\varepsilon}{3}$ (and such that $\text{fv}(\chi) \subseteq \text{fv}(\varphi)$), and let $\psi = \chi - \frac{\varepsilon}{2}$, which is a restricted $\mathcal{L}(V)$ -formula that satisfies the required properties.

The \leq case follows by applying the \geq case to $-\varphi$. \square

Definition 1.4.11. A formula is in *prenex form* if it is a string of quantifiers followed by a quantifier-free formula.

A formula is in *prenex maximal affine normal form*, or *p.max.a.n.f.*, if it is of the form $Q_0x_0Q_1x_1\dots\max_{n<N}\min_{m<M}\varphi_{nm}$, with each φ_{nm} an affine atomic formula.

A formula is in *prenex minimal affine normal form*, or *p.min.a.n.f.*, if it is of the form $Q_0x_0Q_1x_1\dots\min_{n<N}\max_{m<M}\varphi_{nm}$, with each φ_{nm} an affine atomic formula. \triangleleft

Proposition 1.4.12. *Every restricted $\mathcal{L}(V)$ -formula is logically equivalent to a p.max.a.n.f. restricted $\mathcal{L}(V)$ -formula (and to a p.min.a.n.f. restricted $\mathcal{L}(V)$ -formula).*

Proof. For this proof we will write $x \uparrow y = \max\{x, y\}$ and $x \downarrow y = \min\{x, y\}$.

We'll say that a restricted formula is in *scaling normal form*, or *s.n.f.*, if whenever it has a subformula of the form $r \cdot \varphi$, φ is either 1 or an atomic formula. A simple inductive argument shows that every restricted formula is logically equivalent to a restricted formula in s.n.f. We'll say that a restricted formula is in *prenex scaling normal form*, or *p.s.n.f.*, if it is in s.n.f. and furthermore is of the form $Q_0x_0Q_1x_1\dots Q_nx_n\varphi$ where φ is quantifier-free.

We need to show that every restricted formula in s.n.f. is logically equivalent to one in p.s.n.f. To show this, it is sufficient to show that one can move quantifiers out of expressions involving binary connectives. The rest follows by induction. For $\square \in \{+, \uparrow, \downarrow\}$, $\varphi \square Q\dot{v}_i\psi$ is logically equivalent to $Q\dot{v}_j\varphi\square(\psi[\frac{\dot{v}_j}{\dot{v}_i}])$, where j is the smallest index such that \dot{v}_j does not occur in φ or ψ . The other cases follow by symmetry. A restricted formula in s.n.f. is still in s.n.f. after applying these operations, so we have shown that every restricted formula is equivalent to a p.s.n.f. formula.

Finally, we just need to show that a quantifier-free s.n.f. restricted formula is always logically equivalent to one of the form $\max_{n<N}\min_{m<M}\varphi_{nm}$, with each φ_{nm} a restricted affine atomic formula. The proof of this is essentially identical to the proof in discrete

propositional logic that a sentence is always equivalent to one in conjunctive normal form. The extra wrinkle is the presence of the $+$ connective. The relevant facts are these: $x \uparrow (y \downarrow z) = (x \uparrow y) \downarrow (x \uparrow z)$, $x \downarrow (y \uparrow z) = (x \downarrow y) \uparrow (x \downarrow z)$, $x + (y \uparrow z) = (x + y) \uparrow (x + z)$, and $x + (y \downarrow z) = (x + y) \downarrow (x + z)$.

Applying this last step to the quantifier free part of a p.s.n.f. formula and noting that if two prenex formulas have the same quantifier string and have logically equivalent quantifier-free parts, then they are logically equivalent, we get that every restricted formula is logically equivalent to a p.max.a.n.f. formula.

The proof of the p.min.a.n.f. case is essentially the same. \square

1.5 Compactness I: Type Spaces

Definition 1.5.1. A set of closed $\mathcal{L}(V)$ -formulas is called a *partial $\mathcal{L}(V)$ -type*, or just a *partial V -type* if \mathcal{L} is clear by context. A *partial type* is a partial $\mathcal{L}(V)$ -type for some V . An n -type for some ordinal n is a V -type for some set of variables enumerated by n . An *(\mathcal{L} -)theory* is a partial $\mathcal{L}(\emptyset)$ -type.

We say that a partial type Σ *extends* a partial type Π if Σ is a superset of Π .

A V -type Σ is *satisfiable* if there is a structure \mathfrak{M} and a V -assignment ι such that $\mathfrak{M} \models \Sigma(\iota)$. A type is *finitely satisfiable* if every finite subset of it is satisfiable.

A *complete $\mathcal{L}(V)$ -type* is an $\mathcal{L}(V)$ -type that is maximally finitely satisfiable, i.e. it is finitely satisfiable and no larger set of $\mathcal{L}(V)$ -formulas is finitely satisfiable. \triangleleft

Lemma 1.5.2. *Let p be a complete V -type. For every pair of $\mathcal{L}(V)$ -formulas, φ and ψ , either $(\varphi \leq \psi) \in p$ or $(\varphi \geq \psi) \in p$.*

Proof. Assume that for some $\mathcal{L}(V)$ -formulas φ and ψ , neither $p \cup \{\varphi \leq \psi\}$ nor $p \cup \{\varphi \geq \psi\}$

are finitely satisfiable. This implies that there are finite $\Sigma_0, \Sigma_1 \subseteq p$ such that $\Sigma_0 \cup \{\varphi \leq \psi\}$ and $\Sigma_1 \cup \{\varphi \geq \psi\}$ are each not satisfiable. $\Sigma_0 \cup \Sigma_1$ is finite and so is satisfiable by assumption. Let $\mathfrak{M} \models \Sigma_0 \cup \Sigma_1$. Now either $\mathfrak{M} \models \varphi \leq \psi$ or $\mathfrak{M} \models \varphi \geq \psi$. In the first case, we contradict that $\Sigma_0 \cup \{\varphi \leq \psi\}$ is not satisfiable, and in the second case we contradict that $\Sigma_1 \cup \{\varphi \geq \psi\}$ is not satisfiable. So we have a contradiction, and one of $p \cup \{\varphi \leq \psi\}$ and $p \cup \{\varphi \geq \psi\}$ must be finitely satisfiable. By maximality it must be the case that either $(\varphi \leq \psi) \in p$ or $(\varphi \geq \psi) \in p$. \square

Corollary 1.5.3. *Let p be a complete V -type. For every $\mathcal{L}(V)$ -formula φ , there is a unique $r \in I(\varphi)$ such that for every $s < r$, $(\varphi \geq s) \in p$, and for every $s > r$, $(\varphi \leq s) \in p$.*

Proof. By the lemma, for every $s \in I(\varphi)$, either $(\varphi \leq s)$ or $(\varphi \geq s)$ is in p . By maximal finite satisfiability the set $A = \{s \in I(\varphi) : (\varphi \geq s) \in p\}$ must be an initial segment of $I(\varphi)$ and the set $B = \{s \in I(\varphi) : (\varphi \leq s) \in p\}$ must be a final segment of $I(\varphi)$. Furthermore, $A \cap B$ must be at most one point and in particular $\sup A = \inf B$. Let $r = \sup A = \inf B$. \square

Definition 1.5.4. For any complete V -type p and any V -formula φ , let $\varphi(p)$ denote the unique r from Corollary 1.5.3. \triangleleft

Note that if $\varphi(p) = \varphi(q)$ for all $\mathcal{L}(V)$ -formulas, then $p = q$.

Definition 1.5.5. • Given a partial $\mathcal{L}(V)$ -type Σ , the *space of (complete) $\mathcal{L}(V)$ -types (over Σ)* also called a *type space* or a *Stone space*,⁹ written $S_V(\Sigma, \mathcal{L})$ or $S_V(\Sigma)$ if \mathcal{L} is clear by context, is the set of all complete $\mathcal{L}(V)$ -types p such that

⁹By analogy with the Stone spaces in discrete logic, despite the fact that $S_V(\Sigma)$ is not necessarily the Stone space of any Boolean algebra. By rights they should be called Gelfand spaces, in light of Fact 1.5.6, but c'est la vie.

$p \supseteq \Sigma$. If $\Sigma = \emptyset$, we will write $S_V(\mathcal{L})$ or S_V for $S_V(\emptyset, \mathcal{L})$. In this notation V will typically be an ordinal or an explicit tuple of variables.

- If \mathfrak{M} is a structure and $A \subseteq \mathfrak{M}$ is some set of elements, then $S_V(A)$ is $S_V(\text{Th}(\mathfrak{M}_A))$.
- For a closed formula $\varphi \square \psi$, $\llbracket \varphi \square \psi \rrbracket_{V, \Sigma, \mathcal{L}} = \{p \in S_V(\Sigma, \mathcal{L}) : (\varphi \square \psi) \in p\}$. For an open formula $\varphi \square \psi$, $\llbracket \varphi \square \psi \rrbracket_{V, \Sigma, \mathcal{L}} = S_V(\Sigma, \mathcal{L}) \setminus \llbracket \neg(\varphi \square \psi) \rrbracket$. When V , Σ , and \mathcal{L} are clear from context we will omit them. We will also write expressions of the form $\llbracket \varphi \leq \psi < \chi \rrbracket$, with the obvious meaning.
- Sets of the form $\llbracket \varphi = 0 \rrbracket$ (as well as their sets of realizations in structures) for formulas φ are called *zerosets*. All closed formulas can be written in this form.
- Each type space $S_V(\Sigma)$ is endowed with a *logic topology*. This is the topology generated by sets of the form $\llbracket \varphi < \psi \rrbracket$. ◁

Typically we will consider type spaces over an \mathcal{L} -theory, rather than an arbitrary partial $\mathcal{L}(V)$ -type.

We will present the following fact without proof, as it will not be used in this thesis.

Fact 1.5.6. *$S_V(\Sigma)$ is the Gelfand spectrum of the Banach algebra of $\mathcal{L}(V)$ -formulas modulo logical equivalence modulo Σ .*

Lemma 1.5.7. *For any V and Σ , sets of the form $\llbracket \varphi < \psi \rrbracket$ form a base for the logic topology (not just a sub-base). $\llbracket \varphi > \psi \rrbracket$ and $\llbracket \varphi \neq \psi \rrbracket$ are open sets as well. Furthermore, sets of the form $\llbracket \varphi < \psi \rrbracket$ for restricted φ and ψ form a base for the logic topology.*

Proof. The second statement follows easily from the facts that $\llbracket \varphi > \psi \rrbracket = \llbracket \psi < \varphi \rrbracket$ and $\llbracket \varphi \neq \psi \rrbracket = \llbracket 0 < |\varphi - \psi| \rrbracket$.

For the first statement, all we need to do is show that these sets are closed under finite intersections. Since $\llbracket \varphi < \psi \rrbracket = \llbracket 0 < \psi - \varphi \rrbracket$, it's sufficient to consider sets of the form $\llbracket 0 < \varphi \rrbracket$. $\llbracket 0 < \varphi \rrbracket \cap \llbracket 0 < \psi \rrbracket = \llbracket 0 < \min\{\varphi, \psi\} \rrbracket$, so we have the required result.

To see that restricted open formulas form a base for the logic topology, note that we just showed that the intersection of two restricted open formulas is still equivalent to a restricted open formula, so they are closed under intersection. The fact that they form a sub-base (and therefore a base) follows from Proposition 1.4.6. \square

Proposition 1.5.8 (Compactness, Part I). *For any V and Σ , $S_V(\Sigma)$ is a compact Hausdorff space under the logic topology. Furthermore, the topological weight of $S_V(\Sigma)$ is at most $\aleph_0 + |\mathcal{L}| + |V|$.*

Proof. To see that $S_V(\Sigma)$ is Hausdorff, let p and q be distinct complete types. Since p and q are distinct we have that $\varphi(p) \neq \varphi(q)$ for some $\mathcal{L}(V)$ -formula φ . Assume without loss that $\varphi(p) < \varphi(q)$, and let r and s be such that $\varphi(p) < r < s < \varphi(q)$. Then we have that $\llbracket \varphi < r \rrbracket$ and $\llbracket \varphi > s \rrbracket$ are disjoint neighborhoods of p and q . Therefore the space is Hausdorff.

For a closed formula F , $\llbracket F \rrbracket$ is non-empty if and only if it is satisfiable. By Lemma 1.5.7, it's sufficient to show that if Π is a family of closed $\mathcal{L}(V)$ -formulas such that $\Sigma \cup \Pi$ is finitely satisfiable, then there is a $p \in S_V(\Sigma)$ such that $p \supseteq \Pi$, but this follows immediately from Zorn's lemma.

To compute the topological weight of $S_V(\Sigma)$, note that restricted open formulas form a base of the logic topology, by Lemma 1.5.7. \square

Proposition 1.5.9. *Let V be a set of variables, and let Σ be a partial V -type.*

(i) *For any $\mathcal{L}(V)$ -formula φ , the function $p \mapsto \varphi(p) : S_V(\Sigma) \rightarrow \mathbb{R}$ is continuous.*

(ii) If $f : S_V(\Sigma) \rightarrow \mathbb{R}$ is continuous, then there is an $\mathcal{L}(V)$ -formula φ such that $f(p) = \varphi(p)$ for all $p \in S_V(\Sigma)$.

Proof. (i) It is sufficient to check that the preimages of half-infinite intervals in \mathbb{R} are open in $S_V(\Sigma)$, but $\varphi^{-1}((-\infty, r)) = \llbracket \varphi < r \rrbracket$ and $\varphi^{-1}((r, \infty)) = \llbracket \varphi > r \rrbracket$, so this is trivial.

(ii) By the Stone–Weierstrass theorem (Fact A.2.5), for any continuous f , there is a sequence of formulas $\{\psi_i\}_{i < \omega}$ such that $\sup\{|\psi_i(p) - f(p)| : p \in S_V(\Sigma)\} \rightarrow 0$ as $i \rightarrow \infty$. By thinning the sequence we may assume that $\sup\{|\psi_i(p) - \psi_{i+1}(p)| : p \in S_V(\Sigma)\} \leq 2^{-i}$ for all $i < \omega$. Now $\psi_0 + \sum_{i < \omega} [\psi_{i+1} - \psi_i]_{-2^{-i}}^{2^{-i}}$ is the required formula. \square

In light of this proposition we will often conflate $\mathcal{L}(V)$ -formulas (especially modulo logical equivalence modulo Σ) and continuous functions on $S_V(\Sigma)$.

Corollary 1.5.10. *If U is an open subset of some type space $S_V(\Sigma)$, then $U = \llbracket 0 < \varphi \rrbracket$ for some formula φ if and only if U is an F_σ set (i.e. a countable union of closed sets).*

Likewise if F is a closed subset of $S_V(\Sigma)$, then $F = \llbracket 0 \leq \varphi \rrbracket$ for some formula φ if and only if F is a G_δ set (i.e. a countable intersection of open sets).

Proof. Proposition 1.5.9 and Fact A.2.9. \square

The proof of the following proposition technically needs full compactness (Corollary 1.6.5), but we have included it here because of its relevance to this subsection.

Proposition 1.5.11 (Maps between Type Spaces). *Let \mathcal{L} and \mathcal{L}' be signatures, V and W be sets of variable symbols, and Σ and Π be sets of $\mathcal{L}(V)$ -formulas.*

(i) *If $\Sigma \subseteq \Pi$, then $S_V(\Pi) \subseteq S_V(\Sigma)$. The inclusion map is continuous.*

(ii) If $f : W \rightarrow V$ is a function, then for any complete type $p \in S_V(\Sigma)$, the set

$$f^*(p) = \left\{ F \in C\mathcal{L}(W) : F \left[\frac{f(\bar{w})}{\bar{w}} \right] \in p \right\},$$

is a complete type in $S_W(\Sigma')$ where \bar{w} is an enumeration of W , $f(\bar{w})_i = f(w_i)$, $(\varphi \square \psi) \left[\frac{\bar{a}}{\bar{x}} \right] = (\varphi \left[\frac{\bar{a}}{\bar{x}} \right] \square \psi \left[\frac{\bar{a}}{\bar{x}} \right])$, and $\Sigma' = \left\{ F \in C\mathcal{L}(W) : F \left[\frac{f(\bar{w})}{\bar{w}} \right] \in \Sigma \right\}$.

Furthermore, the function $f^* : S_V(\Sigma) \rightarrow S_W(\Sigma')$ is continuous, and if $W \subseteq V$, $f : W \rightarrow V$ is the inclusion map, and Σ is a partial $\mathcal{L}(W)$ -type (in particular, if Σ is a theory), then f^* is an open map (image of an open set is open).

(iii) If \mathcal{L}' is a reduct of \mathcal{L} , then for any complete type $p \in S_V(\Sigma)$, $p \upharpoonright \mathcal{L}'$ is a complete type in $S_V(\Sigma \upharpoonright \mathcal{L}')$.

Proof. (i) This is trivial.

(ii) Firstly, to see that $f^*(p)$ is finitely satisfiable, let $q_0 \subseteq f^*(p)$ be a finite set of formulas. Let \mathfrak{M} be a structure and ι a V -assignment such that $\mathfrak{M} \models q_0 \left[\frac{f(\bar{w})}{\bar{w}} \right] (\iota)$, which exists since p is finitely satisfiable. Then $\mathfrak{M} \models q_0(\iota \circ f)$, so $f^*(p)$ is finitely satisfiable.

Let φ and ψ be $\mathcal{L}(V)$ -formulas. $\chi \left[\frac{f(\bar{w})}{\bar{w}} \right] (p) = r_\chi$ for $\chi \in \{\varphi, \psi\}$. There are three cases—

- $r_\varphi < r_\psi$, in which case $(\varphi \leq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \in p$ and $(\varphi = \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right], (\varphi \geq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \notin p$,
- $r_\varphi = r_\psi$, in which case $(\varphi \leq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right], (\varphi = \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right], (\varphi \geq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \in p$,¹⁰ and
- $r_\varphi > r_\psi$, in which case $(\varphi \geq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \in p$ and $(\varphi \leq \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right], (\varphi = \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \notin p$.

¹⁰This case needs full compactness.

This implies that no proper superset of $f^*(p)$ is finitely satisfiable. Therefore $f^*(p)$ is a complete V -type.

To see that f^* is continuous, note that if $\llbracket \varphi < \psi \rrbracket$ is an $\mathcal{L}(W)$ -formula, then $\llbracket U \rrbracket = \llbracket (\varphi < \psi) \left[\frac{f(\bar{w})}{\bar{w}} \right] \rrbracket$ is an $\mathcal{L}(V)$ -formula such that $\llbracket U \rrbracket = (f^*)^{-1}(\llbracket \varphi < \psi \rrbracket)$. Therefore f^* is continuous.

Now assume that $W \subseteq V$, that f is the inclusion map and that Σ is a partial $\mathcal{L}(W)$ -type. It is sufficient to show that $f^*(\llbracket \varphi > 0 \rrbracket)$ is open for any real $\mathcal{L}(V)$ -formula φ . Assume that $p \in f^*(\llbracket \varphi > 0 \rrbracket)$, then there exists $q \in S_V(\Sigma)$ such that $p = f(q)$ and $q \in \llbracket \varphi > 0 \rrbracket$. Let \bar{w} be an enumeration of $\text{fv}(\varphi) \setminus V$. Let \mathfrak{M} and $\iota : W \rightarrow M$ be such that $\text{tp}(\iota) = q$. Clearly we have that $\mathfrak{M} \models \sup_{\bar{w}}(\iota \upharpoonright V)\varphi > 0$, so we have that $p \in \llbracket \sup_{\bar{w}} \varphi > 0 \rrbracket$. So we have shown that $f^*(\llbracket \varphi > 0 \rrbracket_{V,\Sigma}) \subseteq \llbracket \sup_{\bar{w}} \varphi > 0 \rrbracket_{W,\Sigma}$.

Now assume that $p \in \llbracket \sup_{\bar{w}} \varphi > 0 \rrbracket$. Let \mathfrak{M} and $\iota : W \rightarrow M$ be such that $\text{tp}(\iota) = p$. Since $\mathfrak{M} \models \sup_{\bar{w}} \varphi(\iota) > 0$, we can find an extension $\iota' : V \rightarrow M$ of ι such that $\mathfrak{M} \models \varphi(\iota') > 0$. Let $q = \text{tp}(\iota')$. Since Σ is an $\mathcal{L}(W)$ -type, we have that $\text{tp}(\iota') \in S_W(\Sigma)$, so we have that $p \in f^*(\llbracket \varphi > 0 \rrbracket)$. Therefore $\llbracket \sup_{\bar{w}} \varphi > 0 \rrbracket_{W,\Sigma} \subseteq f^*(\llbracket \varphi > 0 \rrbracket_{V,\Sigma})$, and we have that $f^*(\llbracket \varphi > 0 \rrbracket)_{V,\Sigma} = \llbracket \sup_{\bar{w}} \varphi > 0 \rrbracket_{W,\Sigma}$, hence $f^*(\llbracket \varphi > 0 \rrbracket)_{V,\Sigma}$ is an open set.

(iii) This follows immediately from the fact that any open $\mathcal{L}'(V)$ -formula is also an open $\mathcal{L}(V)$ -formula. \square

1.6 Compactness II: Ultraproducts

There is a high brow definition of ultraproduct (in discrete logic) that makes Łoś's theorem compatible with empty structures [Bar86, appendix]. While it does generalize to continuous logic, we will opt for an uglier equivalent definition.

Definition 1.6.1. Let $\{\mathfrak{M}_i\}_{i \in I}$ be a family of \mathcal{L} -pre-structures. Let \mathcal{U} be an ultrafilter on the index set I . The *pre-ultraproduct* of $\{\mathfrak{M}_i\}_{i \in I}$ (with regards to \mathcal{U}), written $\mathfrak{M}_{\mathcal{U}}^0$, is the \mathcal{L} -pre-structure whose underlying set is $M_{\mathcal{U}}$ where:

- If $\{i \in I : M_i = \emptyset\} \in \mathcal{U}$, then $M_{\mathcal{U}} = \emptyset$.
- If $\{i \in I : M_i \neq \emptyset\} = I' \in \mathcal{U}$, $M_{\mathcal{U}}^0 = \prod_{i \in I'} M_i$.

For any predicate symbol P and any tuple $\bar{a} \in (M_{\mathcal{U}}^0)^{a(P)}$, we have $P^{\mathfrak{M}_{\mathcal{U}}^0}(\bar{a}) = \lim_{\mathcal{U}} P^{\mathfrak{M}_i}(\bar{a}_i)$, where $\lim_{\mathcal{U}}$ is the ultrafilter assisted limit, i.e. $P^{\mathfrak{M}_{\mathcal{U}}^0}(\bar{a})$ is the unique r such that for every open neighborhood $V \ni r$, $\{i \in I : P^{\mathfrak{M}_i}(\bar{a}_i) \in V\} \in \mathcal{U}$. Since the values of $P^{\mathfrak{M}_i}(\bar{a}_i)$ are contained in a compact interval, this is well defined.

For any function symbol f and any tuple $\bar{a} \in (M_{\mathcal{U}}^0)^{a(f)}$, we have $f^{\mathfrak{M}_{\mathcal{U}}^0}(\bar{a})(i) = f^{\mathfrak{M}_i}(\bar{a}_i)$.

For any family of \mathcal{L} -pre-structures $\{\mathfrak{M}_i\}_{i \in I}$, the *ultraproduct* of $\{\mathfrak{M}_i\}_{i \in I}$ (with regards to \mathcal{U}), written $\mathfrak{M}_{\mathcal{U}}$, is ${}^r\mathfrak{M}_{\mathcal{U}}^0$.

We extend (pre-)ultraproducts to families of the form $\{\mathfrak{M}_i, \iota_i\}_{i \in I}$ (\mathcal{L} -(pre-)structures and corresponding V -assignments for some fixed set of variables V) in the obvious way. We will write $\mathfrak{M}_{\mathcal{U}}^0$ and $\iota_{\mathcal{U}}$ for this.

A (pre-)ultraproduct in which all factors are the same structure is called a (pre-)ultrapower. We will write the pre-ultrapower $\mathfrak{M}^{0, \mathcal{U}}$ and the ultrapower $\mathfrak{M}^{\mathcal{U}}$. \triangleleft

Proposition 1.6.2 (Łoś's theorem). *For any set I , any family of \mathcal{L} -pre-structures $\{\mathfrak{M}_i\}_{i \in I}$, any any family of V -assignments $\{\iota_i\}_{i \in I}$, we have that $\mathfrak{M}_{\mathcal{U}}^0$ is actually a pre-structure (i.e. $d^{\mathfrak{M}_{\mathcal{U}}^0}$ is actually a metric, for each $s \in \mathcal{P} \cup \mathcal{F}$, $s^{\mathfrak{M}_{\mathcal{U}}^0}$ is ω_s -uniformly continuous, and for each $P \in \mathcal{P}$ and $\bar{a} \in M_{\mathcal{U}}^0$, $P^{\mathfrak{M}_{\mathcal{U}}^0} \in I(P)$) and for any $\mathcal{L}(V)$ -formula φ we have $\varphi^{\mathfrak{M}_{\mathcal{U}}^0}(\iota_{\mathcal{U}}) = \lim_{\mathcal{U}} \varphi^{\mathfrak{M}_i}(\iota_i)$ (and the same holds for $\mathfrak{M}_{\mathcal{U}}$).*

Proof. First note that by the compactness of $I(P)$ we clearly have $P^{\mathfrak{M}_U^0}(\bar{a}) \in I(P)$ for each $P \in \mathcal{P}$ and $\bar{a} \in M_U^0$.

Now, for any $a, b, c \in M_U^0$, we clearly have $d^{\mathfrak{M}_U^0}(a, a) = 0$, $d^{\mathfrak{M}_U^0}(a, b) = d^{\mathfrak{M}_U^0}(b, a) \geq 0$, and $d^{\mathfrak{M}_U^0}(a, c) \leq d^{\mathfrak{M}_U^0}(a, b) + d^{\mathfrak{M}_U^0}(b, c)$, because these correspond to closed subsets of $I(d)$, $I(d)^2$, and $I(d)^3$, respectively, and ultrafilter assisted limits stay inside closed sets. Similarly, for any $P \in \mathcal{P}$ with $a(P)$ finite, the same reasoning gives that $|P^{\mathfrak{M}_U^0}(\bar{a}) - P^{\mathfrak{M}_U^0}(\bar{b})| \leq \alpha_P(d^{\mathfrak{M}_U^0}(\bar{a}, \bar{b}))$, and for any $f \in \mathcal{F}$ with $a(f)$ finite, the same reasoning gives that $d^{\mathfrak{M}_U^0}(f(\bar{a}), f(\bar{b})) \leq \alpha_f(d^{\mathfrak{M}_U^0}(\bar{a}, \bar{b}))$. For P with $a(P)$ infinite, by the uniform continuity of α_P on $[0, \text{db}(\mathcal{L})]$, for any $\varepsilon > 0$ there is a $k < \omega$ such that for any $i \in I$, $\mathfrak{M}_i \models |P(\bar{a}(i)) - P(\bar{b}(i))| \leq \alpha_P(\sup_{j < k} 2^{-j} d(a_j(i), b_j(i))) + \varepsilon$. For this it follows that

$$|P^{\mathfrak{M}_U^0}(\bar{a}) - P^{\mathfrak{M}_U^0}(\bar{b})| \leq \alpha_P(\sup_{j < k} 2^{-j} d^{\mathfrak{M}_U^0}(a_j, b_j)) + \varepsilon \leq \alpha_P(d^{\mathfrak{M}_U^0}(\bar{a}, \bar{b})) + \varepsilon,$$

since α_P is increasing. Since we can do this for any $\varepsilon > 0$, we have that $|P^{\mathfrak{M}_U^0}(\bar{a}) - P^{\mathfrak{M}_U^0}(\bar{b})| \leq \alpha_P(d^{\mathfrak{M}_U^0}(\bar{a}, \bar{b}))$, as required. A similar argument works for f with $a(f)$ infinite. Therefore \mathfrak{M}_U^0 is an \mathcal{L} -pre-structure.

We will prove the second part of the statement by induction on formulas. It is clear by definition that Łoś's theorem holds for any atomic $\mathcal{L}(V)$ -formula φ for any $\{\mathfrak{M}_i, \iota_i\}_{i \in I}$.

We first need to prove that \mathfrak{M}_U^0 is a pre-structure (and therefore that \mathfrak{M}_U is a reduced pre-structure).

If $\varphi = F\bar{\psi}$ and we've shown Łoś's theorem for each of the formulas ψ_k , then it follows for φ by the continuity of the connective F (and the fact that ultrafilter assisted limits commute with continuous functions).

If M_U^0 is empty then the quantifier case follows easily. Otherwise if M_U^0 is non-empty,

then if $\varphi = \inf_x \psi$ and we've shown Loś's theorem for the formula ψ , then for each $\varepsilon > 0$, let a_i be an element of M_i such that $\psi^{\mathfrak{M}_i}(l_i[\frac{a_i}{x}]) \leq \varphi^{\mathfrak{M}_i}(l_i) + \varepsilon$. Now let a be the corresponding element of $M_{\mathcal{U}}^0$. By the induction hypothesis we have that $\psi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}[\frac{a}{x}]) \leq \varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}) + \varepsilon$. Since we can do this for any $\varepsilon > 0$, we have that $\varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}) \leq \lim_{\mathcal{U}} \varphi^{\mathfrak{M}_i}(l_i)$.

Now suppose that $\varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}) < r$ for some r , then there exists a $b \in M_{\mathcal{U}}^0$ such that $\psi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}[\frac{b}{x}]) < r$. By the induction hypothesis this implies that $\{i \in I : \psi^{\mathfrak{M}_i}(l_i[\frac{b(i)}{x}]) < r\} \in \mathcal{U}$, therefore for any $\varepsilon > 0$, $\{i \in I : \varphi^{\mathfrak{M}_i}(l_i) < r + \varepsilon\} \in \mathcal{U}$. Since we can do this for any $r > \varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}})$ and any $\varepsilon > 0$, we have that $\varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}) \geq \lim_{\mathcal{U}} \varphi^{\mathfrak{M}_i}(l_i)$.

Therefore we have that $\varphi^{\mathfrak{M}_{\mathcal{U}}^0}(l_{\mathcal{U}}) = \lim_{\mathcal{U}} \varphi^{\mathfrak{M}_i}(l_i)$, as required.

So by induction Loś's theorem holds for all formulas. \square

Definition 1.6.3. For any \mathcal{L} -pre-structure \mathfrak{M} , there is a canonical map $\mathfrak{M} \rightarrow \mathfrak{M}^{0,\mathcal{U}}$, called the *diagonal embedding*, defined by $f(a)(i) = a$. This extends naturally to a map to $\mathfrak{M}^{\mathcal{U}}$. \triangleleft

Corollary 1.6.4. For any \mathcal{L} -pre-structure \mathfrak{M} , the diagonal embeddings into $\mathfrak{M}^{0,\mathcal{U}}$ and $\mathfrak{M}^{\mathcal{U}}$ are elementary maps. \square

In light of this corollary, if \mathfrak{M} is a reduced pre-structure, we will regard it as an elementary sub-pre-structure of $\mathfrak{M}^{0,\mathcal{U}}$ and $\mathfrak{M}^{\mathcal{U}}$ in this canonical way.

Corollary 1.6.5 (Compactness, Part II). *A type is satisfiable if and only if it is finitely satisfiable. In particular, every complete type is satisfiable.* \square

The following proposition relies on definitions and results from further on, as well as some infinitary model theory that is beyond the purview of this thesis, but given

its limited relevance elsewhere we have included it here in order to keep it close to the definition of ultraproduct. We have also elected to merely sketch the proof.

Proposition 1.6.6. *If $\{\mathfrak{M}_i\}_{i \in I}$ is a family of \mathcal{L} -pre-structures and \mathcal{U} is an ultrafilter on I , then $\mathfrak{M}_{\mathcal{U}}$ is an \mathcal{L} -structure (i.e. is metrically complete) if and only if either*

- \mathcal{U} is not ω_1 -complete (i.e. there exists $\{X_k\}_{k < \omega}$ such that $X_0 \supseteq X_1 \supseteq X_2 \supseteq \dots$ and $X_k \in \mathcal{U}$ for all k but $\bigcap_{k < \omega} X_k = \emptyset$), or
- $\{i \in I : M_i \text{ is a complete metric space}\} \in \mathcal{U}$.

Proof. Since this is really a statement about the metric, we may assume that the signature is countable (empty, in fact).

Either \mathcal{U} is not ω_1 -complete or it is ω_1 -complete. If it is not ω_1 -complete, then $\mathfrak{M}_{\mathcal{U}}$ is \aleph_1 -saturated, since the signature is countable. \aleph_1 -saturated reduced pre-structures are automatically structures, so we have that $\mathfrak{M}_{\mathcal{U}}$ is complete.

If \mathcal{U} is ω_1 -complete, then it preserves discrete $\mathcal{L}_{\omega_1\omega_1}$ -sentences [CK90, Thm. 4.2.11]. If we ‘discretize’ a metric signature by adding a predicate P_r for every $P \in \mathcal{P}$ and $r \in \mathbb{Q}$ (and keeping the same function symbols), with the intent of encoding a metric structure as a discrete structure wherein $\mathfrak{M}_{\text{dis}} \models P_r(\bar{a})$ if and only if $\mathfrak{M} \models P(\bar{a}) \leq r$, and we let this new signature be \mathcal{L}_{dis} , then the class of discrete \mathcal{L}_{dis} -structures arising from the discretization of a reduced \mathcal{L} -pre-structure is axiomatizable by an $\mathcal{L}_{\omega_1\omega}$ -sentence and there is an $\mathcal{L}_{\omega_1\omega_1}$ -sentence φ such that $\mathfrak{M}_{\text{dis}} \models \varphi$ if and only if \mathfrak{M} is metrically complete (where $\mathfrak{M}_{\text{dis}}$ is the discretization of \mathfrak{M}) for any reduced \mathcal{L} -pre-structure \mathfrak{M} .

Finally, definition chasing (and the fact that ω_1 -complete ultrafilters preserve $\mathcal{L}_{\omega_1\omega_1}$ - and therefore $\mathcal{L}_{\omega_1\omega}$ -sentences) shows that this discretization operation commutes with

ultraproducts, as long as the ultrafilter in question is ω_1 -complete, so we have that if \mathcal{U} is ω_1 -complete, then $\mathfrak{M}_{\mathcal{U}}$ is metrically complete if and only if \mathfrak{M}_i is metrically complete on a \mathcal{U} -large set of indices, and the full result follows. \square

1.7 Logical Operations on Topological Formulas

Often we will find that we want to manipulate topological formulas directly, in a manner analogous to the typical logical operations on discrete formulas. We have already seen some of this. We defined the negation of a topological formula in Definition 1.3.9, and we had to compute the intersection of two topological formulas as a subset of type space—which is logically speaking the conjunction of those two formulas—in the proof of Lemma 1.5.7.

Definition 1.7.1. Let φ , ψ , and $\{\chi_i\}_{i<\omega}$ be formulas, and let t and s be terms. Furthermore, let $\lambda = \inf_x \frac{2dxx}{\text{db}(\mathcal{L})} - 1$. Note that $\mathfrak{M} \models \lambda = 1$ if and only if \mathfrak{M} is empty and $\mathfrak{M} \models \lambda = -1$ if and only if \mathfrak{M} is non-empty. For closed formulas we define the following operations:

- $\forall \bar{x}(0 \leq \varphi) = (0 \leq \max\{\inf_{\bar{x}} \varphi, \lambda\})$,
- $(0 \leq \varphi) \wedge (0 \leq \psi) = (0 \leq \min\{\varphi, \psi\})$,
- $\exists \bar{x}(0 \leq \varphi) = (0 \leq \min\{\sup_{\bar{x}} \varphi, -\lambda\})$,¹¹
- $(0 \leq \varphi) \vee (0 \leq \psi) = (0 \leq \max\{\varphi, \psi\})$,
- $(t = s) = (0 \leq -dts)$, and
- $\bigwedge_{i<\omega} (0 \leq \chi_i) = (0 \leq \sum_{i<\omega} [\chi_i]_{-2^{-i}}^0)$.

¹¹This is *weak existential quantification*. The symbol \exists is a slightly enlarged schwa, also known as the letter ‘e’ upside down. Its L^AT_EX definition is `\newcommand{\wexists}{\textnormal{\larger[2]\textschwa}}`. The schwa requires the `tipa` package, and the `\larger` command requires the `relsize` package.

For open formulas we define the following operations:

- $\forall \bar{x}(0 < \varphi) = (0 < \max\{\inf_{\bar{x}} \varphi, \lambda\})$,¹²
- $(0 < \varphi) \wedge (0 < \psi) = (0 < \min\{\varphi, \psi\})$,
- $\exists \bar{x}(0 < \varphi) = (0 < \min\{\sup_{\bar{x}} \varphi, -\lambda\})$,
- $(0 < \varphi) \vee (0 < \psi) = (0 < \max\{\varphi, \psi\})$,
- $(t \neq s) = (0 < dt s)$, and
- $\bigvee_{i < \omega} (0 < \chi_i) = \left(0 < \sum_{i < \omega} [\chi_i]_0^{2^{-i}}\right)$.

We have only stated these definitions for expressions of the form $(0 \leq \varphi)$ and $(0 < \varphi)$.

To extend these to other expressions we use the following rules:

- $(\varphi \leq \psi) \mapsto (0 \leq \psi - \varphi)$,
- $(\varphi < \psi) \mapsto (0 < \psi - \varphi)$,
- $(\varphi \geq \psi) \mapsto (0 \leq \varphi - \psi)$,
- $(\varphi > \psi) \mapsto (0 < \varphi - \psi)$,
- $(\varphi = \psi) \mapsto (0 \leq -|\varphi - \psi|)$, and
- $(\varphi \neq \psi) \mapsto (0 < |\varphi - \psi|)$.

For the sake of restricted open or closed formulas, by $|\varphi - \psi|$ we technically mean $\max\{\varphi + -1 \cdot \psi, \psi + -1 \cdot \varphi\}$. Finally if F and U are a closed formula and an open formula, respectively, we will write $F \rightarrow U$ for $\neg F \vee U$, an open formula, and $U \rightarrow F$ for $\neg U \vee F$, a closed formula. ◁

We only need λ so that quantification has the correct vacuous behavior regardless of $I(\varphi)$; in an empty structure, $\forall x 0 \leq -1$ should be true, regardless of the fact that as we've defined it $0 > \sup_x -1$. Another way to fix this would be to artificially require that $I(\varphi)$ always be a superset of $[-1, 1]$, but that seems far more disruptive than what we've done here. Note that any of the operations in Definition 1.7.1 (other than

¹²This is *strong universal quantification*. The L^AT_EX definitions of the \forall symbol is `\newcommand{\sforall}{\forall\mkern-7.2mu\forall}`.

the infinitary conjunction and disjunction) applied to restricted formulas results in a restricted formulas.

Proposition 1.7.2 (Correctness of Standard Logical Symbols). *All standard logical symbols in Definition 1.7.1 are literally valid. Specifically: Let Q , R , and $\{S_i\}_{i < \omega}$ be conditions (either open or closed).*

- $\mathfrak{M} \models (Q \wedge R)(\iota)$ if and only if $\mathfrak{M} \models Q(\iota)$ and $\mathfrak{M} \models R(\iota)$.
- $\mathfrak{M} \models (Q \vee R)(\iota)$ if and only if $\mathfrak{M} \models Q(\iota)$ or $\mathfrak{M} \models R(\iota)$.
- $\mathfrak{M} \models \bigwedge_{i < \omega} S_i$ if and only if $\mathfrak{M} \models S_i$ for every $i < \omega$.
- $\mathfrak{M} \models \bigvee_{i < \omega} S_i$ if and only if $\mathfrak{M} \models S_i$ for some $i < \omega$.
- $\mathfrak{M} \models \neg Q(\iota)$ if and only if $\mathfrak{M} \not\models Q(\iota)$.
- $\mathfrak{M} \models (Q \rightarrow R)(\iota)$ if and only if whenever $\mathfrak{M} \models Q(\iota)$, $\mathfrak{M} \models R(\iota)$ as well.
- $\mathfrak{M} \models (\exists x Q)(\iota)$ if and only if there is an $a \in M$ such that $\mathfrak{M} \models Q(\iota[\frac{a}{x}])$.
- $\mathfrak{M} \models (\forall x Q)(\iota)$ if and only if for every $a \in M$, $\mathfrak{M} \models Q(\iota[\frac{a}{x}])$.
- $\mathfrak{M} \models (t = s)(\iota)$ if and only if $t^{\mathfrak{M}}(\iota) = s^{\mathfrak{M}}(\iota)$. (\mathfrak{M} a reduced pre-structure.)

Proposition 1.7.3 (Strong Universal and Weak Existential Quantification). *Let Q be a topological formula (either open or closed).*

- $\mathfrak{M} \models (\forall x Q)(\iota)$ if and only if for every $\mathfrak{N} \succeq \mathfrak{M}$ and every $a \in N$, $\mathfrak{N} \models Q(\iota[\frac{a}{x}])$.
- $\mathfrak{M} \models (\exists x Q)(\iota)$ if and only if there is an $\mathfrak{N} \succeq \mathfrak{M}$ and an $a \in N$ such that $\mathfrak{N} \models Q(\iota[\frac{a}{x}])$.

Proof. Compactness. □

Fact 1.7.4 ((Co-)preservation of Open and Closed Sentences). *Let F be a closed \mathcal{L} -sentence, and let U be an open \mathcal{L} -sentence. Let $\{\mathfrak{M}_i\}_{i \in I}$ be a family of \mathcal{L} -structures, and let \mathcal{U} be an ultrafilter on I .*

- U is co-preserved by \mathcal{U} : If $\mathfrak{M}_{\mathcal{U}} \models U$, then $\{i \in I : \mathfrak{M}_i \models U\} \in \mathcal{U}$.
- F is preserved by \mathcal{U} : If $\{i \in I : \mathfrak{M}_i \models F\} \in \mathcal{U}$, then $\mathfrak{M}_{\mathcal{U}} \models F$.

Proposition 1.7.5 (Logical Completeness of Limited Closed Formula). *Call a closed formula limited if it is in the smallest class of closed formulas containing $P\bar{t} \leq r$ and $P\bar{t} \geq r$ for each predicate symbol P , each tuple of terms \bar{t} , and each rational $r \in \mathbb{Q}$, and closed under \wedge , \vee , $\forall x$, and $\exists x$ (but not infinitary conjunction or infinite quantifier strings). Call an open formula limited if it is the negation of a limited closed formula.*

The class of limited closed formulas are logically complete, i.e. for any type $p \in S_V(\Sigma)$, $\{p\} = \bigcap \{\llbracket F \rrbracket : F \in p, F \text{ a limited closed formula}\}$.

Proof. It is sufficient to show that for any p.max.a.n.f. restricted formula φ such that $p \in [0 < \varphi]$, there is a limited closed formula F such that $p \in \text{int}\llbracket F \rrbracket \subseteq \llbracket F \rrbracket \subseteq [0 < \varphi]$ (since limited closed formulas are closed under finite conjunctions and disjunctions). We will prove this by induction on the number of quantifiers.

Let p be a complete type, and assume that $p \in \llbracket 0 < \varphi \rrbracket$ for some quantifier-free φ . Let $\bar{\psi}$ be a list of the atomic formulas occurring in φ , let $C = \prod_{i < |\bar{\psi}|} I(\psi_i)$, let $p_0 \in C$ be the point $(\psi_0(p), \psi_1(p), \dots, \psi_{|\bar{\psi}|-1}(p))$, and let $U = \{\bar{r} \in C : \varphi[\frac{\bar{r}}{\bar{\psi}}] > 0\}$, where $\varphi[\frac{\bar{r}}{\bar{\psi}}]$ is φ , with each instance of ψ_i replaced with r_i evaluated as a real number. U is an open neighborhood of p_0 , so there exists a sequence $\{a_i, b_i\}_{i < |\bar{\psi}|}$ of pairs of rational numbers

with $a_i < b_i$ such that $p \in \prod_{i < |\bar{\psi}|} (a_i, b_i) \subset \prod_{i < |\bar{\psi}|} [a_i, b_i] \subseteq U$. Now we clearly have that $\bigwedge_{i < |\bar{\psi}|} (a_i \leq \psi) \wedge (\psi \leq b_i)$ is the required closed formula.

Now assume we've shown the statement for all p and φ with n quantifiers.

Let $\chi = \sup_x \varphi$ be a real $\mathcal{L}(V)$ -formula, and assume that $q \in \llbracket 0 < \chi \rrbracket$. Definition chasing shows that a type q is in $\llbracket 0 < \chi \rrbracket$ if and only if there is a type $p(Vx)$ such that $p \in \llbracket 0 < \varphi \rrbracket$ and q is the image of p under the projection map $\pi : S_{Vx} \rightarrow S_V$. So let F be a closed $\mathcal{L}(Vx)$ -formula such that $p \in \text{int} \llbracket F \rrbracket \subseteq \llbracket F \rrbracket \subseteq \llbracket 0 < \varphi \rrbracket$. Now we have that $q \in \pi(\text{int} \llbracket F \rrbracket) \subseteq \text{int} \llbracket \exists x F \rrbracket \subseteq \llbracket \exists x F \rrbracket \subseteq \llbracket 0 < \varphi \rrbracket$ (because the projection map is open).

Let $\chi = \inf_x \varphi$ be a real $\mathcal{L}(V)$ -formula, and assume that $q \in \llbracket 0 < \chi \rrbracket$. Definition chasing and compactness show that a type q is in $\llbracket 0 < \chi \rrbracket$ if and only if for every $p(Vx) \in \pi^{-1}(q)$, $p \in \llbracket 0 < \varphi \rrbracket$. By the induction hypothesis, for each $p \in \pi^{-1}(q)$, we can find a limited closed formula F_p such that $p \in \text{int} \llbracket F_p \rrbracket \subseteq \llbracket F_p \rrbracket \subseteq \llbracket 0 < \varphi \rrbracket$. $\pi^{-1}(q)$ is a closed set, so by compactness there is a finite set $P \subseteq \pi^{-1}(q)$ such that $\bigcup_{p \in P} \text{int} \llbracket F_p \rrbracket \supseteq \pi^{-1}(q)$. Therefore we have that $q \in \text{int} \llbracket \forall x \bigvee_{p \in P} F_p \rrbracket \subseteq \llbracket \forall x \bigvee_{p \in P} F_p \rrbracket \subseteq \llbracket 0 < \chi \rrbracket$.

So by induction the result holds for all p and φ , and limited closed formulas are logically complete. \square

Remark 1.7.6. In light of Proposition 1.7.5, it might be tempting to try to formalize continuous logic in terms of (limited) open and closed formulas directly (similarly to Iovino's positive bound formulas for Banach spaces [Iov99], a precursor of continuous logic). While it is possible, formalizing \mathbb{W} and \exists directly is unpleasant, and real valued formulas, while less intuitive, are fundamental in continuous logic (in particular it is difficult to formalize definable sets and definable (partial) functions in terms of open

and closed formulas directly). Furthermore, it is useful to have both open and closed formulas and real valued formulas at our disposal; a good example of this is Proposition 2.3.19, which is most naturally stated using both kinds of formulas.

For an explicit example of this, say that a real formula $\varphi(\bar{x}, y)$ *defines a singleton at \bar{a}* if there is a unique b such that $\varphi(\bar{a}, b) = 0$ and $d(b, y) \leq \varphi(\bar{a}, y)$ (in any model containing \bar{a}), and say that a closed formula $F(\bar{x}, y)$ *defines a singleton at \bar{a}* if there is a unique b such that $F(\bar{a}, b)$ (in any model containing \bar{a}). Just as in discrete logic, any (real or closed) formula manages to define a partial function on some (possibly empty) set of inputs. The difficulty is with the complexity of the set of inputs for which the output is defined. “ $\varphi(\bar{x}, y)$ defines a singleton at \bar{x} ” is equivalent to a certain closed formula:

$$\exists y(\varphi(\bar{x}, y) = 0) \wedge \forall z(d(y, z) \leq \varphi(\bar{x}, z)). \quad (*)$$

(Although a proof is required to show that that y actually always exists.) On the other hand, “ $F(\bar{x}, y)$ defines a singleton at \bar{x} ” is not equivalent to an open or closed formula. Attempting to write it out literally gives something like:

$$\exists y F(\bar{x}, y) \wedge \forall z(F(\bar{x}, z) \rightarrow y = z).$$

$(F(\bar{x}, z) \rightarrow y = z)$ is neither a closed nor an open formula, but rather the union of a closed and open formula. A priori, the set of types p such that if $\bar{a} \models p$, then $F(\bar{a}, y)$ defines a singleton could fail to even be Borel. This is a good example of how open and closed formulas are useful but also slightly dangerous. Thinking in terms of them recovers much of the intuitive expressiveness of discrete first-order logic; the closed formula (*)

is easier for a typical logician to read and comprehend than

$$\inf_y \max\{\varphi(\bar{x}, y), \sup_z dyz \div \varphi(\bar{x}, z)\} = 0.^{13}$$

But open and closed formulas do not and cannot fully recover discrete expressiveness. Continuous logic is not discrete logic and many things one would like to say cannot be said.

Finally, we will occasionally want to allow ourselves unrestricted logical operations on formulas (we have already done this implicitly in Remark 1.7.6). Equivalently, we will want to be able to think of arbitrary sets of types as a kind of formula.

Definition 1.7.7. For any signature \mathcal{L} and set of variable symbols V , the class of *type-set $\mathcal{L}(V)$ -formulas* is the smallest class containing topological $\mathcal{L}(V)$ -formulas and closed under the following:

- If $\{X_i\}_{i \in I}$ is a set of type-set $\mathcal{L}(V)$ -formulas, then $\bigwedge_{i \in I} X_i$ is a type-set $\mathcal{L}(V)$ -formula.
- If X is a type-set $\mathcal{L}(V)$ -formula, then $\neg X$ is a type-set $\mathcal{L}(V)$ -formula.
- If \bar{v} is a tuple of distinct variables and X is a type-set $\mathcal{L}(V\bar{v})$ -formula, then $\exists \bar{v} X$ is a type-set $\mathcal{L}(V)$ -formulas.

¹³Although some of this is down to notation. Imagine if discrete logic was written like

$$\text{exi}_y \text{and} \{\varphi(\bar{x}, y), \text{all}_z e(y, z) \leftarrow \varphi(\bar{x}, z)\},$$

and somehow the symbol \leftarrow was symmetric, so that nothing like \rightarrow could be written in an immediately intuitive way.

$\bigvee_{i \in I} X_i$ is defined as $\neg \bigwedge_{i \in I} \neg X_i$, $X \rightarrow Y$ is defined as $\neg X \vee Y$, and $\forall \bar{v} X$ is defined as $\neg \exists \bar{v} \neg X$. Satisfaction is given by the following:

- $\mathfrak{M} \models (\bigwedge_{i \in I} X_i)(\iota)$ if and only if $\mathfrak{M} \models X_i(\iota)$ for each $i \in I$,
- $\mathfrak{M} \models \neg X(\iota)$ if and only if $\mathfrak{M} \not\models X(\iota)$, and
- $\mathfrak{M} \models (\exists \bar{v} X)(\iota)$ if and only if there is $\mathfrak{N} \succeq \mathfrak{M}$ and some tuple $\bar{a} \in N$ of the same length as \bar{v} such that $\mathfrak{N} \models X(\iota[\frac{\bar{a}}{\bar{v}}])$.

Free variables are defined in the obvious way.

Open type-set formulas are type-set formulas generated from open formulas using finite conjunctions, arbitrary disjunctions, \exists , and \forall . *Closed type-set formulas* are type-set formulas generated from closed formulas using arbitrary conjunctions, finite disjunctions, \exists , and \forall .

If X and Y are type-set formulas, we write $X \models Y$ to mean $\mathfrak{M} \models Y(\iota)$ for any (\mathfrak{M}, ι) such that $\mathfrak{M} \models X(\iota)$. If T is a (possibly incomplete) theory, we say X and Y are *logically equivalent modulo T* if $T, X \models Y$ and $T, Y \models X$. We say that X and Y are *logically equivalent* if they are logically equivalent modulo \emptyset .

If $X(\bar{x}, \bar{y})$ is a type-set formula, \mathfrak{M} is a structure, and $\bar{a} \in M$ is a tuple assigned to \bar{y} , then we write $X(\mathfrak{M}, \bar{y})$ for the set of tuples $\bar{b} \in \mathfrak{M}$ such that $\mathfrak{M} \models X(\bar{b}, \bar{a})$.

If Σ is a partial $\mathcal{L}(V)$ -type, we regard it as the type-set formula $\bigwedge_{F \in \Sigma} F$. If A is a set of partial types (typically a set of complete types in some particular type space), we regard it as the type-set formula $\bigvee_{\Sigma \in A} \Sigma$. Finally we set $\llbracket X \rrbracket_{\Sigma, \mathcal{L}, V} = \{p \in S_V(\Sigma, \mathcal{L}) : p \models X\}$. \triangleleft

Its not hard to show that $\mathfrak{M} \models X(\iota)$ if and only if $\text{tp}_{\mathfrak{M}}(\iota) \in \llbracket X \rrbracket$ and that $\llbracket X \rrbracket$ interpreted as a type-set formula is logically equivalent to X .

Remark 1.7.8. There is an awkward terminological issue which only becomes relevant in uncountable signatures. Open and closed formulas do not correspond to open and closed sets of types. They actually correspond to open F_σ and closed G_δ sets of types, respectively. Open and closed type-set formulas correspond to open and closed sets of types. Nevertheless, open and closed formulas are the more useful notions, so I have decided to reserve the shorter terms for them, despite the mismatch between topological and formulaic terminology.

Proposition 1.7.9. *If U is an open type-set formula, then it is logically equivalent to one of the form $\bigvee_{i \in I} V_i$, with V_i open formulas. In particular, $\llbracket U \rrbracket$ is a topologically open set.*

If F is a closed type-set formula, then it is logically equivalent to one of the form $\bigwedge_{i \in I} G_i$, with G_i closed formulas. In particular, $\llbracket F \rrbracket$ is a topologically closed set.

Proof. The second part clearly follows from the first. We will prove the first part by induction.

Let $\{U_i\}_{i \in I}$ be a family of open type-set formulas each of which is equivalent to $\bigvee_{j \in J_i} V_j^i$, with V_j^i open formulas. $\bigvee_{i \in I} U_i$ is clearly logically equivalent to a type-set formula of the required form.

Assume that $I = \{0, \dots, n\}$ is finite. Then $\bigwedge_{i \in I}$ is logically equivalent to

$$\bigvee_{j_0 \in J_0, \dots, j_n \in J_n} V_{j_0}^0 \wedge \dots \wedge V_{j_n}^n.$$

Let U be an open type-set formula logically equivalent to $\bigvee_{i \in I} V_i$, with V_i open

formulas. $\exists \bar{v}U$ is logically equivalent to $\bigvee_{i \in I} \exists \bar{v}_i V_i$, where $\bar{v}_i = \bar{v} \upharpoonright \text{fv}(V_i)$. By Proposition 1.7.2 and elementarity, this is logically equivalent to $\bigvee_{i \in I} \exists \bar{v}_i V_i$.

Now assume that $U = \bigvee_{i \in I} V_i$ is a $\mathcal{L}(\bar{x}\bar{v})$ -formula and that for some \mathfrak{M} and some ι , a \bar{x} -assignment with $\text{tp}(\iota) = p(\bar{x})$, $\mathfrak{M} \models (\forall \bar{v}U)(\iota)$.

By compactness there is a closed formula $F_p(\bar{x}\bar{v})$ and an open formula $W_p(\bar{x}\bar{v})$ such that $p(\bar{x}) \models W_p(\bar{x}\bar{v}) \models F_p(\bar{x}\bar{v}) \models U(\bar{x}\bar{v})$ and for any $\mathfrak{N} \succeq \mathfrak{M}$ and any $\bar{a} \in N$, $\mathfrak{N} \models W_p(\iota[\frac{\bar{a}}{\bar{v}}])$. Now consider

$$O(\bar{x}) = \bigvee_{p(\bar{x}) \in \llbracket \forall \bar{v}U \rrbracket_{\bar{x}}} V_{i_p} \wedge \forall \bar{v}_p W_p,$$

where i_p is chosen so that $p \in \llbracket V_{i_p} \rrbracket$ and $\bar{v}_p = \bar{v} \upharpoonright \text{fv}(V_p)$. By construction $O \models U$, and we have just established that $\forall \bar{v}U \models W$, so we have that U is logically equivalent to a type-set formula of the required form. \square

Corollary 1.7.10. *If U is an open type-set formula, then $\mathfrak{M} \models (\exists \bar{v}U)(\iota)$ if and only if there exists $\bar{a} \in M$ of the same length as \bar{v} such that $\mathfrak{M} \models U(\iota[\frac{\bar{a}}{\bar{v}}])$.*

If F is a closed type-set formula, then $\mathfrak{M} \models (\forall \bar{v}U)(\iota)$ if and only if for every $\bar{a} \in M$ of the same length as \bar{v} , $\mathfrak{M} \models F(\iota[\frac{\bar{a}}{\bar{v}}])$.

Proof. Closed type-set formulas are clearly logically equivalent to negations of open type-set formulas, so by duality it is enough to show the first statement. Let U be logically equivalent to $\bigvee_{i \in I} V_i$, where V_i are open formulas. Then $\mathfrak{M} \models \exists \bar{v}U$ if and only if $\mathfrak{M} \models \exists \bar{v}V_i$ for some $i \in I$. This clearly only depends on at most countably many variables in \bar{v} , so if we let \bar{v}_0 be the variables actually occurring in V_i , then we have that $\mathfrak{M} \models \exists \bar{v}_0 V_i$, so we can set the other elements of \bar{v} arbitrarily and we have a witness in \mathfrak{M} . \square

In light of this corollary we may use notation such as $\forall xF(x, y)$ and $\exists xU(x, y)$ for closed and open type-set formulas.

1.8 Saturated, Special, and Monster Models

Just as in discrete logic, sufficiently saturated and special models can be used as a shortcut in many proofs. As such, we will introduce the concepts here, in anticipation of some of the proofs in the next chapter. Our treatment of saturated and special models follows [Hod93] closely.

Definition 1.8.1. • For any cardinal κ , a pre-structure \mathfrak{M} is κ -saturated if for any set $A \subseteq M$ with $|A| < \kappa$, for every type $p(x, A) \in S_x(A)$, p is realized in \mathfrak{M} (i.e. there is $b \in M$ such that $\mathfrak{M} \models p(b, A)$).

- A pre-structure \mathfrak{M} is *saturated* if it is $\aleph^{\text{dc}}\mathfrak{M}$ -saturated.
- A pre-structure \mathfrak{M} is κ -saturated over a dense sub-pre-structure if there is a dense sub-pre-structure $\mathfrak{M}_0 \subseteq \mathfrak{M}$ such that for any set $A \subseteq M_0$ with $|A| < \kappa$, for every type $p(x, A) \in S_x(A)$, there is $b \in M$ such that $\mathfrak{M} \models p(b, A)$.
- A pre-structure \mathfrak{M} is *approximately ω -saturated* if for any finite $\bar{a} \in \mathfrak{M}$, any type $p(x, \bar{a}) \in S_x(\bar{a})$, and any $\varepsilon > 0$, there is $\bar{b}, c \in \mathfrak{M}$ such that $d(\bar{a}, \bar{b}) < \varepsilon$, $\bar{a} \equiv \bar{b}$, and $\mathfrak{M} \models p(c, \bar{b})$.
- A pre-structure \mathfrak{M} is *strongly κ -homogeneous* if for any $\bar{a}, \bar{b} \in \mathfrak{M}$ with $|\bar{a}| = |\bar{b}| < \kappa$ and $\bar{a} \equiv \bar{b}$, there is an automorphism σ of \mathfrak{M} such that $\bar{a} = \sigma(\bar{b})$.

- A *specializing chain* for a structure \mathfrak{M} of density character κ is an elementary chain $\{\mathfrak{M}_\lambda\}_{\lambda < \kappa}$ (indexed by cardinals) such that $\mathfrak{M} = \overline{\bigcup_{\lambda < \kappa} \mathfrak{M}_\lambda}$ and for each $\lambda < \kappa$, \mathfrak{M}_λ is λ^+ -saturated. \mathfrak{M} is *special* if there is a specializing chain for \mathfrak{M} . \triangleleft

One immediate use of sufficiently saturated models is computing quantifiers of the form \forall and \exists without passing to an elementary expansion.

Proposition 1.8.2 (Strong Universal and Weak Existential Quantification in Sufficiently Saturated Models). *Let \mathfrak{M} be an \mathcal{L} -structure that is \aleph_1 -saturated. Let Q be a topological formula (either open or closed), possibly with parameters (note that by definition topological formulas may only have at most countably many parameters).*

(i) $\mathfrak{M} \models (\forall x Q)(\iota)$ if and only if for every $a \in M$, $\mathfrak{M} \models Q(\iota[\frac{a}{x}])$.

(ii) $\mathfrak{M} \models (\exists x Q)(\iota)$ if and only if there is an $a \in M$ such that $\mathfrak{M} \models Q(\iota[\frac{a}{x}])$.

More generally, the same holds for Q an arbitrary type-set formula, as long as \mathfrak{M} is κ^+ -saturated, where Q has at most κ many parameters.

Proof. (i) This follows from dualizing (ii).

(ii) The \Leftarrow direction is immediate. Suppose that $\mathfrak{M} \models (\exists x Q)(\iota)$, then there is some $\mathfrak{N} \succeq \mathfrak{M}$ and $a \in N$ such that $\mathfrak{N} \models Q(\iota[\frac{a}{x}])$. The type of a over the parameters in Q is realized by $b \in M$ by \aleph_1 -saturation, so $\mathfrak{M} \models Q(\iota[\frac{b}{x}])$.

The proof of the final statement is the same. \square

Now we will see that the concept of κ -saturation over a dense sub-pre-structure is equivalent to κ -saturation when $\kappa \geq \aleph_1$.

Proposition 1.8.3. *For any $\kappa \geq \aleph_1$, if \mathfrak{M} is κ -saturated over a dense sub-pre-structure, then it is κ -saturated.*

Proof. Let the approximate κ -saturation of \mathfrak{M} be witnessed by \mathfrak{M}_0 . Let $A \subset M$ be a subset with $|A| < \kappa$. For each $a \in A$, find $\{b_a^i\}_{i < \omega} \subseteq M_0$ such that $d^{\mathfrak{M}}(a, b_a^i) \leq 2^{-i}$. Now let $B = \{b_a^i : a \in A, i < \omega\}$. Note that since $\kappa \geq \aleph_0$, $|B| < \kappa$ still holds. By assumption, \mathfrak{M} realizes every type in $S_1(B)$. Let p be a type in $S_1(A)$. Construct a partial type Σ as follows: for each restricted formula $\varphi(x, \bar{a})$ such that $p(x, A) \models \varphi(x, \bar{a}) \leq 0$ (note that closed formulas of this form completely determine p), let $\varphi(x, b_{a_0}^i, \dots, b_{a_{|a|-1}}^i) \leq \alpha_\varphi(2^{-i})$ be in Σ . It's not hard to show that $\Sigma(x, B) \models p(x, A)$ and that $\Sigma(x, B)$ is consistent. Let $q(x, B)$ be a completion of $\Sigma(x, B)$. By assumption, there is $c \in \mathfrak{M}$ such that $\mathfrak{M} \models q(c, B)$. Therefore $\mathfrak{M} \models p$. \square

Corollary 1.8.4. *For any $\lambda \geq \aleph_0$, \mathfrak{M} fails to be λ -saturated if and only if there is some $A \subseteq \mathfrak{M}$ with $|A| \leq \lambda$ and a type $p \in S_1(A)$ such that for some $\varepsilon > 0$, \mathfrak{M} ($> \varepsilon$)-omits p (i.e. for every q realized in \mathfrak{M} , $d(p, q) > \varepsilon$).*

Remark 1.8.5. Note that Proposition 1.8.3 is false in general for $\kappa = \omega$. A simple example is the theory of ω^ω with the string metric (i.e. the distance between two points is 2^{-n} where n is the index of first disagreement) together with a function f which takes $a(i)$ to $a(i+1)$. This theory is ω -categorical, but for any $a(i)$ such that $a(i) \neq a(j)$ for all $i \neq j$,¹⁴ there is a type for some $b(i)$ such that $a(i) \neq b(j)$ for all i and j (i.e. a and b have disjoint image), but this type is not realized over $a(i) = i$. Note also that this is an example of an ω -categorical theory which fails to be ω -categorical after the addition of a constant (specifically a constant selecting out $a(i)$).

¹⁴This property is expressible in continuous first-order logic.

In Section 2.5, we will show that for $\kappa = \omega$, κ -saturation over a dense sub-pre-structure implies approximate ω -saturation. For now, we will collect a couple more facts regarding κ -saturation.

Proposition 1.8.6. *Let \mathcal{L} be a signature.*

(i) *If a pre-structure \mathfrak{M} is κ -saturated for $\kappa \geq \aleph_1$, then ${}^r\mathfrak{M}$ is κ -saturated and $\overline{\mathfrak{M}}$ is κ -saturated over a dense sub-pre-structure.*

(ii) *A structure \mathfrak{M} is κ -saturated if and only if for every $A \subseteq M$ with $|A| < \kappa$, the set $\{p \in S_1(A) : p \text{ realized in } \mathfrak{M}\}$ is metrically dense in $S_1(A)$.*

Proof. (i) This follows immediately from the fact that the natural quotient map $\mathfrak{M} \rightarrow {}^r\mathfrak{M}$ and the natural inclusion map ${}^r\mathfrak{M} \rightarrow \overline{\mathfrak{M}}$ are both elementary maps, as well as the fact that the image of \mathfrak{M} in $\overline{\mathfrak{M}}$ is dense.

(ii) Fix $A \subseteq M$ with $|A| < \kappa$. Let $p \in S_1(A)$ be some type. Let $p_0 = p$. Find type $q_0 \in S_1(A)$ such that $d(p_0, q_0) < 2^{-0}$ and such that q_0 is realized by $b_0 \in M$. Find $p_1 \in S_1(Ab_0)$ extending p_0 such that $p_1 \in \llbracket dx b_0 < 2^{-0} \rrbracket$.

For any $0 < i < \omega$, given $p_i \in S_1(Ab_0 \dots b_{i-1})$ with $p_i \in \llbracket dx b_i < 2^{-i} \rrbracket$, find a type $q_i \in S_1(Ab_0 \dots b_{i-1})$ such that $d(p_i, q_i) < 2^{-i}$ and such that q_i is realized by b_i in \mathfrak{M} . Find $p_{i+1} \in S_1(Ab_0 \dots b_i)$ extending p_i such that $p_{i+1} \in \llbracket dx b_{i+1} < 2^{-i-1} \rrbracket$.

Note that by construction we have that $d(b_i, b_{i+1}) < 2^{-i} + 2^{-i-1}$, so $\{b_i\}_{i < \omega}$ is a Cauchy sequence with limit $b_\omega \in M$ (since \mathfrak{M} is a structure and not a pre-structure). We also have that $\text{tp}(b_i/A) \rightarrow p$ as $i \rightarrow \infty$, so b_ω realizes p in \mathfrak{M} , as required. \square

The following proposition is largely a direct generalization of the analogous facts in discrete logic. The most notable change is the lower bound on the cofinality in parts (v) and (viii).

Proposition 1.8.7. *Let \mathcal{L} be a signature.*

- (i) *If \mathfrak{M} is a compact structure, then it is κ -saturated for every κ .*
- (ii) *If \mathfrak{M} is a non-compact κ -saturated structure, then $\#^{\text{dc}}\mathfrak{M} \geq \kappa$.*
- (iii) *For any \mathcal{L} -pre-structure \mathfrak{M} , if $\kappa > \#^{\text{dc}}\mathfrak{M} + |\mathcal{L}|$ is strongly inaccessible, then \mathfrak{M} has a saturated elementary extension of density character κ .*
- (iv) *If \mathfrak{M} is κ -saturated then for every $A \subseteq M$ with $\#^{\text{dc}}A < \kappa$ and every $\lambda < \kappa$, \mathfrak{M} realizes every type in $S_\lambda(A)$.*
- (v) *If κ is a regular cardinal with $\kappa > |\mathcal{L}|$ and $\text{cf}(\kappa) > \aleph_0$, then for any \mathcal{L} -structure \mathfrak{M} , there is a κ -saturated elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ with $\#^{\text{dc}}\mathfrak{N} \leq (\#^{\text{dc}}\mathfrak{M})^{<\kappa}$ (where $\lambda^{<\delta} = \sum_{\gamma < \delta} \lambda^\gamma$).*
- (vi) *For every non-compact \mathcal{L} -structure \mathfrak{M} and every strong limit cardinal¹⁵ $\kappa > \#^{\text{dc}}\mathfrak{M} + |\mathcal{L}|$, \mathfrak{M} has a special elementary extension \mathfrak{N} of density character κ . If $\text{cf}(\kappa) > \aleph_0$, then $|\mathfrak{N}| = \kappa$ as well.*
- (vii) *If \mathfrak{M} is special with specializing chain $\{\mathfrak{M}_\lambda\}_{\lambda < \kappa}$ and $\bar{a} \in \bigcup_{\lambda < \kappa} \mathfrak{M}_\lambda$ is a tuple of parameters with $|\bar{a}| < \text{cf}(\kappa)$, then (\mathfrak{M}, \bar{a}) is special as well.*
- (viii) *If \mathfrak{M} is special with $\text{cf}(\#^{\text{dc}}\mathfrak{M}) > \aleph_0$, then for any $\bar{a} \in \mathfrak{M}$ with $|\bar{a}| < \text{cf}(\kappa)$, (\mathfrak{M}, \bar{a}) is special as well.*
- (ix) *If \mathfrak{M} is saturated, then it is special.*
- (x) *If two structures $\mathfrak{M}, \mathfrak{N}$ have $\mathfrak{M} \equiv \mathfrak{N}$, $\#^{\text{dc}}\mathfrak{M} = \#^{\text{dc}}\mathfrak{N}$, and are both special, then there is an isomorphism $\mathfrak{M} \cong \mathfrak{N}$.*

¹⁵For every $\lambda < \kappa$, $2^\lambda < \kappa$.

(xi) If \mathfrak{M} is special with $\text{cf}(\#\text{dc}\mathfrak{M}) > \aleph_0$, then it is strongly $\text{cf}(\#\text{dc}\mathfrak{M})$ -homogeneous.

Proof. (i) \mathfrak{M} (regarded as a set of types) is topologically dense as a subset of $S_1(\mathfrak{M})$. Since it is metrically compact, it is topologically closed, therefore $\mathfrak{M} = S_1(\mathfrak{M})$ and \mathfrak{M} realizes every type over any subset of itself.

(ii) Since \mathfrak{M} is non-compact, there is some $\varepsilon > 0$ such that for any set $A \subseteq \mathfrak{M}$, the partial type $\{dxa \geq \varepsilon\}_{a \in A}$ is consistent. By induction, this implies that \mathfrak{M} has a $(\geq \varepsilon)$ -separated set of cardinality κ , therefore $\#\text{dc}\mathfrak{M} \geq \kappa$.

(iii) The proof is the same as it is in discrete logic.

(iv) Let p be a type in $S_\lambda(A)$. Let $A_0 \subseteq A$ be metrically dense with $|A_0| < \kappa$.

Claim. The reduct map $r : S_\lambda(A) \rightarrow S_\lambda(A_0)$ is a homeomorphism.

Proof of claim. Let $p, q \in S_\lambda(A)$ be distinct types. Since they are distinct, there is a restricted formula $\varphi(\bar{x}, \bar{a})$ such that $p \models \varphi(\bar{x}, \bar{a}) = 0$ and $q \models \varphi(\bar{x}, \bar{a}) = 1$. We can find $\bar{a}_0 \in A_0$ close enough to \bar{a} such that $p \models \varphi(\bar{x}, \bar{a}_0) \leq \frac{1}{3}$ and $q \models \varphi(\bar{x}, \bar{a}_0) \geq \frac{2}{3}$. Therefore $r(p)$ and $r(q)$ are distinct, so r is a continuous bijection between compact Hausdorff spaces and is therefore a homeomorphism (see Fact A.2.11). \square_{claim}

So for any $p \in S_\lambda(A)$, $q = p \upharpoonright A_0 \models p$, so it is enough to show that q is realized in \mathfrak{M} . For each $i < \lambda$, let $q \upharpoonright i$ be q restricted to the first i variables. Now the proof is as it is in discrete logic, i.e. find $b_0 \models q_0(x)$, find $b_1 \models q_1(b_0, x)$, etc. Then the sequence $\{b_i\}_{i < \lambda}$ realizes q and therefore realizes p .

(v) If \mathfrak{M} is compact then this is trivial, so assume that \mathfrak{M} is non-compact.

Let $\mathfrak{M}_0 = \mathfrak{M}$. Clearly $\#\text{dc}\mathfrak{M}_0 \leq (\#\text{dc}\mathfrak{M})^{\aleph_0}$.

For ordinal $i < \kappa$, given \mathfrak{M}_i , let \mathfrak{M}_i^0 be a dense sub-pre-structure of \mathfrak{M}_i such that $|\mathfrak{M}_i^0| = \#\text{dc}\mathfrak{M}_i \leq (\#\text{dc}\mathfrak{M})^{\aleph_0 + |i| + |\mathcal{L}|}$. Let \mathfrak{M}_{i+1} be an elementary extension of \mathfrak{M}_i such that for every $A \subseteq \mathfrak{M}_i^0$ with $|A| \leq \aleph_0 + |i + 1|$ (where $|i|$ is the cardinality of i as an

ordinal), \mathfrak{M}_{i+1} realizes every type in $S_1(A)$. $|S_1(A)|$ is at most $2^{\aleph_0+|A|+|\mathcal{L}|}$, so by the Löwenheim-Skolem theorem (Proposition B.1.2) we can guarantee that

$$\begin{aligned} \#^{\text{dc}}\mathfrak{M}_{i+1} &\leq \#^{\text{dc}}\mathfrak{M}_i + |\mathcal{L}| + (\#^{\text{dc}}\mathfrak{M}_i)^{\aleph_0+|i+1|} \cdot 2^{\aleph_0+|i+1|+|\mathcal{L}|} \\ &\leq (\#^{\text{dc}}\mathfrak{M})^{\aleph_0+|i+1|+|\mathcal{L}|}. \end{aligned}$$

For each limit ordinal $i \leq \kappa$, let $\mathfrak{M}_i = \overline{\bigcup_{j < i} \mathfrak{M}_j}$. By induction we have that

$$\begin{aligned} \#^{\text{dc}}\mathfrak{M}_i &\leq \sum_{j < i} (\#^{\text{dc}}\mathfrak{M})^{\aleph_0+|j|+|\mathcal{L}|} \\ &\leq (\aleph_0 + |i|) \cdot (\#^{\text{dc}}\mathfrak{M})^{\aleph_0+|i|+|\mathcal{L}|} = (\#^{\text{dc}}\mathfrak{M})^{\aleph_0+|i|+|\mathcal{L}|}. \end{aligned}$$

Finally, let $\mathfrak{N} = \mathfrak{M}_\kappa$. Since $\text{cf}(\kappa) > \aleph_0$, \mathfrak{N} is automatically complete. Furthermore, we have that

$$\begin{aligned} \#^{\text{dc}}\mathfrak{N} &\leq \sum_{i < \kappa} (\#^{\text{dc}}\mathfrak{M})^{\aleph_0+|i|+|\mathcal{L}|} \\ &= \sum_{\lambda < \kappa} \lambda \cdot (\#^{\text{dc}}\mathfrak{M})^{\lambda+|\mathcal{L}|} \\ &= \sum_{\lambda < \kappa} (\#^{\text{dc}}\mathfrak{M})^\lambda = (\#^{\text{dc}}\mathfrak{M})^{<\kappa}, \end{aligned}$$

where λ ranges over infinite cardinals.

Let $A \subset \mathfrak{N}$ have $|A| < \kappa$. Since κ is a regular cardinal, $A \subseteq \mathfrak{M}_i$ for some $i < \kappa$. Therefore for any $p \in S_1(A)$, there is $b \in \mathfrak{M}_{i+1}$ such that $\mathfrak{M}_{i+1} \models p(b, A)$, so p is realized in \mathfrak{N} , and \mathfrak{N} is κ -saturated.

(vi) We will need the follow cardinal arithmetic fact from [Hod93, Fact 10.4.1]. If

α is a limit ordinal, then there is a strictly increasing sequence $\{\mu_i\}_{i < \text{cf}(\alpha)}$ of regular cardinals such that $\beth_\alpha = \sum_{i < \text{cf}(\alpha)} \mu_i = \sum_{i < \text{cf}(\alpha)} 2^{\mu_i}$. Also note that $\text{cf}(\alpha) = \text{cf}(\beth_\alpha)$. Let $\{\mu_i\}_{i < \text{cf}(\kappa)}$ be such a sequence.

Let $\delta_0 = \#^{\text{dc}}\mathfrak{M} + |\mathcal{L}| + \mu_0$. Let \mathfrak{N}_0 be an elementary extension of \mathfrak{M} with $\#^{\text{dc}}\mathfrak{N}_0 = \delta_0$. Then, since $\text{cf}(\delta_0^+) = \delta_0^+ > \aleph_0$, we can use part (v) to get \mathfrak{M}_0 , a δ_0^+ -saturated elementary extension of \mathfrak{N}_0 with $\#^{\text{dc}}\mathfrak{M}_0 \leq (\#^{\text{dc}}\mathfrak{N}_0)^{<\delta_0^+} = \delta_0^{\delta_0} = 2^{\delta_0}$. Note that since κ is a strong limit cardinal, $2^{\delta_0} < \kappa$.

Now for each $i < \text{cf}(\kappa)$, given \mathfrak{M}_i with $\#^{\text{dc}}\mathfrak{M}_i \leq 2^{\delta_i}$, let $\delta_{i+1} = 2^{\delta_i} + \mu_i + 1$, and let \mathfrak{N}_{i+1} be an elementary extension of \mathfrak{M}_i with $\#^{\text{dc}}\mathfrak{N}_{i+1} = \delta_{i+1}$, and, again since $\text{cf}(\delta_{i+1}^+) = \delta_{i+1}^+ > \aleph_0$, we can let \mathfrak{M}_{i+1} be a δ_{i+1}^+ -saturated elementary extension of \mathfrak{N}_{i+1} with $\#^{\text{dc}}\mathfrak{M}_{i+1} \leq \delta_{i+1}^{<\delta_{i+1}^+} = 2^{\delta_{i+1}}$. Note that κ is a strong limit cardinal, $2^{\delta_{i+1}} < \kappa$.

For each limit ordinal $i \leq \text{cf}(\kappa)$, let $\mathfrak{M}_i = \bigcup_{j < i} \mathfrak{M}_j$, and let $\delta_i = \mu_i + \sum_{j < i} \delta_j$. Note that $\#^{\text{dc}}\mathfrak{M}_i \leq \delta_i < 2^{\delta_i}$. Note that since $i < \text{cf}(\kappa)$, $2^{\delta_i} < \kappa$.

Finally let $\mathfrak{N} = \mathfrak{M}_{\text{cf}(\kappa)}$. By construction we have that $\#^{\text{dc}}\mathfrak{M}_{\text{cf}(\kappa)} = \kappa$. If $\text{cf}(\kappa) > \aleph_0$, then by Fact A.1.5, $|\mathfrak{N}| = \kappa$ as well.

To see that \mathfrak{N} is special, for each infinite cardinal $\lambda < \kappa$, let $i(\lambda)$ be the smallest i such that $\delta_i \geq \lambda$. Now we have that $\{\mathfrak{M}_{i(\lambda)}\}_{\lambda < \kappa}$ is a specializing chain for \mathfrak{N} , since $\mathfrak{M}_{i(\lambda)}$ is $\delta_{i(\lambda)}^+$ -saturated, and therefore also λ^+ -saturated.

(vii) This follows easily from the definition of special and the fact that $\bar{a} \in \mathfrak{M}_\lambda$ for some $\lambda < \kappa$.

(viii) This follows from part (vii) and the fact that the union of a chain of structures of uncountable cofinality is metrically complete.

(ix) This is obvious.

(x) If \mathfrak{M} is compact, then \mathfrak{N} is as well, and we have $\mathfrak{M} \cong \mathfrak{N}$, so assume that \mathfrak{M} and

\mathfrak{N} are not compact.

Let $\{\mathfrak{M}_\lambda\}_{\lambda < \kappa}$ and $\{\mathfrak{N}_\lambda\}_{\lambda < \kappa}$ be specializing chains for \mathfrak{M} and \mathfrak{N} , respectively. Let $\{m_i\}_{i < \kappa}$ be an enumeration of $\bigcup_{\lambda < \kappa} \mathfrak{M}_\lambda$, and let $\{n_i\}_{i < \kappa}$ be an enumeration of $\bigcup_{\lambda < \kappa} \mathfrak{N}_\lambda$ (note that these exist, since for any infinite cardinal λ , if X is a metric space with density character λ , then $|X| \leq \lambda^{\aleph_0} \leq 2^{\lambda \cdot \aleph_0} = 2^\lambda$).

For each even $j < \kappa$, let a_j be the first element of $\{m_i\}_{i < \kappa}$ such that $m_i \neq a_k$ for any $k < j$ and such that $a_j \in \mathfrak{M}_{\aleph_0 + |j|}$. Note that by $(\aleph_0 + |j|)^+$ -saturation of $\mathfrak{M}_{\aleph_0 + |j|}$ (and the fact that \mathfrak{M} is not compact), such an $a_j = m_i$ always exists. Now let b_j be an element of $\mathfrak{N}_{\aleph_0 + |j|}$ such that $a_{< j} a_j \equiv b_{< j} b_j$. Such a b_j always exists by $(\aleph_0 + |j|)^+$ -saturation of $\mathfrak{N}_{\aleph_0 + |j|}$.

For each odd $j < \kappa$, let b_j be the first element of $\{n_i\}_{i < \kappa}$ such that $n_i \neq b_k$ for any $k < j$ and such that $b_j \in \mathfrak{N}_{\aleph_0 + |j|}$. Again note that such an element always exists. Let a_j be an element of $\mathfrak{M}_{\aleph_0 + |j|}$ such that $a_{< j} a_j \equiv b_{< j} b_j$. Again, such a b_j always exists.

Clearly we have that $a_{< \kappa} \equiv b_{< \kappa}$. We just need to argue that $a_{< \kappa}$ enumerates $\bigcup_{\lambda < \kappa} \mathfrak{M}_\lambda$ and likewise for $b_{< \kappa}$ and $\bigcup_{\lambda < \kappa} \mathfrak{N}_\lambda$. Assume that m_j is not equal to a_i for any $i < \kappa$, and let $\lambda < \kappa$ be the smallest infinite cardinality such that $m_j \in \mathfrak{M}_\lambda$. This implies that for every even ℓ with $\lambda \leq \ell < \kappa$, $a_\ell = m_o$ with $o < j$, but there are κ many such ℓ and $|[0, j]| < \kappa$, so this is a contradiction. The same argument works for b_i , so we have an elementary map between dense sub-pre-structures of \mathfrak{M} and \mathfrak{N} . Therefore $\mathfrak{M} \cong \mathfrak{N}$.

(xi) This follows immediately from parts (viii) and (x). \square

Definition 1.8.8. For a complete theory T , the *monster model of T* , written \mathfrak{C}_T , is a special model of T with $\text{cf}(\#\text{dc}\mathfrak{C}_T)$ much greater than any cardinalities we are explicitly considering (in particular greater than \aleph_0). Typically we will omit the subscript T when it is clear from context.

A cardinality is *small* if it is smaller than $\text{cf}(\#\text{dc}\mathfrak{C}_T)$. A set is *small* if its cardinality is small. A metric space or pre-structure is *small* if its cardinality is small.

A *global V-type* is a type in $S_V(\mathfrak{C}_T)$. A *global type* is a global V -type for some V (typically a single variable). ◁

By the typical convention, we will regard all of our structures as being elementary substructures of \mathfrak{C}_T (an easy argument shows that any $\mathfrak{M} \models T$ with $\#\text{dc}\mathfrak{M} < \text{cf}(\#\text{dc}\mathfrak{C}_T)$ can be elementarily embedded in \mathfrak{C}_T). For a topological or type-set formula X and some $\bar{a} \in \mathfrak{C}_T$, we write $\models X(\bar{a})$ for $\mathfrak{C}_T \models X(\bar{a})$.

Chapter 2

Topometry and Definability

Type spaces in continuous logic have additional natural structure, beyond the logic topology. Specifically, there is a natural notion of the distance between two types.

Definition 2.0.1. For any $n \leq \omega$, n -tuple of variables \bar{v} , and arbitrary tuple of variables \bar{x} , the *induced metric on $S_{\bar{v}\bar{x}}(\mathcal{L})$ over \bar{x}* is given by

$$d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) = \inf\{d^{\mathfrak{M}}(\bar{a}, \bar{b}) : \mathfrak{M} \models p(\bar{a}, \bar{c}), q(\bar{b}, \bar{c})\},$$

where $\inf \emptyset = \text{db}(\mathcal{L})$.¹ We will drop the $/\bar{x}$ if \bar{x} is empty. We will also typically drop the \mathcal{L} and \bar{v} subscripts.

Unless otherwise stated, we take the topometric on $S_{\bar{v}}(\mathcal{L})$ to be $d_{\mathcal{L}\bar{v}}$.

If Σ is a partial $\mathcal{L}(\bar{v}\bar{x})$ -type, the *induced metric on $S_{\bar{v}\bar{x}}(\mathcal{L}, \Sigma)$* is the restriction of $d_{\mathcal{L}\bar{v}/\bar{x}}$ to $S_{\bar{v}\bar{x}}(\mathcal{L}, \Sigma) \subseteq S_{\bar{v}\bar{x}}(\mathcal{L})$. ◁

We will defer verifying that this is actually a metric to Proposition 2.1.4.

The $/\bar{x}$ subscript allows us to treat variables as if they were constants or parameters, specifically, if \bar{x} is a tuple of variables and \bar{c} is a tuple of fresh constants of the same length, then there is a natural bijection f between $S_{\bar{v}\bar{x}}(\Sigma)$ and $S_{\bar{v}}(\Sigma[\frac{\bar{c}}{\bar{x}}])$. Moreover, f is

¹We could make this anything in $[\frac{1}{2}\text{db}(\mathcal{L}), \infty]$ while maintaining that $d_{\Sigma, \mathcal{L}/\bar{x}}$ is an (extended) metric (although we need a choice in $[\text{db}(\mathcal{L}), \infty]$ to ensure that it is in fact a topometric), but this choice is consistent with our treatment of suprema and infima of empty sets in other cases.

a topometric isomorphism between $(S_{\bar{v}\bar{x}}(\Sigma), d_{/\bar{x}})$ and $(S_{\bar{v}}(\Sigma[\frac{\bar{c}}{\bar{x}}]), d)$.

Notation 2.0.2. For any $n \leq \omega$, if $A \subseteq \mathfrak{M}$ is a set of parameters and \bar{b}, \bar{c} are n -tuples in \mathfrak{M} , then $d^{\text{tp}}(\bar{b}, \bar{c}/A) = d(\text{tp}(\bar{b}/A), \text{tp}(\bar{c}/A))$. If we need to write out the elements of the tuples explicitly we will write either $d^{\text{tp}}(b_0 b_1 \dots b_{n-1}, c_0 c_1 \dots c_{n-1}/A)$ or $d^{\text{tp}}(b_0, b_1, \dots, b_{n-1}; c_0, c_1, \dots, c_{n-1}/A)$. \triangleleft

The following lemmas will be used frequently and implicitly.

Lemma 2.0.3. *Let \mathcal{L} be a signature and $\Sigma(\bar{x})$ be a partial \bar{x} -type. For any $p, q \in S_{\bar{v}\bar{x}}(\Sigma)$ if p and q contain the same $\mathcal{L}(\bar{x})$ -formulas, then there exists \mathfrak{M} and $\bar{a}, \bar{b}, \bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models p(\bar{a}\bar{c}), q(\bar{b}\bar{c})$ and $d^{\mathfrak{M}}(\bar{a}, \bar{b}) = d_{/\bar{x}}(p, q)$.*

Proof. Compactness. \square

Lemma 2.0.4. *Let \bar{v} be an at most countable tuple of variables, and let \bar{x} be a tuple of variables. For any open type-set $\mathcal{L}(\bar{v}\bar{x})$ -formula, $U(\bar{v}, \bar{x})$, for any pre-structure \mathfrak{M} and $\bar{b}\bar{a} \in M$, with \bar{b} assigned to \bar{v} and \bar{a} assigned to \bar{x} , if $d_{/\bar{x}, \text{inf}}(\text{tp}(\bar{b}, \bar{a}), \llbracket U \rrbracket) < \varepsilon$, then there exists $\bar{c} \in M$ such that $d(\bar{b}, \bar{c}) < \varepsilon$ and such that $\mathfrak{M} \models U(\bar{c}, \bar{a})$.*

Proof. Since $d_{/\bar{x}, \text{inf}}(\text{tp}(\bar{b}, \bar{a}), \llbracket U \rrbracket) < \varepsilon$, there exists a $p \in \llbracket U \rrbracket$ such that $d_{/\bar{x}, \text{inf}}(\text{tp}(\bar{b}, \bar{a}), p) < \varepsilon$. Therefore there exists an open formula V such that $V \models U$ and $p \in \llbracket V \rrbracket$. Finally, we have that $\mathfrak{M} \models \exists \bar{w} U(\bar{w}, \bar{a}) \wedge d(\bar{w}, \bar{b}) < \varepsilon$, so the required \bar{c} exists. \square

2.1 Topometric Spaces

The following definition was introduced in [BY08c], motivated by type spaces in continuous logic as well as the interaction between the natural topology and metric on the automorphism groups of metric structures.

Definition 2.1.1. A *topometric space*, (X, τ, d) , is a topological space (X, τ) together with a metric d (or possibly an extended metric, allowing $d(x, y) = \infty$) on X satisfying the following compatibility conditions:

- d is lower semi-continuous, i.e. $\{(x, y) \in X^2 : d(x, y) > \varepsilon\}$ is open for every $\varepsilon \geq 0$.
- The topology induced by the metric refines τ .

If X is a set, a *topometry on X* is a pair (τ, d) such that (X, τ, d) is a topometric space. If (X, τ) is a topological space, a *topometric for (X, τ)* is a metric d on X such that (X, τ, d) is a topometric space. If X is a topological space with an understood topology, we may write (X, d) instead of (X, τ, d) .

If $Y \subseteq X$, with (X, τ, d) a topometric space, then $(Y, \tau \upharpoonright Y, d \upharpoonright Y)$ is a topometric space as well. This is referred to as the *induced topometry on Y* .

If (X, τ) is a Hausdorff space, then the *discrete topometric on X* , written δ , is the topometric such that $\delta(x, y) = 1$ if and only if $x \neq y$. If (X, τ) is a topology induced by a metric d , we say that d is the *trivial topometric for X* . \triangleleft

Typically in our applications the topology of a space will be fixed and there will be one or occasionally more topometrics for that topology.

While the metric in a topometric space does induce a topology, purely topological properties of (X, d) are typically not as important as metric properties. As such, topological words such as open, closed, continuous, and compact refer to the topology τ unless otherwise specified by the adjective *metric* or the adverb *metrically* (e.g. metrically compact). If we need to emphasize that we are talking about a property of (X, τ) , we will use the adjective *topological* or the adverb *topologically*. If a term applies exclusively to metric spaces (e.g. 1-Lipschitz, open or closed balls) then it applies to (X, d) .

Here are a few important facts about topometric spaces. Most of these were originally proven in [BY08c]. Part (vii) was proven in [BY10b], and the proof we give here is essentially the same, although we are able to simplify the presentation somewhat as we are assuming compactness.

Proposition 2.1.2. *Let (X, τ, d) be a topometric space.*

(i) (X, τ) is Hausdorff.

(ii) If $F \subseteq X$ is closed, then it is metrically closed.

Now assume that (X, τ) is compact.

(iii) If $F \subseteq X$ is closed, then $x \mapsto d_{\text{inf}}(x, F)$ is a lower semi-continuous function, i.e. for any $\varepsilon \geq 0$, $F^{\leq \varepsilon} = \{x \in X : d_{\text{inf}}(x, F) \leq \varepsilon\}$ is closed. In particular, closed balls are topologically closed.

(iv) (X, d) is spherically complete (i.e. the intersection of a descending sequence of closed balls is non-empty) and therefore metrically complete.

(v) If $f : X \rightarrow Y$ is a continuous function into a metric space (Y, ρ) , then f is metrically uniformly continuous.

(vi) For any non-empty closed F and G , there exists $x \in F$ and $y \in G$ such that $d(x, y) = d_{\text{inf}}(F, G)$.

(vii) If $F, G \subseteq X$ are closed, then for any $r > 0$ with $r < d_{\text{inf}}(F, G)$ there is a continuous 1-Lipschitz function $f : X \rightarrow [0, r]$ such that $F \subseteq f^{-1}(0)$ and $G \subseteq f^{-1}(r)$. In particular, $d(x, y) = \sup\{|f(x) - f(y)|; x, y \in X, f : X \rightarrow \mathbb{R} \text{ cont. 1-Lip.}\}$.

Proof. (i) A space X is Hausdorff if and only if the diagonal is closed. The fact that the diagonal is closed in a topometric space follows from the fact that the metric is a metric (and not a pseudo-metric) and the lower semi-continuity of the metric: $\{(x, y) : d(x, y) = 0\}$ is the diagonal and is closed.

(ii) This is equivalent to the fact that the metric refines the topology.

(iii) This follows from the fact that the projection maps $\pi_{0,1} : X^2 \rightarrow X$ are closed whenever X is compact and from $F^{\leq \varepsilon} = \pi_1(\pi_0^{-1}(F) \cap \{(x, y) : d(x, y) \leq \varepsilon\})$.

(iv) This follows from the fact that closed balls are topologically closed.

(v) For each $\varepsilon > 0$, let $U_\varepsilon = \{(x, y) \in X^2 : |f(x) - f(y)| < \varepsilon\}$. Each U_ε is an open subset of X^2 that contains the diagonal. Since $\bigcap_{\delta > 0} \{(x, y) \in X^2 : d(x, y) \leq \delta\}$ is the diagonal, by compactness, there must exist a $\delta > 0$ such that $\{(x, y) \in X^2 : d(x, y) \leq \delta\} \subseteq U_\varepsilon$. In other words, for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $x, y \in X$, $d(x, y) \leq \delta \rightarrow |f(x) - f(y)| < \varepsilon$, which is precisely metric uniform continuity.

(vi) Let $r = d_{\inf}(F, G)$. For each $\varepsilon > 0$, we have that $F \cap G^{\leq r+\varepsilon}$ is a non-empty closed set. Therefore $F \cap \bigcap_{\varepsilon > 0} G^{\leq r+\varepsilon}$ is non-empty. Let x be some element of that set. By the symmetry of the metric, for each $\varepsilon > 0$, $B_{\leq r+\varepsilon}(x) \cap G$ is a non-empty closed set, so $\bigcap_{0 < \varepsilon} B_{\leq r+\varepsilon}(x) \cap G$ is non-empty. Let y be an element of that set. By assumption, $d(x, y) \geq r$, but we also have that $d(x, y) \leq r + \varepsilon$ for every $\varepsilon > 0$, so $d(x, y) = r = d_{\inf}(F, G)$.

(vii) Let $\{r_i\}_{0 < i < \omega}$ be an enumeration of all rational multiples of r in the interval $(0, r)$. Let $L_0 = \{(F, 0), (X, r)\}$ and $U_0 = \{(G, r), (X, 0)\}$. Note that L_0 and U_0 satisfy the following condition (for $i = 0$):

$(*)_i$ For every $(A, a) \in L_i$ and $(B, b) \in U_i$, A and B are closed and if $a < b$, then $d_{\inf}(A, B) > b - a$. For every $(A, a), (A', a') \in L_i$ with $a < a'$, $A \subseteq A'$. For every

$(B, b), (B', b') \in U_i$ with $b < b'$, $B \supseteq B'$. For every $(A, c) \in L_i$ and $(B, c) \in U_i$,
 $A \cup B = X$.

Given L_i and U_i , satisfying $(*)_i$, construct L_{i+1} and U_{i+1} as follows: By $(*)_i$ and by the triangle inequality, the sets

$$\begin{aligned}\mathbb{L}_i &= \bigcup \{A^{\leq r_k - r_{i+1}} : (A, r_k) \in L_i, r_k < r_{i+1}\} \text{ and} \\ \mathbb{U}_i &= \bigcup \{B^{\leq r_{i+1} - r_k} : (B, r_k) \in U_i, r_k > r_{i+1}\}\end{aligned}$$

are disjoint. The sets in these unions are closed by part (iii). Since \mathbb{L}_i and \mathbb{U}_i are finite unions of closed sets they are closed. Since X is compact and Hausdorff, we can find open $V_i \supseteq \mathbb{L}_i$ and $W_i \supseteq \mathbb{U}_i$ such that $V_i \cap W_i = \emptyset$. Let $L_{i+1} = L_i \cup \{(X \setminus W_i, r_{i+1})\}$, and let $U_{i+1} = U_i \cup \{(X \setminus V_i, r_{i+1})\}$.

Now to verify that L_{i+1} and U_{i+1} satisfy $(*)_{i+1}$. Clearly $X \setminus W_i$ and $X \setminus V_i$ are closed. By construction, if $(B, b) \in U_i$ with $b > r_{i+1}$, then $X \setminus W_i$ and $B^{\leq b - r_k}$ are disjoint, so by compactness $d_{\inf}(X \setminus W_i, B) > b - r_k$. The same argument works for $(A, a) \in L_i$ with $a < r_{i+1}$. Clearly $X \setminus W_i \supseteq \mathbb{L}_i \supseteq A$ for any A with $(A, a) \in L_i$ and $a < r_{i+1}$, and likewise for $(B, b) \in U_i$ with $b > r_{i+1}$. And finally, since V_i and W_i are disjoint, we have that $(X \setminus V_i) \cup (X \setminus W_i) = X$.

Let $L = \bigcup_{i < \omega} L_i$ and $U = \bigcup_{i < \omega} U_i$. Set $f(x) = \inf\{a \in (0, r) : x \in A, (A, a) \in L\}$, with $\inf \emptyset = r$. Note that for any $s \in [0, r]$, $\{x : f(x) \leq s\} = \bigcap \{A : (A, a) \in L, a > s\}$ (although note that even if there is an $(A, s) \in L$, it may not be the case that $\{x : f(x) \leq s\} = A$). Therefore $\{x : f(x) \leq s\}$ is a closed set. By construction, if $a < b < a'$ are rational multiples of r in $(0, r)$ and $(A, a), (A', a') \in L$ and $(B, b) \in U$, then $A \subseteq X \setminus B \subseteq A'$. Therefore if $f(x) < s$, for some $s \in [0, r]$, then we can find

$a < b < a'$ (rational multiples of r in $(0, r)$) such that $f(x) < a < b < a' < s$, and we have that for any $y \in X \setminus B$, $f(y) < a' < s$. Therefore, since we can do this for any x such that $f(x) < s$, we have that $\{x : f(x) < s\}$ is an open set. Therefore $f(x)$ is a continuous function. Also, note that by construction, if $x \in F$, then $f(x) = 0$, and if $x \in G$, then $f(x) = r$, and so $F \subseteq f^{-1}(0)$ and $G \subseteq f^{-1}(r)$.

So now we just need to verify that $f(x)$ is 1-Lipschitz. Let x and y be such that $|f(x) - f(y)| > \varepsilon$ with $f(x) < f(y)$. For any $a < b$, rational multiples of r in $[0, r]$, such that $f(x) < a < b < f(y)$, we have that $x \in A$ with $(A, a) \in L$ and $y \in B$ with $(B, b) \in U$. By $(*)_i$ for sufficiently large i , we have that $d(x, y) \geq d_{\text{inf}}(A, B) > b - a$. Since we can do this for any a, b satisfying the above conditions, we have that $d(x, y) \geq |f(x) - f(y)|$. Since we can do this for any x and y , we have that $f(x)$ is 1-Lipschitz. \square

Notation 2.1.3. For a topometric space (X, d, τ) with subsets $Y, Q \subseteq X$ with $Q \subseteq Y$, we will write $\text{int}_Y Q$ and $\text{cl}_Y Q$ for the topological interior and closure, respectively, of Q in the subspace Y . We will write $\text{ext}_Y U$ for the topological exterior (i.e. $\text{int}_Y(Y \setminus U)$). We will write Q_Y° and \overline{Q}^Y for the metric interior and closure, respectively, of Q in the subspace Y . If $Y = X$ or if Y is otherwise clear from context we will typically drop the subscript.

If we need to specify the relevant topology or metric, we will write $\text{int}_{Y,\tau} Q$, $\text{cl}_{Y,\tau} Q$, $\text{ext}_{Y,\tau} Q$, $Q_{Y,d}^\circ$, or $\overline{Q}^{Y,d}$. If we only need to specify the particular topology or metric, we may drop the Y . \triangleleft

Proposition 2.1.4. *For any \mathcal{L} , at most countable tuple \bar{v} , tuple \bar{x} , and partial $\mathcal{L}(\bar{v}\bar{x})$ -type, $d_{\mathcal{L}\bar{v}/\bar{x}}$ is a metric on $S_{\bar{v}\bar{x}}(\Sigma, \mathcal{L})$. Furthermore, $d_{\mathcal{L}\bar{v}/\bar{x}}$ is a topometric for $S_n(\Sigma, \mathcal{L})$.*

Proof. To see that $d_{\mathcal{L}\bar{v}/\bar{x}}$ is a metric, first note that $d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) \geq 0$, $d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) =$

$d_{\mathcal{L}\bar{v}/\bar{x}}(q, p)$, and $d_{\mathcal{L}\bar{v}/\bar{x}}(p, p) = 0$ all clearly hold. If we assume that $d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) = 0$, then for each $k < \omega$, let $(\mathfrak{M}_k, \bar{a}_k, \bar{b}_k, \bar{c}_k)$ be a structure with tuples such that $\mathfrak{M}_k \models p(\bar{a}_k, \bar{c}_k)$, $\mathfrak{M}_k \models q(\bar{b}_k, \bar{c}_k)$, and $d(\bar{a}_k, \bar{b}_k) < 2^{-k}$. Then if \mathcal{U} is a non-principal ultrafilter on ω , we have that $\mathfrak{M}_{\mathcal{U}} \models d(\bar{a}_{\mathcal{U}}, \bar{b}_{\mathcal{U}}) = 0$, so $p = q$.

For any three types $p, q, r \in S_{\bar{v}\bar{x}}(\mathcal{L})$ which imply the same $\mathcal{L}(\bar{x})$ -type, let $d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) = s$ and $d_{\mathcal{L}\bar{v}/\bar{x}}(q, r) = t$. We have that

$$\Sigma = \{p(\bar{v}, \bar{x}), q(\bar{w}, \bar{x}), d\bar{v}\bar{w} = s\} \text{ and}$$

$$\Pi = \{q(\bar{w}, \bar{x}), r(\bar{u}, \bar{x}), d\bar{w}\bar{u} = t\}$$

are both satisfiable by compactness. The Craig interpolation theorem (Proposition B.1.8) implies that $\Sigma \cup \Pi$ is satisfiable (since q is a consistent type), and therefore

$$\{p(\bar{v}, \bar{x}), r(\bar{u}, \bar{x}), d\bar{v}\bar{u} \leq s + t\}$$

is satisfiable, so $d_{\mathcal{L}\bar{v}/\bar{x}}(p, r) \leq s + t = d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) + d_{\mathcal{L}\bar{v}/\bar{x}}(q, r)$, as required.

To show that $d_{\mathcal{L}\bar{v}/\bar{x}}$ is a topometric on $S_{\bar{v}\bar{x}}(\mathcal{L})$, first note that if \mathcal{F} is a filter of points in $S_{\bar{v}\bar{x}}(\mathcal{L})^2$ such that $\{(p, q) \in S_{\bar{v}\bar{x}}(\mathcal{L})^2 : d_{\mathcal{L}\bar{v}/\bar{x}}(p, q) \leq \varepsilon\} \in \mathcal{F}$, then for any ultrafilter $\mathcal{U} \supseteq \mathcal{F}$, if \mathcal{U} converges to (r, s) , then for each $(p, q) \in S_{\bar{v}\bar{x}}(\mathcal{L})^2$ we can find a structure with tuples $(\mathfrak{M}_{pq}, \bar{a}_{pq}, \bar{b}_{pq}, \bar{c}_{pq})$ such that $\mathfrak{M}_{pq} \models p(\bar{a}_{pq}, \bar{c}_{pq}) \wedge q(\bar{b}_{pq}, \bar{c}_{pq}) \wedge d\bar{a}_{pq}\bar{b}_{pq} = d_{\mathcal{L}\bar{v}/\bar{x}}(p, q)$. Taking an ultraproduct with \mathcal{U} gives $\mathfrak{M}_{\mathcal{U}} \models r(\bar{a}_{\mathcal{U}}, \bar{c}_{\mathcal{U}}) \wedge s(\bar{b}_{\mathcal{U}}, \bar{c}_{\mathcal{U}}) \wedge d\bar{a}_{\mathcal{U}}\bar{b}_{\mathcal{U}} \leq \varepsilon$, so $d_{\mathcal{L}\bar{v}/\bar{x}}$ is lower semi-continuous.

To see that $d_{\mathcal{L}\bar{v}/\bar{x}}$ refines the topology, let p be an $\mathcal{L}(\bar{v}\bar{x})$ -type, and let $[\varphi(\bar{v}, \bar{x}) > 0]$ be an open neighborhood of it. Let $\varphi(p) = r > 0$. Find $\varepsilon > 0$ small enough that if

$d(\bar{a}, \bar{b}) < \varepsilon$, then $|\varphi(\bar{a}, \bar{c}) - \varphi(\bar{b}, \bar{c})| < r$. Now by construction we have that $B_{<\varepsilon}^{d_{\mathcal{L}^{\bar{v}/\bar{x}}}}(p) \subseteq \llbracket \varphi > 0 \rrbracket$, so since we can do this for any p and any open neighborhood in a basis of the topology, we have that the metric topology refines the logic topology. \square

There is a special property that the metrics in type spaces (over theories, rather than arbitrary types) have over and above arbitrary compact topometric spaces. This property was identified in [BY08c].

Definition 2.1.5. A topometric d for the space X is *open* if for every open set $U \subseteq X$ and every $\varepsilon > 0$, the set $U^{<\varepsilon} = \{x \in X : (\exists y \in U)d(x, y) < \varepsilon\}$ is open as well. \triangleleft

Note that this property is not preserved under passing to subspaces (see Counterexample C.1.3).

Proposition 2.1.6. *For any at most countable tuple of variables \bar{v} , any tuple of variables \bar{x} , and any partial $\mathcal{L}(\bar{x})$ -type Σ , the metric $d_{/\bar{x}}$ on $S_{\bar{v}\bar{x}}(\Sigma)$ is open.*

Proof. Let $U(\bar{v}\bar{x})$ be an open formula. We want to argue that

$$\llbracket U(\bar{v}\bar{x}) \rrbracket_{\bar{v}\bar{x}}^{d_{/\bar{x}} < \varepsilon} = \llbracket \exists \bar{w}(U(\bar{w}\bar{x}) \wedge d_{\bar{v}\bar{x}} < \varepsilon) \rrbracket_{\bar{v}\bar{x}}.$$

If $\varepsilon > \text{db}\mathcal{L}$, then $\llbracket U(\bar{v}\bar{x}) \rrbracket^{<\varepsilon} = S_{\bar{v}\bar{x}}(\Sigma)$, so assume that $\varepsilon \leq \text{db}\mathcal{L}$.

Assume that $p \in \llbracket U(\bar{v}\bar{x}) \rrbracket^{<\varepsilon}$, so there exists a $q \in \llbracket U(\bar{v}\bar{x}) \rrbracket$ such that $d_{/\bar{x}}(p, q) < \varepsilon$. Note that since $d_{/\bar{x}}(p, q) < \text{db}(\mathcal{L})$, p and q imply the same complete $\mathcal{L}(\bar{x})$ -type $r(\bar{x})$. Let \mathfrak{M} be a sufficiently saturated structure such that $\mathfrak{M} \models r(\bar{c})$. Find some $\bar{a} \in M$ such that $\mathfrak{M} \models p(\bar{a}\bar{c})$. By saturation, there is a $\bar{b} \in M$ such that $\mathfrak{M} \models q(\bar{b}\bar{c})$ and $d(\bar{a}, \bar{b}) = d^{\text{tp}}(\bar{a}, \bar{b}/\bar{c}) < \varepsilon$. Therefore $\mathfrak{M} \models \exists \bar{w}(U(\bar{w}\bar{x}) \wedge d_{\bar{v}\bar{x}} < \varepsilon)$ and so $p \in \llbracket \exists \bar{w}(U(\bar{w}\bar{x}) \wedge d_{\bar{v}\bar{x}} < \varepsilon) \rrbracket$.

Assume that $p \in \llbracket \exists \bar{w}(U(\bar{w}\bar{x}) \wedge d\bar{v}\bar{w} < \varepsilon) \rrbracket$. Let $\mathfrak{M} \models p(\bar{a}\bar{c})$. Now we have that there is some $\bar{b} \in M$ such that $\mathfrak{M} \models U(\bar{a}\bar{c}) \wedge d\bar{a}\bar{b} < \varepsilon$. Therefore $d^{tp}(\bar{a}, \bar{b}/\bar{c}) < \varepsilon$ and hence $\text{tp}(\bar{b}\bar{c}) \in \llbracket U(\bar{v}\bar{x}) \rrbracket$ (this is where we use that Σ is an $\mathcal{L}(\bar{x})$ -type and not an $\mathcal{L}(\bar{v}\bar{x})$ -type) and $p \in \llbracket U(\bar{v}\bar{x}) \rrbracket^{<\varepsilon}$. \square

Non-open metrics will be relevant to continuous logic in the context of approximate isomorphism and approximate categoricity (Chapter 6). The induced metric on a type space over a partial type (rather than just a partial theory) is also typically not open.

2.2 Example: The Halo




Figure 1: \mathfrak{H} and a, b, c




Figure 2: Minimizing $d(ac, b'c')$ for $b'c' \equiv bc$

We should pause for a moment to highlight what is perhaps the subtlest aspect of the induced metric on type spaces, namely, its interaction with parameters. Given a parameter a , there is a natural embedding $i : S_x(a) \rightarrow S_{xy}(\emptyset)$. Specifically, given $r \in S_x(a)$, we have $i(r) = \{F[\frac{y}{a}] : F \in r\}$. This would more commonly be described as taking types of the form $p(x, a)$ to types of the form $p(x, y)$. The subtlety is that while this map is a topological embedding, it is not a metric embedding (although it is always 1-Lipschitz).

Example 2.2.1 (The Halo Structure). *Elements a, b, c in some structure such that $d^{\text{tp}}(ac, bc) < d^{\text{tp}}(a, b/c)$.*

Description. Let \mathfrak{H} be a structure in the empty signature with $\text{db}(\mathcal{L}) = 4$. Let S^1 be the unit circle in \mathbb{R}^2 , and let M be $S^1 \times \{0, 1\}$ with the following metric:

- $d((x, 0), (y, 0))$ is the distance in \mathbb{R}^2 .
- $d((x, 0), (y, 1)) = d((x, 0), (y, 0)) + 1$.
- For $x \neq y$, $d((x, 1), (y, 1)) = d((x, 0), (y, 0)) + 2$.

Now let x and y be points in S^1 such that $d^{\mathbb{R}^2}(x, y) = \varepsilon > 0$, and let $a = (x, 1)$, $b = (y, 1)$, $c = (x, 0)$. Clearly we have that $d(ac, bc) = 2 + \varepsilon$. It's not hard to show that $d^{\text{tp}}(a, b/c) = 2 + \varepsilon$ as well.

On the other hand, when computing $d^{\text{tp}}(ac, bc)$, we are allowed to 'shift' the pair bc so as to line up a and b . If y' is the point on S^1 such that $d(x, y') = \varepsilon$, but $y \neq y'$, and we set $b' = a$ and $c' = (y', 0)$, then we have $bc \equiv b'c'$ (by an explicit automorphism) and $d(ac, b'c') = \varepsilon$.

So we have $d^{\text{tp}}(ac, bc) \leq d(ac, b'c') = \varepsilon < 2 + \varepsilon$, as required. In fact it is actually the case that $d^{\text{tp}}(ac, bc) = \varepsilon$. □

Let $T = \text{Th}(\mathcal{H})$, then the type spaces $S_{xy}(T)$ and $S_x(T_c) = S_x(c)$ can be described explicitly and are pictured in Figure 3. $S_{xy}(T)$ is 4 disjoint closed semicircles and a topologically isolated point. $S_x(T_c)$ is 2 disjoint closed semicircles. The solid lines indicate regions where the type space is topologically locally metrically compact, so in particular the logic and metric topologies agree. The dotted lines indicate regions where the metric is uniformly discrete, i.e. there is an $\varepsilon > 0$ such that if $x \neq y$, then $d(x, y) > \varepsilon$.




Figure 3: The type spaces $S_{xy}(T)$ and $S_x(T_c)$

Note, however, that the dotted lines are not topologically discrete. The inclusion map i takes each semicircle in $S_x(T_c)$ to its respective semicircle on the top row of $S_{xy}(T)$. In the $x \in S^1 \times \{1\}$ region i is mapping a uniformly discrete metric bijectively onto a compact metric, which is the furthest an injection could possibly be from being an isometric embedding.

2.3 Locatable and Definable Sets

Definable sets are perhaps the most notorious concept in continuous logic. They do not match the immediate intuition that a formula $\varphi(x)$ in discrete logic is like a closed formula $\varphi(x) = 0$, which would imply that a definable set in a structure \mathfrak{M} should be the set of elements a such that $\varphi^{\mathfrak{M}}(a) = 0$. The confusion is not helped here by the fact that while the term ‘definable set’ in discrete logic typically refers to a subset of a particular

structure, in the context of this thesis the term ‘definable set’ refers mostly to a certain class of closed formulas, specifically those admitting relative existential quantification for open formulas.

2.3.1 Locatable Sets

Some forms of relative quantification are free. For any closed formulas $F(x, y)$ and $G(x)$ and open formulas $U(x, y)$ and $V(x)$, we have

$$(\forall x \in G)U(x, y) \equiv \forall x(G(x) \rightarrow U(x, y)),$$

$$(\exists x \in G)F(x, y) \equiv \exists x(G(x) \wedge F(x, y)),$$

$$(\exists x \in V)U(x, y) \equiv \exists x(V(x) \wedge U(x, y)), \text{ and}$$

$$(\forall x \in V)F(x, y) \equiv \forall x(V(x) \rightarrow F(x, y)).$$

Moreover, we can do this for open and closed type-set formulas as well. As it turns out this is a characterization of closed type-set formulas up to logical equivalence:

Proposition 2.3.1. *A type-set formula $X(\bar{x})$ is logically equivalent to a closed type-set formula if and only if for every closed type-set formula $F(x, y)$, the type-set formula $\exists \bar{x}(X(\bar{x}) \wedge F(\bar{x}, \bar{y}))$ is logically equivalent to a closed type-set formula.*

Proof. $X(\bar{x})$ is logically equivalent to $\exists \bar{z}(X(\bar{z}) \wedge \bigwedge_{i < |\bar{x}|} x_i = z_i)$. □

This characterization does not extend to open type-set formulas, because $x = y$ is a closed formula, not an open formula.

Definition 2.3.2. Let T be a (possibly incomplete) theory and \bar{v} an n -tuple of variables for $n \leq \omega$. A set $L \subseteq S_{\bar{v}}(T)$ is *locatable* if for every open formula $U(\bar{v}, \bar{w})$, the type-set formula $\exists \bar{v} L(\bar{v}) \wedge U(\bar{v}, \bar{w})$ is logically equivalent to an open type-set formula modulo T .

If \bar{v} is an n -tuple of variables for $n \leq \omega$ and \bar{x} is an arbitrary tuple of variables, we say that $L \subseteq S_{\bar{v}\bar{x}}(T)$ is *\bar{x} -uniformly locatable* if for every open formula $U(\bar{v}, \bar{w}, \bar{x})$, the type-set formula $\exists \bar{v} L(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$ is logically equivalent to an open type-set formula modulo T .

For any \mathcal{L} -theory T , a type-set $\mathcal{L}(\bar{v}\bar{x})$ -formula $L(\bar{v}, \bar{x})$ is *\bar{x} -uniformly locatable over T* if $\llbracket L \rrbracket_{T, \bar{v}\bar{x}}$ is uniformly locatable. $L(\bar{v})$ is *locatable over T* if it is uniformly locatable in the empty tuple of variables over T . \triangleleft

Note that this easily implies the same for open type-set formula (i.e. if $L(\bar{v}, \bar{x})$ is uniformly locatable, $\exists \bar{v} L(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$ is logically equivalent to an open type-set formula for any open type-set formula $U(\bar{v}, \bar{w})$), since relative existential quantification distributes over arbitrary disjunctions. Also note that open sets are clearly locatable. Furthermore, (uniform) locatability is preserved in expansions and extensions:

Proposition 2.3.3. *Let T be an \mathcal{S} -theory. Let $L \subseteq S_{\bar{v}\bar{x}}(T, \mathcal{S})$ be \bar{x} -uniformly locatable. Let $\mathcal{S}' \supseteq \mathcal{S}$ be an expansion, and let $T' \supseteq T$ be an \mathcal{S}' -theory extending T . Then $L' = \{p \in S_{\bar{v}\bar{x}}(T', \mathcal{S}') : p \upharpoonright \mathcal{S} \in L\}$ is \bar{x} -uniformly locatable.*

Proof. Let $U(\bar{v}, \bar{w}, \bar{x}) = (\varphi(\bar{v}, \bar{w}, \bar{x}) > 0)$. For each $i < \omega$, find $\varepsilon_i > 0$ such that if $d(\bar{a}, \bar{a}') < \varepsilon_i$ then $|\varphi(\bar{a}, \bar{b}, \bar{c}) - \varphi(\bar{a}', \bar{b}, \bar{c})| \leq 2^{-i}$. Since the metric is part of \mathcal{S} , we have that for any $i < \omega$, $\exists \bar{v} L(\bar{v}, \bar{x}) \wedge d(\bar{v}, \bar{u}) < \varepsilon_i$ is logically equivalent to an open type-set $\mathcal{S}(\bar{u}\bar{x})$ -formula, $V_i(\bar{u}, \bar{x})$ modulo T . So now let $W(\bar{w}, \bar{x}) = \bigvee_{i < \omega} \exists \bar{u} V_i(\bar{u}, \bar{x}) \wedge \varphi(\bar{u}, \bar{w}, \bar{x}) > 2^{-i}$. Clearly $T, \exists \bar{v} L(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x}) \models W(\bar{w}, \bar{x})$, so the same is true with T' in place of T .

By construction we also have $T, W(\bar{w}, \bar{x}) \models \exists \bar{v} L(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$, so again the same is true with T' in place of T . Therefore we have that $W(\bar{w}, \bar{x})$ is logically equivalent to $\exists \bar{v} L(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$ modulo T' . \square

The set of parameters over which a uniformly locatable family is non-empty is always open.

Proposition 2.3.4. *If $L(\bar{v}, \bar{x}) \subseteq S_{\bar{v}\bar{x}}(T)$ is \bar{x} -uniformly locatable, then the set*

$$X(\bar{x}) = \{p \in S_{\bar{x}}(T) : \llbracket L \rrbracket_{T, \bar{v}\bar{x}} \cap \llbracket p \rrbracket_{T, \bar{v}\bar{x}} \neq \emptyset\}$$

is open.

Proof. $X(\bar{x})$ is logically equivalent to $\exists \bar{v} L(\bar{v}, \bar{x}) \wedge \top(\bar{v}, \bar{x})$, where \top is the open formula that is always true. \square

We can give a topometric characterization of locatable sets.

Definition 2.3.5. Let (X, τ, d) be a topometric space. We say that a set $L \subseteq X$ is *locatable* if for every $\varepsilon > 0$, $L \subseteq \text{int}_X L^{<\varepsilon}$. \triangleleft

Proposition 2.3.6. *For any (possibly incomplete) theory T , any n -tuple \bar{v} with $n \leq \omega$, any \bar{x} , and any set $X \subseteq S_{\bar{v}\bar{x}}(T)$, X is uniformly locatable over \bar{x} in the sense of Definition 2.3.2 if and only if X is locatable with regards to the topometric $d_{/\bar{x}}$ in the sense of Definition 2.3.5.*

Proof. Assume that \bar{x} is empty and that X is locatable in the first sense. Then for any $\varepsilon > 0$, we have that $L^{<\varepsilon} = \llbracket \exists \bar{w} L(\bar{w}) \wedge d\bar{w}\bar{v} < \varepsilon \rrbracket_{\bar{v}}$, so $\text{int} L^{<\varepsilon} = L^{<\varepsilon} \supseteq L$, and L is locatable in the second sense.

Now assume that L is locatable in the second sense. Let $U(\bar{v}, \bar{w}) = (\varphi(\bar{v}, \bar{w}) > 0)$ be an open formula. For each $i < \omega$, find $\varepsilon_i > 0$ such that if $\varphi(\bar{a}, \bar{b}) > 2^{-i}$ and $d(\bar{a}, \bar{a}') < \varepsilon_i$,

then $\varphi(\bar{a}', \bar{b}) > 0$ (this always exists by uniform continuity). Let $V_i = \text{int } L^{<2^{-i}}$, and consider the open formula

$$O(\bar{w}) = \bigvee_{i < \omega} \exists \bar{v} V_i(\bar{v}) \wedge \varphi(\bar{v}, \bar{w}) > 2^{-i}.$$

Note that since $T, L(\bar{v}) \models V_i(\bar{v})$ for each $i < \omega$, we clearly have that

$$T, \exists \bar{v} L(\bar{v}) \wedge U(\bar{v}, \bar{w}) \models O(\bar{w}).$$

Now assume that $\mathfrak{M} \models T, O(\bar{b})$. Then for some $i < \omega$,

$$\mathfrak{M} \models \exists \bar{v} V_i(\bar{v}) \wedge \varphi(\bar{v}, \bar{b}) > 2^{-i}.$$

So there is $\bar{a} \in M$ such that $\mathfrak{M} \models V_i(\bar{a}) \wedge \varphi(\bar{a}, \bar{b}) > 2^{-i}$. In an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ there is \bar{a}' such that $d(\bar{a}, \bar{a}') < 2^{-i}$. Therefore by construction, $\varphi(\bar{a}', \bar{b}) > 0$, and we have that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}) \wedge U(\bar{v}, \bar{b})$. Therefore $T, O(\bar{w}) \models \exists \bar{v} L(\bar{v}) \wedge U(\bar{v}, \bar{b})$, as required.

For the second statement, the entire proof relativizes to parameters. The easiest way to see this is to note that we proved the first part for arbitrary incomplete theories and to convert the variables \bar{x} to new constants. \square

Now we collect some facts about locatable sets in arbitrary topometric spaces. In [BY10b] the concept of a normal topometric space is defined. The definition given there is not the same as the definition we will give here, but was proven equivalent in [BY10b].

Definition 2.3.7. A topometric space (D, τ, d) is *normal* if for any closed sets $F, G \subseteq X$ with $d_{\text{inf}}(F, G) > r > 0$, there is a 1-Lipschitz continuous function $f : X \rightarrow [0, r]$ such

that $F \subseteq f^{-1}(0)$ and $G \subseteq f^{-1}r$. \triangleleft

As was mentioned in Proposition 2.1.2, compact topometric spaces satisfy this property, so compact topometric spaces, and in particular type spaces, are normal.

Proposition 2.3.8. *Let (X, τ, d) be a topometric space.*

(i) *Open sets are locatable.*

(ii) *L is locatable if and only if \bar{L} (the metric closure) is locatable.*

(iii) *If L is locatable, then \bar{L} is G_δ (countable intersection of open sets).*

(iv) *An arbitrary union of locatable sets is locatable, so in particular the locatable sets form a complete lattice under \subseteq , with \sqcup given by union.*

(v) *If L is locatable, $Y \subseteq X$ has $L \subseteq Y$, $\tau' \supseteq \tau \upharpoonright Y$, and $d' \leq d \upharpoonright Y$, then L is locatable in (Y, τ', d') .*

Assume that X is normal (as a topometric space).

(vi) *If L is locatable and U is open, then $L \cap U$ is locatable.*

(vii) *Locatability is local to the boundary: L is locatable if and only if for every $x \in \partial L$ ($= \text{cl}L \setminus \text{int}L$, the topological boundary of L) and every neighborhood $U \ni x$, $x \in \text{int}(L \cap U)^{<\varepsilon}$ for every $\varepsilon > 0$.*

(viii) *If L_0 is locatable and $L_1 \subseteq L_0$ is locatable in $(L_0, \tau \upharpoonright L_0, d \upharpoonright L_0)$, then L_1 is locatable in (X, τ, d) .*

Assume that d is an open metric and that X is possibly not normal.

(ix) L is locatable if and only if it is a Hausdorff metric limit of open sets (i.e. for every $\varepsilon > 0$, there is an open set U such that $d_H(L, U) < \varepsilon$, where $d_H(A, B) = \inf\{\varepsilon : A \subseteq B^{<\varepsilon}, B \subseteq A^{<\varepsilon}\}$).

(x) L is locatable if and only if $d_{\text{inf}}(x, L)$ is upper semi-continuous (i.e. $L^{<\varepsilon} = \{x : d_{\text{inf}}(x, L) < \varepsilon\}$ is open for every $\varepsilon > 0$).

Assume that X is compact.

(xi) If F and G are closed and non-empty, then there exists $x \in F$ and $y \in G$ such that $d(x, y) = d_{\text{inf}}(F, G)$.

Proof. (i)–(iv) are straightforward.

(v) For any $\varepsilon > 0$, we have $Y \cap L^{d < \varepsilon} \subseteq L^{d' < \varepsilon}$, and so

$$L \subseteq \text{int}_{X, \tau} L^{d < \varepsilon} \subseteq \text{int}_{Y, \tau'} (Y \cap L^{d < \varepsilon}) \subseteq \text{int}_{Y, \tau'} L^{d' < \varepsilon}.$$

Hence L is locatable in (Y, τ', d') .

(vi) Consider $x \in L \cap U$, with U open. We have that $d_{\text{inf}}(x, X \setminus U) > r > 0$ for some r (since the metric refines the topology), so let $f : X \rightarrow [0, r]$ be a 1-Lipschitz continuous function such that $f(x) = 0$ and $X \setminus U \subseteq f^{-1}(r)$.

Now consider $\varepsilon > 0$ with $\varepsilon < \frac{r}{2}$. By assumption, $x \in \text{int } L^{<\varepsilon}$. Find open $V \subseteq \text{int } L^{<\varepsilon}$ such that $x \in V$. Now consider $V \cap \{f < \varepsilon\}$. This is also an open neighborhood of x . Now consider $L \cap (V \cap \{f < \varepsilon\})^{<\varepsilon}$. Note that

$$(V \cap \{f < \varepsilon\})^{<\varepsilon} \subseteq \{f < \varepsilon\}^{<\varepsilon} \subseteq \{f < 2\varepsilon\} \subseteq U,$$

hence

$$L \cap (V \cap \{f < \varepsilon\})^{<\varepsilon} \subseteq U,$$

as well. So now let $y \in V \cap \{f < \varepsilon\}$. By the choice of V , there is a $z \in L$ such that $d(y, z) < \varepsilon$. This implies that z is in $(V \cap \{f < \varepsilon\})^{<\varepsilon}$ and therefore $z \in U$. Finally, this gives that $V \cap \{f < \varepsilon\} \subseteq (L \cap U)^{<\varepsilon}$, and so we have $x \in \text{int}(L \cap U)^{<\varepsilon}$. Since we can do this for any $x \in L \cap U$, we have that $L \cap U \subseteq \text{int}(L \cap U)^{<\varepsilon}$. This implies the same for any $\delta > \varepsilon$, so since we can do this for arbitrarily small $\varepsilon > 0$, we have that $L \cap U$ is locatable.

(vii) The \Leftarrow direction is straightforward. The \Rightarrow direction follows from (vi).

(viii) Pick $\varepsilon > 0$. Since L_1 is locatable in L_0 , we have that $\text{int}_{L_0} L_1^{<\varepsilon/2} \supseteq L_1$. By the definition of relative topology, $\text{int}_{L_0} L_1^{<\varepsilon/2} = L_0 \cap U$ for some open-in- X set U . Therefore $\text{int}_{L_0} L_1^{<\varepsilon/2}$ is locatable in X , by (vi). This implies that

$$L_1 \subseteq \text{int}_{L_0} L_1^{<\varepsilon/2} \subseteq \text{int}_X \left(\text{int}_{L_0} L_1^{<\varepsilon/2} \right)^{<\varepsilon/2} \subseteq L_1^{<\varepsilon},$$

and hence $L_1 \subseteq \text{int}_X L_1^{<\varepsilon}$. Since we can do this for any $\varepsilon > 0$, L_1 is locatable in X , as required.

(ix) The \Rightarrow direction is straightforward, so assume that L is a Hausdorff limit of open sets. For each $\varepsilon > 0$, let U_ε be an open set such that $d_H(L, U_\varepsilon) < \varepsilon$. By openness, $U_{\varepsilon/2}^{<\varepsilon/2}$ is an open set. By the definition of the Hausdorff metric we have that $L \subseteq U_{\varepsilon/2}^{<\varepsilon/2}$ and $U_{\varepsilon/2} \subseteq L^{<\varepsilon/2}$. Therefore by the triangle inequality we have $U_{\varepsilon/2}^{<\varepsilon/2} \subseteq L^{<\varepsilon}$, so

$$L \subseteq U_{\varepsilon/2}^{<\varepsilon/2} \subseteq \text{int} L^{<\varepsilon},$$

as required. So L is locatable.

(x) The \Leftarrow direction is obvious. Assume that L is locatable. By the triangle inequality we have that for any $\varepsilon > \delta > 0$, $(\text{int } L^{<\delta})^{<\varepsilon-\delta} \subseteq L^{<\varepsilon}$. Furthermore, by openness, $(\text{int } L^{<\delta})^{<\varepsilon-\delta}$ is an open set. Therefore $(\text{int } L^{<\delta})^{<\varepsilon-\delta} \subseteq \text{int } L^{<\varepsilon}$. Since L is locatable, we also have that $L^{<\varepsilon-\delta} \subseteq (\text{int } L^{<\delta})^{<\varepsilon-\delta}$, and hence for a fixed $\varepsilon > 0$,

$$L^{<\varepsilon} = \bigcup_{0 < \delta < \varepsilon} L^{<\varepsilon-\delta} \subseteq \bigcup_{0 < \delta < \varepsilon} (\text{int } L^{<\delta})^{<\varepsilon-\delta} \subseteq \text{int } L^{<\varepsilon} \subseteq L^{<\varepsilon},$$

as required.

(xi) For any $\varepsilon > d_{\text{inf}}(F, G)$, $F \cap G^{\leq \varepsilon}$ is non-empty. Therefore by compactness, $H = F \cap \bigcap_{\varepsilon > d_{\text{inf}}(F, G)} G^{\leq \varepsilon}$ is non-empty. \square

Part (vi) of Proposition 2.3.8 is completely trivial to prove in the context of type spaces with the machinery we've developed for type-set formulas: If $L(\bar{x})$ is a locatable formula, then $\exists \bar{v}(L(\bar{v}) \wedge U(\bar{v})) \wedge V(\bar{v}, \bar{w})$ is clearly logically equivalent to $\exists \bar{v} L(\bar{v}) \wedge (U(\bar{v}) \wedge V(\bar{v}, \bar{w}))$ (by the associativity of \wedge), which is equivalent to an open type-set formula by the locatability of L .

These seem to be about the only nice properties locatable sets have in general (and in type spaces, for that matter).

Open formulas have a special property over and above arbitrary locatable formulas. Open formulas do not just admit relative weak existential quantification with other open formulas, they admit relative existential quantification. This motivates the following definition:

Definition 2.3.9. A type-set formula $L(\bar{v})$ is *strongly locatable over T* if for any structure $\mathfrak{M} \models T$ and any open formula $U(\bar{v}, \bar{w})$ such that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}) \wedge U(\bar{v}, \bar{a})$, there exists

$\bar{b} \in M$ such that $\mathfrak{M} \models L(\bar{b}) \wedge U(\bar{b}, \bar{a})$.

A type-set formula $L(\bar{v}, \bar{x})$ is *strongly and \bar{x} -uniformly locatable² over T* or *strongly, \bar{x} -uniformly locatable over T* if for any structure $\mathfrak{M} \models T$, any parameters $\bar{a}\bar{b}$, and any open formula $U(\bar{v}, \bar{w}, \bar{x})$ such that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge U(\bar{v}, \bar{b}, \bar{a})$, there exists $\bar{c} \in M$ such that $\mathfrak{M} \models L(\bar{c}, \bar{a}) \wedge U(\bar{c}, \bar{b}, \bar{a})$. \triangleleft

Note that we have defined this for structures \mathfrak{M} , rather than pre-structures. Open formulas are ‘strongly locatable over pre-structures,’ but this requirement is too strong in general.

Proposition 2.3.10. *Let T be a (possibly incomplete) theory, let \bar{v} be an at most countable tuple of variables, and let \bar{x} be a tuple of variables.*

- (i) *If $\{L_i(\bar{v}, \bar{x})\}_{i \in I}$ are strongly and \bar{x} -uniformly locatable, then $\bigcup_{i \in I} L_i(\bar{v}, \bar{x})$ is strongly and \bar{x} -uniformly locatable.*
- (ii) *If $L(\bar{v}, \bar{x})$ is \bar{x} -uniformly locatable and $\llbracket L \rrbracket \subseteq S_{\bar{v}\bar{x}}(T)$ is G_δ (countable intersection of open sets), then it is strongly and \bar{x} -uniformly locatable.*

Proof. (i) This is straightforward.

(ii) Assume that $L(\bar{v}, \bar{x})$ is \bar{x} -uniformly locatable and $\llbracket L \rrbracket$ is G_δ . Let $L(\bar{v}, \bar{x})$ be logically equivalent to $\bigwedge_{i < \omega} W_i(\bar{v}, \bar{x})$, where $\{W_i\}_{i < \omega}$ is a sequence of open type-set formulas. For any $\delta > 0$ and open formula $O(\bar{v}, \bar{x})$, let $V_\delta^O(\bar{v}, \bar{x})$ be an open type-set formula corresponding to $\llbracket L(\bar{v}, \bar{x}) \wedge O(\bar{v}, \bar{x}) \rrbracket^{d/\bar{x} < \delta}$. Let $O_0(\bar{v}, \bar{x})$ be an open formula such that $\llbracket O_0 \rrbracket = S_{\bar{v}\bar{x}}(T)$.

²There is neither a sense in which the uniformity is strong nor a sense in which the strength is particularly uniform.

Fix a structure \mathfrak{M} and $\bar{a}\bar{b} \in M$ and an open formula $U(\bar{v}, \bar{w})$. Assume that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge U(\bar{v}, \bar{b}, \bar{a})$. Let $U(\bar{v}, \bar{w}, \bar{x}) \equiv (\varphi(\bar{v}, \bar{w}, \bar{x}) > 0)$. There is an $r > 0$ such that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge \varphi(\bar{v}, \bar{b}, \bar{a}) > r$. Find $\varepsilon_0 > 0$ small enough that if $d(\bar{c}, \bar{c}') \leq \varepsilon_0$, then $|\varphi(\bar{c}, \bar{b}, \bar{a}) - \varphi(\bar{c}', \bar{b}, \bar{a})| < r$.

Note that $\mathfrak{M} \models \exists \bar{v} V_{2^{-1}\varepsilon_0}^{O_0}(\bar{v}, \bar{a}) \wedge W_0(\bar{v}, \bar{a}) \wedge \varphi(\bar{v}, \bar{b}, \bar{a}) > r$, so since $V_{2^{-1}\varepsilon_0}^{O_0}$ and W_0 are open type-set formulas, there actually is $\bar{c}_0 \in M$ such that $\mathfrak{M} \models V_{2^{-1}\varepsilon_0}^{O_0}(\bar{c}_0, \bar{a}) \wedge W_0(\bar{c}_0, \bar{a}) \wedge \varphi(\bar{c}_0, \bar{b}, \bar{a}) > r$. Note that by construction, $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge d(\bar{v}, \bar{c}_0) < 2^{-1}\varepsilon_0$.

For each $i < \omega$, given \bar{c}_i , $\varepsilon_i > 0$, and $O_i(\bar{v}, \bar{x})$ such that $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge d(\bar{v}, \bar{c}_i) < 2^{-i-1}\varepsilon_i$ and $\mathfrak{M} \models V_{2^{-i-1}\varepsilon_i}^{O_i}(\bar{c}_i, \bar{a}) \wedge W_i(\bar{c}_i, \bar{a})$, let p_i be a type in $\llbracket L \wedge O_i \rrbracket$ such that $d_{/\bar{x}}(\text{tp}(\bar{c}_i/\bar{a}), p_i) < 2^{-i-1}\varepsilon_i$. Find an open formula $O_{i+1}(\bar{v}, \bar{x})$ such that $p_i \in \llbracket O_{i+1} \rrbracket$ and $\text{cl}\llbracket O_{i+1} \rrbracket \subseteq \llbracket W_i \rrbracket$. Find $\varepsilon_{i+1} > 0$ such that $\varepsilon_{i+1} \leq \varepsilon_i$ and $(\text{cl}\llbracket O_{i+1} \rrbracket)^{d_{/\bar{x}} \leq \varepsilon_{i+1}} \subseteq \llbracket W_i \rrbracket$.

Now we have that $\mathfrak{M} \models \exists \bar{v} V_{2^{-i-2}\varepsilon_{i+1}}^{O_{i+1}}(\bar{v}, \bar{a}) \wedge W_{i+1}(\bar{v}, \bar{a}) \wedge d(\bar{v}, \bar{c}_i) < 2^{-i-1}\varepsilon_i$ (because \mathfrak{M} has an elementary extension realizing p_i). Therefore there exists $\bar{c}_{i+1} \in M$ such that $\mathfrak{M} \models V_{2^{-i-2}\varepsilon_{i+1}}^{O_{i+1}}(\bar{v}, \bar{a}) \wedge W_{i+1}(\bar{v}, \bar{a}) \wedge d(\bar{c}_{i+1}, \bar{c}_i) < 2^{-i-1}\varepsilon_i$.

By construction, we have that $\{\bar{c}_i\}_{i < \omega}$ is a Cauchy sequence. Let $\bar{c}_\omega \in \mathfrak{M}$ be its limit. For any $i < \omega$, we have that $d(\bar{c}_i, \bar{c}_\omega) < \sum_{k=i}^{\infty} 2^{-k-1}\varepsilon_k \leq \varepsilon_i$ (since $\{\varepsilon_i\}_{i < \omega}$ is non-increasing). Therefore by construction, we have that $\varphi^{\mathfrak{M}}(\bar{c}_\omega, \bar{b}, \bar{a}) > 0$, or in other words $\mathfrak{M} \models U(\bar{c}_\omega, \bar{b}, \bar{a})$, and $\text{tp}(\bar{c}_\omega/\bar{a}) \in \llbracket W_i \rrbracket$ for each $i < \omega$. Therefore in particular we have that $\mathfrak{M} \models L(\bar{c}_\omega, \bar{a})$, as required.

Therefore L is strongly, \bar{x} -uniformly locatable. □

Unlike locatable sets, there is almost certainly no hope for a precise characterization (topometric or otherwise) of strongly locatable sets, given the complexity of omitting types in structures in continuous logic [FM18].

Corollary 2.3.11. *Let $L(\bar{v}, \bar{x})$ be \bar{x} -uniformly locatable. If L' is the type-set formula corresponding to $\overline{[L]}^{d/\bar{x}}$, then $[L']$ is G_δ , and so in particular L' is strongly and uniformly locatable.*

Proof. This follows from Proposition 2.3.8 part (ii). □

Corollary 2.3.12. *If $L(\bar{v}, \bar{x})$ is \bar{x} -uniformly locatable, \mathfrak{M} is an $|\bar{x}|^+$ -saturated structure, and $\mathfrak{M} \models \exists \bar{v} L(\bar{v}, \bar{a}) \wedge U(\bar{v}, \bar{b}, \bar{a})$ for some open type-set formula $U(\bar{v}, \bar{w}, \bar{x})$, then there exists $\bar{c} \in M$ such that $\mathfrak{M} \models L(\bar{c}, \bar{a}) \wedge U(\bar{c}, \bar{b}, \bar{a})$.*

Proof. If we let L' be the type-set formula corresponding to $\overline{[L]}^{d/\bar{x}}$, then $\mathfrak{M} \models \exists \bar{v} L'(\bar{v}, \bar{a}) \wedge U(\bar{v}, \bar{b}, \bar{a})$. By Proposition 2.3.10, L' is strongly, \bar{x} -uniformly locatable, so there is $\bar{c}' \in \mathfrak{M}$ such that $\mathfrak{M} \models L'(\bar{c}', \bar{a}) \wedge U(\bar{c}', \bar{b}, \bar{a})$. Find $\varepsilon > 0$ small enough that if $d(\bar{c}', \bar{c}) < \varepsilon$, then $\mathfrak{M} \models U(\bar{c}, \bar{b}, \bar{a})$. Find $p(\bar{v}, \bar{a})$ such that $d_{/\bar{x}}(\text{tp}(\bar{c}'\bar{a}), p) < \varepsilon$ and such that $p \in [L]$. By $|\bar{x}|^+$ -saturation, there is \bar{c} such that $\mathfrak{M} \models p(\bar{c}, \bar{a})$ and $d(\bar{c}', \bar{c}) < \varepsilon$. Therefore by construction, $\mathfrak{M} \models L(\bar{c}, \bar{a}) \wedge U(\bar{c}, \bar{b}, \bar{a})$. □

2.3.2 Definable Sets

Locatable sets are a kind of topometric generalization of open sets. The analogous generalization for closed sets does not really produce a larger class of sets.³ Definable sets are the corresponding analog of clopen sets: closed and locatable.

Definition 2.3.13. Given a topometric space (X, τ, d) , we say that $D \subseteq X$ is *definable* or a *definable set* if it is closed and locatable. If X is itself a subspace, we may say that D is *relatively definable in X* .

³Although this depends on precisely how one states it. If (X, τ, d) is a topometric space with an open metric, then a set Y is locatable if and only if $d_{\text{inf}}(x, Y)$ is upper semi-continuous, so one possible analog would be to require that $d_{\text{inf}}(x, Y)$ be lower semi-continuous. If X is compact, we have that $d_{\text{inf}}(x, Y)$ is lower semi-continuous if and only if \bar{Y} is topologically closed.

A point $x \in X$ is called *d-atomic* if $\{x\}$ is locatable (and therefore definable). x is called *weakly d-atomic* if $\text{int } B_{<\varepsilon}^d(x)$ is non-empty for every $\varepsilon > 0$. If there is more than one topometric on X we will use terms such as ρ -atomic or weakly d_0 -atomic.

Given an at most countable tuple of variables \bar{v} and a tuple of variables \bar{x} , a closed type-set formula $D(\bar{v}, \bar{x})$ is *\bar{x} -uniformly definable over T* if $\llbracket D \rrbracket_{T, \bar{v}\bar{x}}$ is locatable with regards to the topometric $d_{/\bar{x}}$. A closed type-set formula is *definable over T* if it is locatable over T . If we need to specify the language, formulas, or theory, we will say that D is *\bar{x} -uniformly $\mathcal{L}(\bar{v}\bar{x})$ -definable over T* or that D is *$\mathcal{L}(\bar{v})$ -definable over T* . \triangleleft

Note that (uniformly) definable formulas are strongly (and uniformly) locatable by Proposition 2.3.10 and the fact that topologically closed sets are metrically closed in topometric spaces.

While technically we should have terms such as *d-definable* or *d-locatable* to specify the metric, situations in which multiple topometrics are relevant occur most frequently with the concepts of *d-atomic* and weakly *d-atomic*.

Remark 2.3.14. [BY08c] introduced the concept of weakly *d-atomic* points, although there they are called *d-isolated* and weakly *d-isolated*. I thought that the term ‘*d-isolated*’ was too easy to interpret as ‘isolated with regards to the metric.’

[BY08c] expressed a desire to maintain a separation between topometric terminology and continuous logic terminology. We have not successfully maintained such a distinction here, but if one did wish to do so I would recommend using the term ‘*clocatable*,’ or perhaps ‘*d-findable*,’ in place of ‘*definable*.’

Proposition 2.3.15. *Let (X, τ, d) be a topometric space.*

- (i) A definable set is G_δ .⁴
- (ii) A finite union of definable sets is definable.
- (iii) Any clopen set is definable.
- (iv) If X is normal (as a topometric space) and D , E , and $\partial D \cap \partial E$ are definable, then $D \cap E$ is definable (in particular if ∂D and ∂E are disjoint, such as when one of D and E is clopen).
- (v) If D is definable in (X, τ, d) , $Y \subseteq X$ has $D \subseteq Y$, $\tau' \supseteq \tau \upharpoonright Y$, and $d' \leq d \upharpoonright Y$, then D is definable in (Y, τ', d') .
- (vi) If X is compact, then a closed set $D \subseteq X$ is definable if and only if there is a continuous function $f : X \rightarrow [0, \infty]$ such that $D = \{x : f(x) = 0\}$ and for all $x \in X$, $d_{\text{inf}}(x, D) \leq f(x)$. (Note that the topology on $[0, \infty]$ is chosen so that it is homeomorphic to $[0, 1]$.) If $0 < \text{diam}(X) < \infty$, then f can be chosen such that $f(x) \in [0, \text{diam}(X)]$.
- (vii) If X is compact and d is an open metric, then a closed set $D \subseteq X$ is definable if and only if $d_{\text{inf}}(x, D)$ is continuous.
- (viii) If d is an open metric, then x is d -atomic if and only if x is weakly d -atomic.
- (ix) If there exists a collection \mathcal{U} , with $|\mathcal{U}| \leq \kappa$, of open sets such that for any closed set $F \subseteq X$ and open neighborhood $V \supseteq F$, there is $U \in \mathcal{U}$ with $V \supseteq U \supseteq F$ (in particular, if X is compact and has a basis of cardinality κ), then the collection of

⁴Compare this with the following fact: A clopen set is open.

definable sets has metric density character at most κ with regards to the Hausdorff metric (i.e. $d_H(X, Y) = \inf\{\varepsilon > 0 : X \subseteq Y^{<\varepsilon}, Y \subseteq X^{<\varepsilon}\}$).

Proof. (i) Straightforward.

(ii) This follows from the facts that finite unions of closed sets are closed and arbitrary unions of locatable sets are locatable.

(iii) Obvious.

(iv) $D \cap E = (D \cap \text{int } E) \cup (\text{int } D \cap E) \cup (\partial D \cap \partial E)$. By Proposition 2.3.8 part (vi), $D \cap \text{int } E$ and $\text{int } D \cap E$ are locatable sets, so we have that $D \cap E$ is a union of locatable sets and is therefore locatable. Furthermore we have that $D \cap E$ is closed, so it is definable.

(v) This follows from Proposition 2.3.8 part (v) and the fact that D is still closed in (Y, τ') .

(vi) If X is empty then this is trivial. If X is a singleton and D is empty, let $f(x) = 1$. If X is a singleton and D is non-empty, let $f(x) = 0$.

Assume that X has more than one point (so in particular $\text{diam}(X) > 0$). Let $U_0 = X$. Given U_i , find $\varepsilon_i > 0$ with $\varepsilon_i < \infty$ such that $\varepsilon_i < 2^{-i}$ and $\varepsilon_i < \varepsilon_{i-1}$ (or $\varepsilon_i < \infty$ if $i = 0$), and such that $D^{\leq \varepsilon_i} \subseteq U_i$ (this exists by compactness, since $D = \bigcap_{\varepsilon > 0} D^{\leq \varepsilon}$). Let $U_{i+1} = \text{int } D^{\leq \varepsilon_i}$. Note that $U_i \supseteq D^{\leq \varepsilon_i} \supseteq U_{i+1}$.

Now for each $i < \omega$ with $i > 1$, let $g_i : X \rightarrow [0, 1]$ be a continuous function such that $D^{\leq \varepsilon_i} \subseteq g_i^{-1}(0)$ and $X \setminus U_i \subseteq g_i^{-1}(1)$. Such a g_i exists by Urysohn's lemma. Now let

$$h(x) = \sum_{1 < i < \omega} (\varepsilon_{i-2} - \varepsilon_{i-1}) g_i(x).$$

Since $\{\varepsilon_i\}_{i < \omega}$ is a strictly decreasing sequence of positive numbers that limits to 0,

this sum converges uniformly and $h(x)$ is a continuous function. Note that $h(x) \in [0, \varepsilon_0]$.

For any x such that $0 < d_{\text{inf}}(x, D) \leq \varepsilon_0$, find $k < \omega$ such that $\varepsilon_k \geq d_{\text{inf}}(x, D) > \varepsilon_{k+1}$. (Note that such a k always exists.) This implies that $x \notin D^{\leq \varepsilon_{k+1}} \supseteq U_{k+2}$. Since $x \notin U_{k+2}$, $x \notin U_\ell$ for any $\ell > k + 1$. Therefore

$$h(x) \geq \sum_{k+1 < i < \omega} \varepsilon_{i-2} - \varepsilon_{i-1} = \varepsilon_k \geq d_{\text{inf}}(x, D).$$

For any x such that $d_{\text{inf}}(x, D) > \varepsilon_0$, we have $x \notin D^{\leq \varepsilon_i}$ for any $i < \omega$ with $i > 1$, so

$$h(x) = \sum_{1 < i < \omega} \varepsilon_{i-2} - \varepsilon_{i-1} = \varepsilon_0.$$

Finally, set

$$f(x) = \frac{\varepsilon_0 h(x)}{\varepsilon_0 - h(x)}$$

if $h(x) < \varepsilon_0$ and $f(x) = \infty$ if $h(x) \geq \varepsilon_0$. Note that $f(x)$ is continuous and that by construction, $f(x) \geq h(x)$ and $f(x) = 0$ if and only if $h(x) = 0$.

For $x \notin D$ with $d_{\text{inf}}(x, D) \leq \varepsilon_0$, we already have that $d_{\text{inf}}(x, D) \leq f(x)$. For x with $d_{\text{inf}}(x, D) > \varepsilon_0$, we have that $h(x) = \infty \geq d_{\text{inf}}(x, D)$.

Now we just need to verify that $D = \{x : f(x) = 0\}$. Clearly we have that if $x \in D$, then $x \in D^{\leq \varepsilon_i}$ for every $i < \omega$, so $h(x) = 0$ and therefore $f(x) = 0$. Conversely, assume that $x \notin D$, then there is some $i < \omega$ with $i > 0$, such that $x \notin D^{\leq \varepsilon_i} \supseteq U_{i+1}$ (since D is closed and the metric refines the topology), so $h(x) \geq \varepsilon_{i-1} - \varepsilon_i > 0$, and therefore $f(x) > 0$. So we have that $D = \{x : f(x) = 0\}$, as required.

If $0 < \text{diam}(X) < \infty$, we can use $\min\{f(x), \text{diam}(X)\}$ instead of $f(x)$, to ensure that the functions output is in $[0, \text{diam}(X)]$.

(vii) The \Leftarrow direction is clear. For the \Rightarrow direction, since X is compact, by Proposition 2.1.2 for any closed set F and any $\varepsilon > 0$, $F^{\leq \varepsilon}$ is closed, or equivalently that $d_{\text{inf}}(x, F)$ is lower semi-continuous. Since D is also locatable, $d_{\text{inf}}(x, D)$ is both lower semi-continuous and upper semi-continuous, and therefore is continuous.

(viii) The \Rightarrow direction is clear. The \Leftarrow direction follows from Proposition 2.3.8 part (ix).

(ix) Let $\{D_i\}_{i < \kappa^+}$ be a collection of definable sets. Fix $\varepsilon > 0$. For each $i < \kappa^+$, find $U_i \in \mathcal{U}$ such that $D \subseteq U_i \subseteq \text{int } D^{\leq \varepsilon}$. By the pigeonhole principle, there are $i < j < \kappa^+$ such that $U_i = U_j$. Therefore $D_i \subseteq D_j^{\leq \varepsilon}$ and $D_j \subseteq D_i^{\leq \varepsilon}$, and so $d_H(D_i, D_j) \leq \varepsilon$. Since we can do this for any $\varepsilon > 0$, there is no κ^+ -sized ($> \varepsilon$)-separated set of definable sets, and the collection of definable sets has Hausdorff metric density character at most κ . \square

Note the conspicuous absence of any particularly strong statements regarding *intersections* in Proposition 2.3.15. In general the intersection of two definable sets may be some frozzy, useless conglomeration.

Corollary 2.3.16. *If $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$, then it is logically equivalent modulo $\Sigma(\bar{x})$ to a closed formula.*

Proof. Proposition 2.3.15 part (vii) implies that $\llbracket D \rrbracket$ is the zeroset of the continuous function $p \mapsto d_{\text{inf}}(p, \llbracket D \rrbracket)$. By Proposition 1.5.9, this implies that D is logically equivalent to a closed formula. \square

Notation 2.3.17. If $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over T , by an abuse of notation we will also write $D(\bar{v}, \bar{x})$ to mean a real formula corresponding to the distance predicate of D (modulo T). If we wish to emphasize that we are talking about the distance predicate and not the closed formula, we will write $d_{\text{inf}}(\bar{v}, D(\cdot, \bar{x}))$ or $d_{\text{inf}}(\bar{v}, D)$. \triangleleft

Proposition 2.3.18 (The Disconnect between Something and Nothing). *If $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$, a partial $\mathcal{L}(\bar{x})$ -type, then $F(\bar{x}) = \exists \bar{v} D(\bar{v}, \bar{x})$ is a clopen formula over Σ (i.e. $\llbracket F(\bar{x}) \rrbracket$ is clopen in $S_{\bar{x}}(\Sigma)$).*

Proof. $F(\bar{x})$ is logically equivalent modulo Σ to $\exists \bar{v}(D(\bar{v}, \bar{x}) < \text{db}(\mathcal{L}))$. \square

The following proposition is largely equivalent to Theorem 9.12 in [BYBHU08], although they don't cover the case of uniformly definable sets and their axioms exclude the possibility of an empty definable set.

Proposition 2.3.19 (Axiom for Definability). *A real formula $\varphi(\bar{v}, \bar{x})$ is the distance predicate of a \bar{x} -uniformly definable set over $\Sigma(\bar{x})$ if and only if*

$$\begin{aligned} \Sigma(\bar{x}) \models \forall \bar{v}(\varphi(\bar{v}, \bar{x}) = \text{db}(\mathcal{L})) \\ \forall \bar{v}[\exists \bar{w}(\varphi(\bar{w}, \bar{x}) = 0 \wedge d\bar{v}\bar{w} = \varphi(\bar{v}, \bar{x})) \\ \wedge \forall \bar{w}(\varphi(\bar{v}, \bar{x}) \leq \varphi(\bar{w}, \bar{x}) + d\bar{v}\bar{w})]. \end{aligned}$$

Proof. Assume that $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$, and let $\varphi(\bar{v}, \bar{x})$ be its distance predicate. Let \mathfrak{M} be a structure and \bar{a} a tuple such that $\mathfrak{M} \models \Sigma(\bar{a})$. If $\mathfrak{M} \not\models \exists \bar{v} D(\bar{v}, \bar{a})$, then by the definition of the induced metric on type space, we have $\mathfrak{M} \models \forall \bar{v} \varphi(\bar{v}, \bar{a})$.

If $\mathfrak{M} \not\models \exists \bar{v} D(\bar{v}, \bar{a})$, then $\mathfrak{M} \models \forall \bar{v}(\varphi(\bar{v}, \bar{x}) = \text{db}(\mathcal{L}))$, since φ is the distance predicate of an empty set.

If $\mathfrak{M} \models \exists \bar{v} D(\bar{v}, \bar{a})$, then we have $\mathfrak{M} \not\models \forall \bar{v} \varphi(\bar{v}, \bar{a}) = \text{db}(\mathcal{L})$, hence for any $\bar{b} \in M$, we have

$$\mathfrak{M} \models \exists \bar{w}(\varphi(\bar{w}, \bar{a}) = 0 \wedge d\bar{b}\bar{w} = d_{/\bar{x}, \text{inf}}(\text{tp}(\bar{b}\bar{a}), \llbracket D \rrbracket)),$$

by the definition of $d_{/\bar{x}}$ and the fact that φ is the distance predicate for D . Therefore in either case we have that φ satisfies the axiom for definability.

Now assume that $\varphi(\bar{v}, \bar{x})$ satisfies the axiom for definability over $\Sigma(\bar{x})$. Pick \mathfrak{M} and $\bar{a} \in M$ such that $\mathfrak{M} \models \Sigma(\bar{a})$. Now we have two cases. First assume that $\mathfrak{M} \models \forall \bar{v}(\varphi(\bar{v}, \bar{a}) = \text{db}(\mathcal{L}))$. Then we have that $(\varphi(\bar{v}, \bar{a}) = 0)$ is empty, so $\varphi(\bar{x}, \bar{a})$ is the distance predicate of its zeroset.

Now assume that

$$\mathfrak{M} \models \forall \bar{v}[\exists \bar{w}(\varphi(\bar{w}, \bar{a}) = 0 \wedge d\bar{v}\bar{w} = \varphi(\bar{v}, \bar{a})) \wedge \forall \bar{w}(\varphi(\bar{v}, \bar{a}) \leq \varphi(\bar{w}, \bar{a}) + d\bar{v}\bar{w})].$$

By looking at sufficiently saturated elementary extensions of \mathfrak{M} , we have that for any $q \in S_{\bar{v}\bar{x}}(\text{tp}(\bar{a}))$ (which is canonically isomorphic to $S_{\bar{v}}(\bar{a})$), $d_{/\bar{x}, \text{inf}}(q, \llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket) \leq \varphi(q, \bar{a})$, therefore $(\varphi(\bar{v}, \bar{a}) = 0)$ is definable. All that remains is to verify that $d_{/\bar{x}, \text{inf}}(q, \llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket) \geq \varphi(q, \bar{a})$ (so that in particular, they are equal). Find a type $r \in \llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket$ such that $d_{/\bar{x}}(q, r) = d_{/\bar{x}, \text{inf}}(q, \llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket)$. By sufficient saturation, for any $\bar{b} \in \mathfrak{M}$ such that $\mathfrak{M} \models q(\bar{b}, \bar{a})$, there is $\bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models r(\bar{c}, \bar{a})$ and such that $d(\bar{b}, \bar{c}) = d_{/\bar{x}}(q, r)$. This implies that $\varphi(\bar{b}, \bar{a}) \leq \varphi(\bar{c}, \bar{a}) + d(\bar{b}, \bar{c}) = 0 + d_{/\bar{x}}(q, r) = d_{/\bar{x}, \text{inf}}(q, \llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket)$, so the required equality holds, and we have that $\varphi(\bar{v}, \bar{a})$ is the distance predicate of $(\varphi(\bar{v}, \bar{a}) = 0)$.

Since this works for any \mathfrak{M} and \bar{a} such that $\mathfrak{M} \models \Sigma(\bar{a})$, we have that $(\varphi(\bar{v}, \bar{x}) = 0)$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$ and $\varphi(\bar{v}, \bar{x})$ is the witnessing distance predicate. \square

Corollary 2.3.20. *For any real formula $\varphi(\bar{v}, \bar{x})$, there is a closed formula $F(\bar{x})$ such that for any structure \mathfrak{M} and $\bar{a} \in M$, $\varphi(\bar{v}, \bar{a})$ is the distance predicate of a definable set over $\text{Th}(\mathfrak{M})$ if and only if $\mathfrak{M} \models F(\bar{a})$.*

Notation 2.3.21. For a real formula $\varphi(\bar{v}, \bar{x})$, we let

$$\text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{x}) = \forall \bar{v}(\varphi(\bar{v}, \bar{x}) = \text{db}(\mathcal{L}))$$

$$\vee \forall \bar{v} [\exists \bar{w}(\varphi(\bar{w}, \bar{x}) = 0 \wedge d\bar{v}\bar{w} = \varphi(\bar{v}, \bar{x})) \wedge \forall \bar{w}(\varphi(\bar{v}, \bar{x}) \leq \varphi(\bar{w}, \bar{x}) + d\bar{v}\bar{w})],$$

$$\text{def}_{\bar{v}}\varphi(\bar{v}, \bar{x}) = \min\{\sup_{\bar{x}} |\varphi(\bar{v}, \bar{x}) - \text{db}(\mathcal{L})|,$$

$$\sup_{\bar{v}} \max \left\{ \inf_{\bar{w}} \max\{|\varphi(\bar{w}, \bar{x})|, |d\bar{v}\bar{w} - \varphi(\bar{v}, \bar{w})|\}, \frac{1}{2} \sup_{\bar{w}} \varphi(\bar{v}, \bar{x}) \div (\varphi(\bar{w}, \bar{x}) + d\bar{v}\bar{w}) \right\} \}. \triangleleft$$

So in particular, $\Sigma(\bar{x}) \models \text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{x})$ if and only if $\varphi(\bar{v}, \bar{x})$ is the distance predicate of a \bar{x} -uniformly definable set over $\Sigma(\bar{x})$. Note that $\text{DEF}_{\bar{v}}$ is, syntactically speaking, a quantifier that takes a real formula and produces a closed formula, which is not something we have seen before. Also note that $\text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{x})$ is logically equivalent to $\text{def}_{\bar{v}}\varphi(\bar{v}, \bar{x}) = 0$, and moreover just like any other quantifiers, the free variables of $\text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{x})$ and $\text{def}_{\bar{v}}\varphi(\bar{v}, \bar{x})$ are contained in \bar{x} . The factor of $\frac{1}{2}$ is only required to make the following statement true.

Lemma 2.3.22. *For any pair of formulas $\varphi(\bar{v}, \bar{x}), \psi(\bar{v}, \bar{x})$ and any partial type $\Sigma(\bar{x})$,*

$$\left\| \text{def}_{\bar{v}}\varphi(\bar{v}, \bar{x}) - \text{def}_{\bar{v}}\psi(\bar{v}, \bar{x}) \right\|_{\Sigma} \leq \|\varphi(\bar{v}, \bar{x}) - \psi(\bar{v}, \bar{x})\|_{\Sigma}.$$

Proof. This follows from the analogous fact for the quantifiers \inf and \sup (Lemma 1.4.2), as well as the facts that $|\min\{a, b\} - \min\{c, d\}| \leq \max\{|a - c|, |b - d|\}$, $|\max\{a, b\} - \max\{c, d\}| \leq \max\{|a - c|, |b - d|\}$, $||a - b| - |a - c|| \leq |b - c|$, and $|(a \div b) - (c \div d)| \leq |a - c| + |b - d| \leq 2 \max\{|a - c|, |b - d|\}$. \square

Proposition 2.3.23 (Relative Quantification of Real Formulas). *A closed formula $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$ if and only if for every real formula $\varphi(\bar{v}, \bar{w}, \bar{x})$,*

there is a real formula $\chi(\bar{w}, \bar{x})$ such that for every model \mathfrak{M} of T , for any \bar{a} and \bar{b} in M ,

$$\chi^{\mathfrak{M}}(\bar{b}, \bar{a}) = \inf\{\varphi^{\mathfrak{M}}(\bar{c}, \bar{b}, \bar{a}) : \bar{c} \in M, \mathfrak{M} \models D(\bar{c}, \bar{a})\}.$$

The same statement holds for sup.

Proof. For the \Leftarrow direction, we have that there is a formula $\chi(\bar{w}, \bar{x})$ such that for any structure \mathfrak{M} of T and $\bar{a}, \bar{b} \in M$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, we have $\chi^{\mathfrak{M}}(\bar{b}, \bar{a}) = \inf\{d(\bar{c}, \bar{b}) : \bar{c} \in M, \mathfrak{M} \models D(\bar{c}, \bar{a})\}$. This implies that for any $p \in S_{\bar{v}, \bar{x}}(T)$, $d_{/\bar{x}, \text{inf}}(p, \llbracket D \rrbracket) = \chi(p)$, so D is definable.

For the \Rightarrow direction, let

$$\chi(\bar{w}, \bar{x}) = \inf_{\bar{v}} \min\{\varphi(\bar{v}, \bar{w}, \bar{x}) + \omega_{\varphi, \bar{v}}(|D(\bar{v}, \bar{x})|), \sup I(\varphi)\},$$

where $\omega_{\varphi, \bar{v}}$ is given in Definition 1.3.15. Note that $\omega_{\varphi, \bar{v}}(\text{db}(\mathcal{L})) = \sup I(\varphi) - \inf I(\varphi)$. (The absolute value bars are so that χ is well-defined as a formula, as $\omega_{\varphi, \bar{v}}$ is only defined for non-negative inputs.) Let \mathfrak{M} be a structure and \bar{a}, \bar{b} be chosen so that $\mathfrak{M} \models \Sigma(\bar{a})$.

Assume that $D(\mathfrak{M}, \bar{a})$ is empty. We have that $\mathfrak{M} \models \forall \bar{v} D(\bar{v}, \bar{a}) = \text{db}(\mathcal{L})$, so $\mathfrak{M} \models \forall \bar{w} \chi(\bar{w}, \bar{a}) = \sup I(\varphi)$ and we're done.

Assume that $D(\mathfrak{M}, \bar{a})$ is non-empty. We have that $s = \inf\{\varphi^{\mathfrak{M}}(\bar{c}, \bar{b}, \bar{a}) : \bar{c} \in M, \mathfrak{M} \models D(\bar{c}, \bar{a})\}$. Clearly by construction we have $\chi^{\mathfrak{M}}(\bar{b}, \bar{a}) \leq s$.

Now assume that $\chi^{\mathfrak{M}}(\bar{b}, \bar{a}) < r$. Then there is some $\bar{c}' \in M$ such that $\mathfrak{M} \models \varphi(\bar{c}', \bar{b}, \bar{a}) + \alpha(D(\bar{c}', \bar{a})) < r$. By strong, uniform locatability of uniformly definable sets, for any $\varepsilon > 0$ there is $\bar{c} \in M$ with $\mathfrak{M} \models D(\bar{c}, \bar{a})$, $d(\bar{c}, \bar{c}') < d(\bar{c}, D(\cdot, \bar{a})) + \varepsilon$. By the choice of α , we have that $\mathfrak{M} \models \varphi(\bar{c}, \bar{b}, \bar{a}) \leq \varphi(\bar{c}', \bar{b}, \bar{a}) + \alpha(D(\bar{c}, \bar{a}))$, so in particular $s \leq \varphi(\bar{c}', \bar{b}, \bar{a}) + \alpha(D(\bar{c}, \bar{a}))$.

Since we can do this for any $\varepsilon > 0$, by continuity of α we have that $s \leq \varphi(\bar{c}', \bar{b}, \bar{a}) + \alpha(D(\bar{c}', \bar{a}))$. Furthermore, since we can do this for any $r > \chi(\bar{b}, \bar{a})$, we have that $s \leq \chi(\bar{b}, \bar{a})$, and hence

$$\chi(\bar{b}, \bar{a}) = s = \inf\{\varphi^{\mathfrak{M}}(\bar{c}, \bar{b}, \bar{a}) : \bar{c} \in M, \mathfrak{M} \models D(\bar{c}, \bar{a})\},$$

as required.

For sup, recall that $\sup \varphi = -\inf -\varphi$. □

Corollary 2.3.24. *If $D(\bar{v}, \bar{x})$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$, then for any open formula $U(\bar{v}, \bar{w}, \bar{x})$, $\exists \bar{v} D(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$ is logically equivalent modulo $\Sigma(\bar{x})$ to an open formula, $V(\bar{w}, \bar{x})$.*

Proof. Let $U(\bar{v}, \bar{w}, \bar{x}) = (\varphi(\bar{v}, \bar{w}, \bar{x}) > 0)$ (as in Proposition 2.3.23). We have that $\exists \bar{v} D(\bar{v}, \bar{x}) \wedge U(\bar{v}, \bar{w}, \bar{x})$ is logically equivalent to $\chi(\bar{w}, \bar{x}) > 0$, with $\chi(\bar{w}, \bar{x})$ given by Proposition 2.3.23 (in the sup case). □

Notation 2.3.25. For any real formula $D(\bar{v}, \bar{x})$, we write:

$$\begin{aligned} \inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) &= \inf_{\bar{v}} \min\{\varphi(\bar{v}, \bar{w}, \bar{x}) + \omega_{\varphi, \bar{v}}(|D(\bar{v}, \bar{x})|), \sup I(\varphi)\} \text{ and} \\ \sup_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) &= \sup_{\bar{v}} \max\{\varphi(\bar{v}, \bar{w}, \bar{x}) - \omega_{\varphi, \bar{v}}(|D(\bar{v}, \bar{x})|), \inf I(\varphi)\} \end{aligned}$$

where $\omega_{\varphi, \bar{v}}$ is given in Definition 1.3.15. We will only use this notation when we are working in the context of a partial $\mathcal{L}(\bar{x})$ -type $\Sigma(\bar{x})$ (or, in particular, a fixed theory T) where $D(\bar{v}, \bar{x})$ is a real formula that is the distance predicate of a \bar{x} -uniformly definable

set over Σ . Let $\Lambda_c(\bar{x}) = \forall \bar{v}(D(\bar{v}, \bar{x}) = \text{db}(\mathcal{L}))$ and $\Lambda_o(\bar{x}) = \forall \bar{v}(D(\bar{v}, \bar{x}) > 0)$. We write

$$\begin{aligned} (\forall \bar{v} \in D(\cdot, \bar{x}))(\varphi(\bar{v}, \bar{w}, \bar{x}) > 0) &= \Lambda_o(\bar{x}) \vee \inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) > 0, \\ (\exists \bar{v} \in D(\cdot, \bar{x}))(\varphi(\bar{v}, \bar{w}, \bar{x}) > 0) &= \neg \Lambda_c(\bar{x}) \wedge \sup_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) > 0, \\ (\forall \bar{v} \in D(\cdot, \bar{x}))(\varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0) &= \Lambda_c(\bar{x}) \vee \inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0, \text{ and} \\ (\exists \bar{v} \in D(\cdot, \bar{x}))(\varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0) &= \neg \Lambda_o(\bar{x}) \wedge \sup_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0. \end{aligned}$$

When \bar{x} is the empty tuple or when we don't need to emphasize the role of \bar{x} , we will write D instead of $D(\bar{x})$. Again, we will only use this notation when $D(\bar{v}, \bar{x})$ is known to be the distance predicate of a \bar{x} -uniformly definable set. \triangleleft

The formulas $\Lambda_c(\bar{x})$ and $\Lambda_o(\bar{x})$ are closed and open formulas, respectively, saying that $D(\bar{x})$ is empty (if it is a definable set).

Remark 2.3.26. To make relative quantification over a real valued formula $D(\bar{v}, \bar{x})$ that may or may not be a distance predicate better behaved, we could introduce notation analogous to the notation we will have later for partial definable functions, namely something to the effect of $(\forall \bar{v} \in D(\cdot, \bar{x}) \downarrow)(\varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0)$ for

$$\text{DEF} \bar{v} D(\bar{v}, \bar{x}) \wedge \left[\left(\inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0 \right) \vee \Lambda_c(\bar{x}) \right],$$

meaning ' $D(\bar{v}, \bar{x})$ is a distance predicate in \bar{v} and for all \bar{v} in $D(\bar{x})$, $\varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0$ holds.'

The meaning of this expression only depends on $(\varphi(\bar{v}, \bar{w}, \bar{x}) \geq 0)$ as a closed formula rather than $\varphi(\bar{v}, \bar{w}, \bar{x})$ as a real valued formula. (Although note that it does still depend

on $D(\bar{x}, \bar{v})$ as a real valued formula and that this cannot generally be avoided as alluded to in Remark 1.7.6.)

With regards to open formulas, since $\text{DEF}_{\bar{v}}D(\bar{v}, \bar{x})$ is a closed formula, we can't say ' $D(\bar{v}, \bar{x})$ is definable in \bar{v} and for all \bar{v} in $D(\bar{x})$, strongly, $\varphi(\bar{v}, \bar{w}, \bar{x}) > 0$ holds' with a topological formula; the closest we can get is 'if $D(\bar{v}, \bar{x})$ is definable, then for all \bar{v} in $D(\bar{x})$, strongly, $\varphi(\bar{v}, \bar{w}, \bar{x}) > 0$ holds,' which would be

$$\text{DEF}_{\bar{v}}D(\bar{v}, \bar{x}) \rightarrow \left(\inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi(\bar{v}, \bar{w}, \bar{x}) > 0 \right) \vee \Lambda_o(\bar{x}).$$

Again the meaning of this expression does not depend on $\varphi(\bar{v}, \bar{w}, \bar{x})$ as a real valued formula, only on $(\varphi(\bar{v}, \bar{w}, \bar{x}) > 0)$ as an open formula. That said, notation of the form $(\forall \bar{v} \in D(\cdot, \bar{x}) \downarrow)(\varphi(\bar{v}, \bar{w}, \bar{x}) > 0)$ does not clearly suggest this meaning.

Definition 2.3.27. If $D(\bar{v}, \bar{x})$ and $E(\bar{v}, \bar{x})$ are \bar{x} -uniformly definable over $\Sigma(\bar{x})$, then the *Hausdorff distance between D and E* is given by $d_H(D(\bar{x}), E(\bar{x})) = \sup_{\bar{v}} |D(\bar{v}, \bar{x}) - E(\bar{v}, \bar{x})|$. This is a real formula with free variables among \bar{x} . If \bar{x} is empty we will write $d_H(D, E)$.

The *logical Hausdorff distance between D and E over Σ* is given by $d_{H, \Sigma}(D(\bar{x}), E(\bar{x})) = \|D(\bar{v}, \bar{x}) - E(\bar{v}, \bar{x})\|_{\Sigma}$. When Σ is empty we will write d_H .⁵ ◁

Note that in particular, $\text{fv}(D) \setminus \bar{v}$ and $\text{fv}(E) \setminus \bar{v}$ may be disjoint.

⁵Note that the first term defined here is a formula and the second is the supremal value of this formula.

Proposition 2.3.28. *For any partial type $\Sigma(\bar{x})$ and $D(\bar{v}, \bar{x})$ and $E(\bar{v}, \bar{x})$ which are \bar{x} -uniformly definable over $\Sigma(\bar{x})$,*

$$d_H(D(\cdot, \bar{x}), E(\cdot, \bar{x})) = \max \left\{ \sup_{\bar{v} \in D(\cdot, \bar{x})} \inf_{\bar{w} \in E(\cdot, \bar{x})} d(\bar{v}, \bar{w}), \sup_{\bar{w} \in E(\cdot, \bar{x})} \inf_{\bar{v} \in D(\cdot, \bar{x})} d(\bar{v}, \bar{w}) \right\} \text{ and}$$

$$d_{H,\Sigma}(D(\cdot, \bar{x}), E(\cdot, \bar{x})) = \sup \{d_H(D(\cdot, p), E(\cdot, p)) : p \in S_{\bar{x}}(\Sigma)\}.$$

In particular, if \mathfrak{M} is a structure such that $\mathfrak{M} \models \Sigma(\bar{a})$, then $d_H^{\mathfrak{M}}(D(\cdot, \bar{a}), E(\cdot, \bar{a}))$ is the Hausdorff distance between $D(\mathfrak{M}, \bar{a})$ and $E(\mathfrak{M}, \bar{a})$.

Proof. The equality regarding $d_{H,\Sigma}$ follows by definition, and the final statement follows from the equality regarding d_H and the definition of the Hausdorff distance, so we only need to show the equality regarding d_H . By Fact A.1.6 we know that the first equality holds in any structure \mathfrak{M} with $\bar{A} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$. Therefore the equality holds. \square

Proposition 2.3.29. *For any partial type $\Sigma(\bar{x})$ and any at most countable tuple of variables \bar{v} , the set $\{D(\bar{v}, \bar{x}) : D \text{ } \bar{x}\text{-uniformly definable over } \Sigma(\bar{x})\}$ is metrically complete under $d_{H,\Sigma}$ and has metric density character at most $\aleph_0 + |\mathcal{L}| + |\bar{x}|$. Furthermore, $d_{H,\Sigma}(D(\cdot, \bar{x}), E(\cdot, \bar{x})) = 0$ if and only if D and E are logically equivalent modulo Σ .*

Proof. Let $\{D_i(\bar{v}, \bar{x})\}_{i < \omega}$ be a sequence of \bar{x} -uniformly definable sets over $\Sigma(\bar{x})$ that is a Cauchy sequence in the metric $d_{H,\Sigma}$. This implies that the corresponding distance predicates are also a Cauchy sequence relative to $\|\cdot\|_{\Sigma}$. So by Proposition 1.4.6 we have that there is a real formula $\varphi(\bar{v}, \bar{x})$ such that $\{D_i(\bar{v}, \bar{x})\}_{i < \omega}$ converges to φ as a sequence of distance predicates under $\|\cdot\|_{\Sigma}$. Now we just need to argue that $\varphi(\bar{v}, \bar{x})$ is the distance predicate of a \bar{x} -uniformly definable set over $\Sigma(\bar{x})$. By Fact A.1.6, in any \mathfrak{M} with $\bar{a} \in \mathfrak{M}$

such that $\mathfrak{M} \models \Sigma(\bar{a})$, the sequence $D_i(\mathfrak{M}, \bar{a})$ converges to some closed set $F \subseteq \mathfrak{M}^{|\bar{v}|}$. By Fact A.1.6 again, $\varphi^{\mathfrak{M}}(\bar{v}, \bar{a})$ must be the distance predicate of F . Therefore we have that $\Sigma(\bar{x}) \models \text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{x})$, as required.

The statement regarding metric density character follows from the fact that the collection of distance predicates of \bar{x} -uniformly definable sets over Σ is a subspace of the collection of all $\mathcal{L}(\bar{v}\bar{x})$ -formulas, which has metric density character at most $\aleph_0 + |\mathcal{L}| + |\bar{v}\bar{x}| = \aleph_0 + |\mathcal{L}| + |\bar{x}|$. The final statement regarding logical equivalence follows by the definition of logical equivalence modulo Σ . \square

Corollary 2.3.30 (Union of Chain of Definable Sets). *Fix a type $\Sigma(\bar{x})$, and let $\{D_i(\bar{v}, \bar{x})\}_{i < \omega}$ be a sequence of \bar{x} -uniformly definable sets over $\Sigma(\bar{x})$ such that for each i , $\Sigma(\bar{x}), D_i(\bar{v}, \bar{x}) \models D_{i+1}(\bar{v}, \bar{x})$. The following are equivalent:*

- (i) *For any \mathfrak{M} and $\bar{b} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{b})$, the set $\overline{\bigcup_{i < \omega} D_i(\mathfrak{M}, \bar{b})}$ is definable by some distance predicate $\varphi(\bar{v}, \bar{b})$.*
- (ii) *The set of types $\overline{\bigcup_{i < \omega} \llbracket D_i(\bar{v}, \bar{x}) \rrbracket_{\Sigma(\bar{x})}}$ is definable in $(S_{\bar{v}\bar{x}}(\Sigma), d_{/\bar{x}})$.*
- (iii) *The set of types $\overline{\bigcup_{i < \omega} \llbracket D_i(\bar{v}, \bar{x}) \rrbracket_{\Sigma(\bar{x})}}$ is type-definable in $(S_{\bar{v}\bar{x}}(\Sigma), d_{/\bar{x}})$.*
- (iv) *$\lim_{i \rightarrow \infty} \sup\{r : \Sigma(\bar{x}) \cup \{d_H(D_i(\bar{x}), D_{i+1}(\bar{x})) \geq r\}$ is consistent} = 0, i.e. $\{D_i(\bar{v}, \bar{x})\}_{i < \omega}$ is a Cauchy sequence in the Hausdorff metric over $\Sigma(\bar{x})$.*

Proof. (iv) \Rightarrow (ii) follows from Proposition 2.3.29, and (ii) \Rightarrow (i) is obvious.

(i) \Rightarrow (iv). Assume that (i) succeeds but (iv) fails. This implies that there is an $\varepsilon > 0$ such that for infinitely many $i < \omega$, $\Sigma(\bar{x}) \cup \{d_H(D_i(\bar{x}), D_{i+1}(\bar{x})) \geq \varepsilon\}$ is consistent. Let $i(n)$ be an enumeration of this subsequence. Since the chain of definable sets is nested, we actually have $\Sigma(\bar{x}) \cup \{d_H(D_{i(n)}(\bar{x}), D_{i(n+1)}(\bar{x})) \geq \varepsilon\}$ is

consistent for every $n < \omega$. For each $n < \omega$, find \mathfrak{M}_n and $\bar{a}_n \bar{b}_n \in \mathfrak{M}_n$ such that $\mathfrak{M}_n \models \Sigma(\bar{b}_n) \wedge D_{i(n+1)}(\bar{a}_n, \bar{b}_n) \wedge D_i(\bar{a}_n, \bar{b}_n) \geq \varepsilon$ (which exist by compactness and the fact that $D_i(\bar{x})$ is a subset of $D_{i+1}(\bar{x})$ over $\Sigma(\bar{x})$). Let $(\mathfrak{M}_\mathcal{U}, \bar{a}_\mathcal{U}, \bar{b}_\mathcal{U})$ be a non-trivial ultraproduct of the sequence $\{\mathfrak{M}_n, \bar{a}_n, \bar{b}_n\}_{n < \omega}$. By (i), there is a formula $\varphi(\bar{v}, \bar{b}_\mathcal{U})$ that is a distance predicate for $\overline{\bigcup_{i < \omega} D_i(\mathfrak{M}_\mathcal{U}, \bar{b}_\mathcal{U})} = \overline{\bigcup_{n < \omega} D_{i(n)}(\mathfrak{M}_\mathcal{U}, \bar{b}_\mathcal{U})}$. This implies that $\text{cl} \bigcup_{n < \omega} \llbracket D_{i(n)}(\bar{v}, \bar{b}_\mathcal{U}) \rrbracket \subseteq \llbracket \varphi(\bar{v}, \bar{b}_\mathcal{U}) = 0 \rrbracket$. So in particular we must have

$$\text{tp}(\bar{a}_\mathcal{U} \bar{b}_\mathcal{U}) \in \text{cl} \bigcup_{n < \omega} \llbracket D_{i(n)}(\bar{v}, \bar{b}_\mathcal{U}) \rrbracket \subseteq \llbracket \varphi(\bar{v}, \bar{b}_\mathcal{U}) = 0 \rrbracket,$$

and hence $\mathfrak{M} \models \varphi(\bar{a}_\mathcal{U}, \bar{b}_\mathcal{U}) = 0$. But by construction,

$$d_{\text{inf}}^{\mathfrak{M}_\mathcal{U}} \left(\bar{a}_\mathcal{U}, \overline{\bigcup_{n < \omega} D_{i(n)}(\mathfrak{M}_\mathcal{U}, \bar{b}_\mathcal{U})} \right) \geq \varepsilon,$$

so in particular, $\mathfrak{M} \models \varphi(\bar{a}_\mathcal{U}, \bar{b}_\mathcal{U}) \geq \varepsilon$, which is a contradiction. Therefore we must have that (i) \Rightarrow (iv), as required.

(ii) \Leftrightarrow (iii). The forward direction is obvious. The backward direction follows from the fact that the metric closure of a locatable set is locatable. \square

2.3.3 Examples: The Discrete Interval and the Polarized Square

At this point it will be instructive to examine a couple of examples of the issues surrounding locatable and definable sets in detail. Some of the issues here are relevant to Section 2.4, but none of the results of that section are necessary to understand this example.

The first example shows how a theory in continuous logic can have a non-trivial type




Figure 4: $S_1(\text{DI})$ with a cartoon of its ‘microscopic’ metric topology.

space that fails to have any non-trivial definable sets.

Definition 2.3.31. The *discrete interval* is the type space $S_1(\text{DI})$ of the theory of the following structure:

Let \mathcal{L} be the signature with a single unary 1-Lipschitz predicate P with $I(P) = [0, 1]$ and $\text{db}(\mathcal{L}) = 1$. Let \mathfrak{D} be the \mathcal{L} -structure whose universe is $[0, 1]$ with $P^{\mathfrak{D}}(x) = x$ and $d^{\mathfrak{D}}(x, y) = 1$ if and only if $x \neq y$. \triangleleft

Proposition 2.3.32. $S_1(\text{DI})$ is homeomorphic to $[0, 1]$, but the only definable subsets of $S_1(\text{DI})$ are \emptyset and $S_1(\text{DI})$.

Proof. A direct computation shows that if \mathcal{U} is a non-principal ultrafilter on ω , then $\mathfrak{D}^{\mathcal{U}}$ is isomorphic to $\mathfrak{S} \times 2^{\aleph_0}$, with the distance between points in distinct copies of \mathfrak{S} equal to 1. Since this is a non-trivial ultrapower and \mathcal{L} is countable, $\mathfrak{D}^{\mathcal{U}}$ realizes all types in $S_1(\text{DI})$. It is easy to see that the automorphism type of any $a \in \mathfrak{D}^{\mathcal{U}}$ depends only on $P^{\mathfrak{D}^{\mathcal{U}}}(a)$, so we have that for any $a, b \in \mathfrak{D}^{\mathcal{U}}$, that $a \equiv b$ if and only if $P^{\mathfrak{D}^{\mathcal{U}}}(a) = P^{\mathfrak{D}^{\mathcal{U}}}(b)$. Therefore $S_1(\text{DI})$ is homeomorphic to $[0, 1]$. Furthermore, we have that $d^{tp}(a, b) = 1$ if and only if $a \neq b$.

Let $F \subseteq S_1(\text{DI})$ be a set. For any $0 < \varepsilon < 1$, $F^{<\varepsilon} = F$. This implies that a set can only be locatable if it is open, and therefore a set can only be definable if it is clopen, but the only clopen sets are \emptyset and $S_1(\text{DI})$. \square

We haven't defined the concept yet, but it turns out that DI is superstable. In fact, more is true.

Proposition 2.3.33. *If T is a non- ω -stable discrete theory, then T interprets DI.*

Proof. Find a sequence of formulas $\varphi_i(x)$ such that for any $\sigma \in 2^{<\omega}$, $\bigwedge_{i < |\sigma|} \pm_{\sigma(i)} \varphi_i(x)$ does not have ordinal Morley rank, where $\pm_1 = \neg$ and \pm_0 is the empty string. Interpret each φ_i as a $\{0, 1\}$ -valued formula, with $\varphi_i(a) = 0$ if and only if $\models \varphi_i(a)$. Now consider $\psi(x) = \sum_{i < \omega} 2^{-i-1} \varphi_i(x)$. The reduct in which we forget everything except $=$ and ψ is inter-definable with DI. \square

This gives us some clue that no stability assumption weaker than ω -stability is going to ensure the presence of non-trivial definable sets.

The second example, the definiendum of the following Definition 2.3.34, is useful for seeing non-trivial examples of different phenomena surrounding locatable and definable sets (see Figure 6). It should also be noted that this second example is actually bi-interpretable with the first.

Definition 2.3.34. The *polarized square* is the type space $S_1(\text{PS})$ of the theory PS of this structure:

Let \mathcal{L} contain two 1-Lipschitz unary predicates, P and Q , with $I(P) = I(Q) = [0, 1]$, and let $\text{db}(\mathcal{L}) = 1$.

Let \mathfrak{S} be an \mathcal{L} -structure whose universe is $[0, 1]^2$. For each $(x, y) \in \mathfrak{S}$, let $P^{\mathfrak{S}}((x, y)) = x$ and $Q^{\mathfrak{S}}((x, y)) = y$. Finally, for $(x, y), (z, w) \in \mathfrak{S}$, we set $d^{\mathfrak{S}}((x, y), (z, w)) = 1$ if $x \neq z$ and $d^{\mathfrak{S}}((x, y), (z, w)) = |y - w|$ if $x = z$. \triangleleft

Proposition 2.3.35. *$S_1(\text{PS})$ is homeomorphic to $[0, 1]^2$. If we let $f : [0, 1]^2 \rightarrow S_1(\text{PS})$*




Figure 5: $S_1(\text{PS})$ with a cartoon of its ‘microscopic’ metric topology.

be this homeomorphism, then $\text{tp}((x, y)) = f((x, y))$ for any $(x, y) \in \mathfrak{S}$. In particular, every type in $S_1(\text{PS})$ is realized in \mathfrak{S} .

Furthermore, for any $(x, y), (z, w) \in \mathfrak{S}$, $d^{\text{tp}}((x, y), (z, w)) = d^{\mathfrak{S}}((x, y), (z, w))$.

Proof. This proof is almost exactly the same as the proof of Proposition 2.3.32. \square

We will represent $\text{tp}_{\mathfrak{S}}((x, y))$ by $p_{x,y}$. We will often conflate $S_1(\text{PS})$ and $[0, 1]^2$.

The point of this example is that it has a lot of definable sets, but it doesn’t have enough definable sets.

Proposition 2.3.36. For any continuous function $f : [0, 1] \rightarrow [0, 1]$, the set

$$\llbracket Q(x) = f(P(x)) \rrbracket$$

is definable.

Proof. The distance predicate of that set is $|Q(x) - f(P(x))|$. \square




Figure 6: A locatable subset L of $S_1(\text{PS})$ that is not definable and has empty interior, together with $L^{<\frac{1}{4}}$. (Note that the endpoints are missing.)

This gives us a good example of two definable sets whose intersection is not definable, namely $\llbracket P(x) = Q(x) \rrbracket$ and $\llbracket P(x) = 1 - Q(x) \rrbracket$. Their intersection is $\{p_{1/2,1/2}\}$, which is not a definable set.

Proposition 2.3.37. *For any definable set $D \subseteq S_1(\text{PS})$ and any $\varepsilon > 0$, the set $D^{\leq \varepsilon}$ is definable as well.*

Proof. If D is empty then this is obvious, otherwise if D is non-empty then for any $p \in S_1(\text{PS})$, we have that $d_{\text{inf}}(p, D^{\leq \varepsilon}) = d_{\text{inf}}(p, D) \div \varepsilon$. \square

Corollary 2.3.38. *For any $p_{a,b} \neq p_{c,e}$, if b and e are not both equal to 0 and not both equal to 1, then there exist disjoint definable sets $D_0, D_1 \subseteq S_1(\text{PS})$ such that $p_{a,b} \in \text{int } D_0$ and $p_{c,e} \in \text{int } D_1$.*

Proof. Assume without loss that $b \geq e$. Consider $f(x) = [b + |x - a|]_0^1$ and $g(x) = [e - |x - c|]_0^1$. Clearly $p_{a,b} \in \llbracket Q = f(P) \rrbracket$ and $p_{c,e} \in \llbracket Q = g(P) \rrbracket$ and these sets




Figure 7: Separating $p_{a,b}$ and $p_{c,e}$ with definable neighborhoods (2.3.38).

are disjoint (this relies on the fact that b and e are not both 0 and not both 1). By Propositions 2.3.36 these sets are definable. By compactness there is some $\varepsilon > 0$ such that $D_0 = \llbracket Q = f(P) \rrbracket^{\leq \varepsilon}$ and $D_1 = \llbracket Q = g(P) \rrbracket^{\leq \varepsilon}$ are disjoint. By Proposition 2.3.37, these are the required definable sets. \square

Corollary 2.3.39. $S_1(\text{PS})$ has a basis of co-definable sets.

Proof. For any $\langle a, b \rangle \in [0, 1] \times (0, 1)$, for $\varepsilon > 0$ small enough that

$$B_{x,y}^\varepsilon = \{ \langle x, y \rangle \in [0, 1]^2 : \sqrt{(x-a)^2 + (y-b)^2} < \varepsilon \} \subseteq [0, 1] \times (0, 1),$$

we have that $B_{x,y}^\varepsilon$ is a co-definable neighborhood. Specifically, $[0, 1]^2 \setminus B_{x,y}^\varepsilon$ is the union of an open set and the graphs of the functions $f(x)$ and $g(x)$, where $f(x) = g(x) = b$ if $|x-a| \geq \varepsilon$ and $f(x) = b + \sqrt{\varepsilon^2 - (x-a)^2}$ and $g(x) = b - \sqrt{\varepsilon^2 - (x-a)^2}$, which is necessarily a locatable set. Since it is also closed, it is a definable set.

For $\langle a, b \rangle \in [0, 1] \times \{0, 1\}$, a similar argument gives that $B_{x,y}^\varepsilon$ is co-definable for any

$\varepsilon < 1$. Finally, sets of these forms form a basis of the topology on $S_1(\mathbf{PS})$. \square

Despite this prevalence of definable and co-definable sets, $S_1(\mathbf{PS})$ does not have what we would really want, a basis or even a network of definable neighborhoods (which we will see in Section 2.4 are actually equivalent), i.e. it is not true that for every p and open neighborhood U , there exists a definable set D such that $p \in D \subseteq U$.

Proposition 2.3.40. *A set $R \subseteq S_1(\mathbf{PS})$ is locatable if and only if for every $p_{x,y} \in R$ and every neighborhood $U \ni p_{x,y}$ the set $\{z \in [0, 1] : (\exists w)p_{z,w} \in R \cap U\}$ contains an open neighborhood of x in $[0, 1]$.*

Proof. For the \Leftarrow direction, pick $\varepsilon > 0$ and $p_{x,y} \in R$. Let

$$U = \left(\left(x - \frac{\varepsilon}{2}, x + \frac{\varepsilon}{2} \right) \times \left(y - \frac{\varepsilon}{2}, y + \frac{\varepsilon}{2} \right) \right) \cap ([0, 1] \times [0, 1]).$$

By assumption, for some open-in- $[0, 1]$ set $V \ni x$, for every $z \in V$, there is w such that $p_{z,w} \in U$. This implies that $R^{<\varepsilon}$ contains $V \times \left(y - \frac{\varepsilon}{2}, y + \frac{\varepsilon}{2} \right) \ni p_{x,y}$, therefore $p_{x,y} \in \text{int } R^{<\varepsilon}$. Since we can do this for any $p_{x,y} \in R$ and $\varepsilon > 0$, we have that R is locatable.

For the \Rightarrow direction, assume that R is locatable. Fix $p_{x,y} \in R$ and $U \ni p_{x,y}$, an open neighborhood. Find $\varepsilon > 0$ small enough that $([x - \varepsilon, x + \varepsilon] \times [y - \varepsilon, y + \varepsilon]) \cap ([0, 1] \times [0, 1]) = V \subseteq U$. By assumption $p_{x,y} \in \text{int } R^{<\varepsilon}$. Therefore there is some $\delta > 0$ such that $([x - \delta, x + \delta] \cap [0, 1]) \times \{y\} \subseteq V \cap \text{int } R^{<\varepsilon}$. By definition, this implies that for every $z \in [x - \delta, x + \delta] \cap [0, 1]$, there is w such that $p_{z,w} \in R^{<\varepsilon}$ and $|y - w| < \varepsilon$. Therefore, by definition, $p_{z,w} \in V$. Therefore $\{z \in [0, 1] : (\exists w)p_{z,w} \in R \cap U\}$ contains all of $[x - \delta, x + \delta] \cap [0, 1]$ and hence contains an open neighborhood of x in $[0, 1]$. \square

Corollary 2.3.41. *If $D \subseteq S_1(\text{PS})$ is non-empty and definable, then for any $x \in [0, 1]$ there is $p_{x,y} \in D$ for some $y \in [0, 1]$.*

Proof. Assume that D has no element $p_{x,y}$ for some x . Then by compactness, there is some $\varepsilon > 0$ such that $D \cap \{p_{z,w} : |x - z| < \varepsilon\} = \emptyset$, furthermore, there is a maximum such ε . Let ε be maximal, so we have that $D \cap \{p_{z,w} : |x - z| = \varepsilon\}$ is non-empty, but by Proposition 2.3.40, this implies that D is not locatable, and therefore not definable either, so we have a contradiction. Therefore no such x can exist. \square

Corollary 2.3.42. *There is no definable set $D \subseteq S_1(\text{PS})$ such that $p_{0,0} \in D$ and D is disjoint from $\{p_{1,x} : x \in [0, 1]\}$.*

Corollary 2.3.43. *For any $b, e \in [0, 1]$, any definable sets D, E with $p_{0,b} \in D$ and $p_{0,e} \in E$ fail to be disjoint. The same holds for $p_{1,b}$ and $p_{1,e}$.*

Proof. If $b = e$ we are done, so assume that this is not the case. Consider the functions $f(x) = \inf\{a \in [0, 1] : p_{a,x} \in D\}$ and $g(x) = \inf\{c \in [0, 1] : p_{c,x} \in E\}$. These functions are well defined by Corollary 2.3.41. It follows from the fact that D and E are closed that these functions are continuous. The required result then follows from the intermediate value theorem.

The argument for $p_{1,b}$ and $p_{1,e}$ is much the same. \square

Corollary 2.3.44. *Distinct types $p_{a,b}, p_{c,e}$ can be separated by disjoint definable sets if and only if a and c are not both 0 and are not both 1.*

Proof. This follows from Corollaries 2.3.38 and 2.3.43. \square

So we see that the question of whether or not a pair of types can be separated by definable sets can be very non-local.

Characterizing the definable sets in $S_1(\mathbf{PS})$ explicitly is trickier. From Proposition 2.3.36 we know that closed sets that are unions of the graphs of continuous functions are definable; the converse of this is only approximately true.

Definition 2.3.45. For any continuous $f : [0, 1] \rightarrow [0, 1]$ and $\varepsilon > 0$, the *tubular neighborhood of f of radius ε* , written $U_{f,\varepsilon}$, is $\{(x, y) \in [0, 1]^2 : |y - f(x)| < \varepsilon\}$. Sets of this form are called *tubular neighborhoods*. An open set is *totally tubular* if it is a (possibly empty) union of tubular neighborhoods. \triangleleft

Proposition 2.3.46. For any definable set D , and any $\varepsilon > 0$, $D^{<\varepsilon}$ is totally tubular.

Proof. If D is empty, then $D^{<\varepsilon}$ is either empty or all of $S_1(\mathbf{PS})$ and the result is trivial. Assume that D is non-empty.

It is sufficient to show that for any $p \in D^{<\varepsilon}$, there is a tubular neighborhood $U_{f,\delta}$ such that $p \in U_{f,\delta} \subseteq D^{<\varepsilon}$. Fix $\delta > 0$ with $\delta < \varepsilon$. By Proposition 2.3.37, we have that $D^{\leq\delta}$ is a definable set, so Proposition 2.3.40 applies to $D^{\leq\delta}$.

Let C be the collection of all sets of the form $B_{z,w}^r = \{p_{u,v} \in S_1(\mathbf{PS}) : (z - u)^2 + (w - v)^2 < r\}$ such that r is some positive real, $p_{z,w} \in D^{\leq\delta}$, and $B_{z,w}^r \subseteq D^{<\varepsilon}$.

Find $s > 0$ such that $B_{x,y}^s \subseteq D^{<\varepsilon}$. If $x = 1$, skip to (), otherwise build a tree whose nodes are labeled by elements of C by the following. Let the root node be $B_{x,y}^s$. Let the children of $B_{z,w}^r$ be the set of all $B_{u,v}^t$ such that $z < u$ and $B_{z,w}^r \cap B_{u,v}^t \neq \emptyset$. Note that in particular, if $z = 1$, then $B_{z,w}^r$ has no children.

(*) Note that, by Proposition 2.3.40, if $z < 1$ and $B_{z,w}^r$ is on the tree, then it has children.

Let $t = \sup\{z : (\exists w, r) B_{z,w}^r \text{ is on the tree}\}$. We want to argue that $t = 1$. Let p_{z_i, w_i} be a sequence of points that are the centers of balls on the tree such that $z_i \rightarrow t$. By

compactness this has a convergent subsequence, converging to some $p_{u,v} \in D^{\leq \delta}$. We can find $t > 0$ such that $B_{u,v}^t \subseteq D^{< \varepsilon}$. Therefore for some i , the ball on the tree whose center was p_{z_i, w_i} has $B_{u,v}^t$ as a child, so by (*) we have a contradiction. The same argument shows that there is some ball on the tree of the form $B_{1,w}^r$.

Let B_0, B_1, \dots, B_n be a path on the tree with $B_0 = B_{x,y}^s$ and $B_n = B_{1,w}^r$. Let $g : [x, 1] \rightarrow [0, 1]$ be the piecewise linear function interpolating between the centers of B_0, B_1, \dots, B_n . Clearly by construction we have that the graph of g is contained in $D^{\leq \varepsilon}$.

(♣) If $x > 0$, by the symmetric procedure we can get $h : [0, x] \rightarrow [0, 1]$, a continuous function satisfying $h(x) = y$ and whose graph is contained in $D^{\leq \varepsilon}$. By setting $f = g \cup h$ (or just g or h if only one of these exists), we have a continuous function $f : [0, 1] \rightarrow [0, 1]$ such that $f(x) = y$. Now, by compactness, there must be some $\gamma > 0$ such that $U_{f,\gamma} \subseteq D^{\leq \varepsilon}$. Since we can do this for every $\delta < \varepsilon$, we have that every $p \in D^{< \varepsilon}$ is contained in some tubular neighborhood contained in $D^{< \varepsilon}$, and hence $D^{< \varepsilon}$ is totally tubular. \square

Lemma 2.3.47. *An open set $V \subseteq [0, 1]^2$ is totally tubular if and only if for every $p \in V$, there exists a continuous function $f : [0, 1] \rightarrow [0, 1]$ whose graph contains p and is contained in V .*

Proof. (\Rightarrow) Assume that V is totally tubular. For any $p_{z,w} \in V$, let $U_{f,\varepsilon}$ be chosen so that $p_{z,w} \in U_{f,\varepsilon} \subseteq V$. Now we clearly have that $p_{z,w}$ is contained in the graph of $g(x) = [f(x) + y - f(z)]_0^1$. Since $|f(x) - g(x)| \leq |y - f(z)| < \varepsilon$ for every $x \in [0, 1]$, we have that the graph of g is contained in $U_{f,\varepsilon}$ and therefore also contained in V as required.

(\Leftarrow) Assume that V is an open set that is equal to a union of graphs of functions.

For any $p_{z,w} \in V$, let $f : [0, 1] \rightarrow [0, 1]$ be a continuous function whose graph contains $p_{z,w}$ and is contained in V . By compactness, there is some $\varepsilon > 0$ such that $U_{f,\varepsilon} \subseteq V$, therefore we have that V is the union of such sets and is totally tubular. \square

Corollary 2.3.48. *A closed set $D \subseteq S_1(\text{PS})$ is definable if and only if there exists a sequence $\{F_i\}_{i < \omega}$ limiting to D in the Hausdorff metric such that each F_i is a finite union of graphs of continuous functions.*

Proof. (\Rightarrow) From Proposition 2.3.46 and Lemma 2.3.47 we have that for any $i < \omega$, D is covered by the graphs of continuous functions contained in $D^{<2^{-i}}$. By compactness, this implies that D is covered by $U_{f_0,2^{-i}}, \dots, U_{f_{n-1},2^{-i}}$ for some finite list of continuous functions f_0, \dots, f_{n-1} . Let F_i be the union of the graphs of these functions. By construction we have that $d_H(D, F_i) \leq 2^{-i}$ (actually $< 2^{-i}$ by compactness).

(\Leftarrow) Any such F_i are definable, so as a closed Hausdorff metric limit of definable sets, D is definable. \square


Despite this, it is not true that every definable $D \subseteq S_1(\text{PS})$ is a union of the graphs of continuous functions.

Proposition 2.3.49. *Let*

$$D = \{p_{x,y} : x = 0 \vee y = 1\} \cup \bigcup_{0 < n < \omega} \llbracket Q = \min\{nP, 1\} \rrbracket.$$

D is a definable set that is not the union of graphs of continuous functions.

Proof. $L = \{p_{x,y} : y = 1\} \cup \bigcup_{0 < n < \omega} \llbracket Q = \min\{nP, 1\} \rrbracket$ is a union of graphs of continuous functions, and so is locatable. All we need to do is show that for every $\varepsilon > 0$, $\{p_{x,y} : x = 0\} \cup L^{<\varepsilon}$ contains an open neighborhood of $\{p_{x,y} : x = 0\}$. For any $\varepsilon > 0$, we have

Figure 8: D from Proposition 2.3.49

that $\{p_{x,y} : x \in (0, 2\varepsilon)\} \subseteq L^{<\varepsilon}$, so $\{p_{x,y} : x = 0\} \cup L^{<\varepsilon}$ contains an open neighborhood of $\{p_{x,y} : x = 0\}$, thus we have that D is definable.

Let $f : [0, 1] \rightarrow [0, 1]$ be a continuous function whose graph is a subset of D . Clearly we have $f(1) = 1$. For each $0 < n < \omega$, if for some $x \in (\frac{1}{n+1}, \frac{1}{n})$, $f(x) < 1$, then $f(x)$ must be equal to $\min\{kx, 1\}$ for some $k \leq n$. Otherwise if that never happens, then $f(x) = 1$.

So in particular, we have that no graph of a continuous function that is a subset of D contains $p_{0,1/2}$, which is an element of D . Hence D is not a union of graphs of continuous functions. \square

The example in Proposition 2.3.49 still has the property that it is the topological closure of a union of graphs of continuous functions. It is easy to construct an example of a family of continuous functions for which the closure of their union is not a definable set, but there is still the question of the converse.

Question 2.3.50. *Is every definable $D \subseteq S_1(\text{PS})$ the topological closure of a union of*

graphs of continuous functions?

On the other hand, it is possible to completely characterize the minimal non-empty definable sets.

Lemma 2.3.51. *For any non-empty definable set $D \subseteq S_1(\text{PS})$, the function $f(x) = \sup\{y : p_{x,y} \in D\}$ is continuous and its graph is a subset of D .*

Proof. First note that by Corollary 2.3.41, for any x , the set in the definition of $f(x)$ is non-empty.

Fix $x \in [0, 1]$, and let $\{x_i\}_{i < \omega}$ be a sequence of elements of $[0, 1]$ limiting to x . By the compactness of D we have that $f(x) \geq \lim_{i \rightarrow \infty} f(x_i)$. Conversely, by Proposition 2.3.40 we have that for any $\varepsilon > 0$, there is an open neighborhood $U \ni x$ such that for all $z \in U$, $f(z) \geq f(x) - \varepsilon$. Therefore it must be the case that $f(x) = \lim_{i \rightarrow \infty} f(x_i)$, and $f(x)$ is continuous. \square

Proposition 2.3.52. *A non-empty definable set $D \subseteq S_1(\text{PS})$ is minimal (i.e. has no non-empty proper subsets that are definable) if and only if it is the graph of a continuous function $f : [0, 1] \rightarrow [0, 1]$.*

Proof. (\Rightarrow) By Lemma 2.3.51, if D is non-empty but also isn't the graph of a definable function, then it has a proper subset that is the graph of a definable function, thus D is not minimal.

(\Leftarrow) By Corollary 2.3.41, no non-empty proper subset of the graph of a definable function can be definable. \square

Turning our attention to arbitrary locatable sets, we know that any union of sets of the form $D \cap U$, with D definable and U open, is itself locatable, but the converse does not hold.

Proposition 2.3.53. *There is a locatable set $L \subseteq S_1(\text{PS})$ that does not contain any non-empty sets of the form $D \cap U$, with D definable and U open.*

Proof. There are continuum many sets that are intersections of a definable set and an open set. Let $\{X_i\}_{i < 2^{\aleph_0}}$ be an enumeration of them. Construct L by transfinite induction. For each X_i , if X_i is non-empty, then there are continuum many $x \in [0, 1]$ such that for some y , $p_{x,y} \in X_i$. Find an $x \in [0, 1]$ that has not been used yet such that for some y , $p_{x,y} \in X_i$. Let (x_i, y_i) be this x and y .

Finally $L = S_1(\text{PS}) \setminus \{p_{x_i, y_i}\}_{i < 2^{\aleph_0}}$. For each $x \in [0, 1]$, L contains all but at most one $p_{x,y}$, so $\bar{L} = S_1(\text{PS})$ and L is locatable, but by construction if $D \cap U \subseteq L$ for definable D and open U , then $D \cap U$ is empty. \square


Of course, in some sense this example isn't very interesting. The real question is whether or not every metrically complete or just G_δ locatable set can be written as a union of sets of the form $D \cap U$, with D definable and U open. The answer to this is also no.

Proposition 2.3.54. *There is a metrically complete locatable set $L \subseteq S_1(\text{PS})$ that is not the union of sets of the form $D \cap U$ with D definable and U open.*

Proof. Let

$$L = \{\langle 0, 0 \rangle\} \cup \bigcup_{i < \omega} (2^{-i-1}, 2^{-i+1}) \times \{2^{-i}\}.$$

$\bigcup_{i < \omega} (2^{-i-1}, 2^{-i+1}) \times \{2^{-i}\}$ is clearly locatable. For any $i < \omega$, we want to show that $[0, 2^{-i-1}) \times [0, 2^{-i-1}) \subseteq \text{int } L^{< 2^{-i}}$. Let $\langle x, y \rangle$ be an element of $[0, 2^{-i-1}) \times [0, 2^{-i-1})$. If $x = 0$, then we have $d(\langle x, y \rangle, \langle 0, 0 \rangle) < 2^{-i-1} < 2^{-i}$. On the other hand if $x > 0$, then $x < 2^{-i-1}$, so there is some $j > i + 1$ such that $x \in (2^{-j-1}, 2^{-j+1})$. This implies that

Figure 9: L from Proposition 2.3.54

$d_{\text{inf}}(\langle x, y \rangle, (2^{-j-1}, 2^{-j+1}) \times \{2^{-j}\}) \leq 2^{-j} < 2^{-i-1}$. So in either case we have $\langle x, y \rangle \in L^{<2^{-i}}$.

Therefore L is locatable.

Now we want to show that if $D \cap U$, with D definable and U open, is a subset of L , then $\langle 0, 0 \rangle \notin D \cap U$. If $\langle 0, 0 \rangle$ is an element of a definable set D , then it is also an element of the graph of $f(x) = \inf\{y : \langle x, y \rangle \in D\}$. So if $\langle 0, 0 \rangle \in D \cap U \subseteq L$, this would imply that $\langle 0, 0 \rangle$'s connected component in L is non-trivial, but it is trivial, so no such D and U can exist. \square

But unlike the example in Proposition 2.3.53, the example in Proposition 2.3.54 contains subsets of the form $D \cap U$ with D definable and U open, so a natural question is this:

Question 2.3.55. *Does every G_δ locatable set in $S_1(\text{PS})$ contain a non-empty subset of the form $D \cap U$ with D definable and U open?*

2.3.4 Finite Counting Quantifiers and Algebraic Sets

There are many ways to ‘approximately count’ the number of elements of a subset of a metric space at a given scale $\varepsilon > 0$, but only a few of them are natural for open and closed formulas.

For the sake of simplicity, we will write these definitions for single variable formulas only. The extension to multiple variables and parameters is obvious.

Definition 2.3.56. Let $F(x)$ be a closed type-set formula, $L(x)$ be a locatable type-set formula (such as in particular an open formula or a definable set), and $D(x)$ be a definable set. Let n be a positive natural number. The *finite metric entropy quantifiers* are defined as

$$\begin{aligned}\exists_{>\varepsilon}^{\text{ent} \geq n} xL(x) &= (\exists \bar{x} \in L) \bigwedge_{i < j < n} dx_i x_j > \varepsilon, \\ \Theta_{\geq \varepsilon}^{\text{ent} \geq n} xF(x) &= (\Theta \bar{x} \in F) \bigwedge_{i < j < n} dx_i x_j \geq \varepsilon, \text{ and} \\ \text{ent}_x^n D(x) &= \sup_{\bar{x} \in D} \min_{i < j < n} dx_i x_j.\end{aligned}$$

The *finite (external) covering number quantifiers* are defined as

$$\begin{aligned}\exists_{\leq \varepsilon}^{\text{cov} \geq n} xL(x) &= (\Lambda > 0) \wedge \forall \bar{x} (\exists y \in L) \bigwedge_{i < n-1} dx_i y > \varepsilon, \\ \Theta_{< \varepsilon}^{\text{cov} \geq n} xF(x) &= (\Lambda = 1) \wedge \forall \bar{x} (\Theta y \in F) \bigwedge_{i < n-1} dx_i y \geq \varepsilon, \text{ and} \\ \text{cov}_x^n D(x) &= \Lambda \cdot \inf_{\bar{x}} \sup_{y \in D} \min_{i < n} dx_i y,\end{aligned}$$

where Λ is the $\{0, 1\}$ -valued sentence such that $\mathfrak{M} \models \Lambda = 0$ if and only if \mathfrak{M} is empty. \triangleleft

Note that in general ε could be a formula, rather than just a fixed real number.

It is also possible to define quantifiers for the (external) packing number,⁶ but we won't need them. Here we are really feeling the pinch of not being able to form statements like $V(x) \rightarrow U(x)$. We can't form upper bounds on metric entropy or covering numbers of open formulas that are equivalent to an open formula or on closed formulas that are equivalent to a closed formula. Furthermore, we can't formalize the internal covering number in an open or closed way, except for definable sets. Beyond these difficulties, most of these are only valid in sufficiently saturated models.

Proposition 2.3.57. *Let \mathfrak{M} be an \mathcal{S} -structure. Let $F(x)$ be an open formula, and let $L(x)$ be a strongly locatable formula (such as in particular an open set or a definable set). Fix $\varepsilon > 0$.*

(i) $\mathfrak{M} \models \exists_{>\varepsilon}^{\text{ent} \geq n} xL(x)$ if and only if $\#_{>\varepsilon}^{\text{ent}} L(\mathfrak{M}) \geq n$.

(ii) If $\#_{\geq\varepsilon}^{\text{ent}} F(\mathfrak{M}) \geq n$, then $\mathfrak{M} \models \exists_{\geq\varepsilon}^{\text{ent} \geq n} xF(x)$.

(iii) If $\mathfrak{M} \models \exists_{\leq\varepsilon}^{\text{cov} \geq n} xL(x)$, then $\#_{\leq\varepsilon}^{\text{cov}, M} L(\mathfrak{M}) \geq n$.

(iv) If $\#_{<\varepsilon}^{\text{cov}, M} F(\mathfrak{M}) \geq n$, then $\mathfrak{M} \models \exists_{<\varepsilon}^{\text{cov} \geq n} xF(x)$.

The converses hold if \mathfrak{M} is \aleph_1 -saturated. Let $D(x)$ be a definable set.

(v) $\mathfrak{M} \models \text{ent}_x^n D(x) > \varepsilon$ if and only if $\mathfrak{M} \models \exists_{>\varepsilon}^{\text{ent} \geq n} xD(x)$.

(vi) $\mathfrak{M} \models \text{ent}_x^n D(x) \geq \varepsilon$ if and only if $\mathfrak{M} \models \exists_{\geq\varepsilon}^{\text{ent} \geq n} xD(x)$.

(vii) $\mathfrak{M} \models \text{cov}_x^n D(x) > \varepsilon$ if and only if $\mathfrak{M} \models \exists_{\leq\varepsilon}^{\text{cov} \geq n} xD(x)$.

⁶Although not the real valued variant for definable sets.

(viii) $\mathfrak{M} \models \text{cov}_x^n D(x) \geq \varepsilon$ if and only if $\mathfrak{M} \models \exists_{<\varepsilon}^{\text{cov} \geq n} x D(x)$.

So in particular,

$$\begin{aligned} (\text{ent}_x^n D(x))^{\mathfrak{M}} &= \sup\{\varepsilon > 0 : \mathfrak{M} \models \exists_{>\varepsilon}^{\text{ent} \geq n} x D(x)\} \\ &= \sup\{\varepsilon > 0 : \mathfrak{M} \models \exists_{\geq\varepsilon}^{\text{ent} \geq n} x D(x)\} \text{ and} \\ (\text{cov}_x^n D(x))^{\mathfrak{M}} &= \sup\{\varepsilon > 0 : \mathfrak{M} \models \exists_{\leq\varepsilon}^{\text{cov} \geq n} x D(x)\} \\ &= \sup\{\varepsilon > 0 : \mathfrak{M} \models \exists_{<\varepsilon}^{\text{cov} \geq n} x D(x)\}. \end{aligned}$$

Proof. These statements follow from the definition of strongly locatable, as well as basic properties of \mathbb{W} and \exists and Proposition 1.8.2. \square

Definition 2.3.58. For $\square \in \{<, \leq\}$ and a theory T , a partial type Σ (or equivalently a closed type-set formula) is called $(\square \varepsilon)$ -algebraic (over T) if for every model $\mathfrak{M} \models T$, $\#_{\square\varepsilon}^{\text{cov}} \Sigma(\mathfrak{M})$ is finite.

Σ is called *algebraic (over T)* if it is $(< \varepsilon)$ -algebraic (over T) for every $\varepsilon > 0$. \triangleleft

A few other characterizations of $(< \varepsilon)$ -algebraic will be useful at some point.

Proposition 2.3.59. *Let T be a complete theory, and let $F(\bar{x})$ be a closed type-set formula. The following are equivalent:*

- (i) $F(\bar{x})$ is $(< \varepsilon)$ -algebraic over T ,
- (ii) for some model \mathfrak{M} , $\llbracket F(\bar{x}) \rrbracket_{\mathfrak{M}} \subseteq (\mathfrak{M}^n)^{<\varepsilon}$,
- (iii) for every model \mathfrak{M} , $\llbracket F(\bar{x}) \rrbracket_{\mathfrak{M}} \subseteq (\mathfrak{M}^n)^{<\varepsilon}$, and
- (iv) $T \models \neg \exists_{<\varepsilon}^{\text{cov} \geq k} \bar{x} F(\bar{x})$ for some $k < \omega$ with $k > 1$,

where we regard \mathfrak{M}^n as a subset of $S_{\bar{x}}(\mathfrak{M})$ (note that $\llbracket F(\bar{x}) \rrbracket_{\mathfrak{M}} \subseteq S_{\bar{x}}(\mathfrak{M})$), specifically the image of \mathfrak{M}^n under the map $\bar{a} \mapsto \text{tp}(\bar{a}/\mathfrak{M})$.

Proof. If T has an empty model then this is trivial, so assume that T has a non-empty model. If $F(\bar{x})$ is inconsistent over T , then this is also trivial, so assume that $F(\bar{x})$ is consistent over T .

(i) \Rightarrow (iv). Assume that for every $k < \omega$, $T \models \exists_{<\varepsilon}^{\text{cov} \geq k} \bar{x}F(\bar{x})$. This implies that in a sufficiently saturated model of T , we have a witness that F is not $(< \varepsilon)$ -algebraic.

(iv) \Rightarrow (i). Obvious.

(iv) \Rightarrow (iii). For $k > 1$, $\neg \exists_{<\varepsilon}^{\text{cov} \geq k} \bar{x}F(\bar{x})$ is logically equivalent (over a theory with non-empty models) to $\exists \bar{x}_0 \dots \bar{x}_{k-2} (\forall \bar{y} \in F) \bigvee_{i < k-1} d(\bar{x}_i, \bar{y}) < \varepsilon$. This implies that there are $\bar{a}_0, \dots, \bar{a}_{k-2} \in \mathfrak{M}$ such that $\llbracket F(\bar{x}) \rrbracket_{\mathfrak{M}} \subseteq \bigcup_{i < k-1} B_{<\varepsilon}(\text{tp}(\bar{a}_i/\mathfrak{M})) \subseteq (\mathfrak{M}^n)^{<\varepsilon}$.

(iii) \Rightarrow (ii). Obvious.

(ii) \Rightarrow (iv). By compactness there is a finite set $M_0 \subseteq \mathfrak{M}^n$ such that $\llbracket F(\bar{x}) \rrbracket_{\mathfrak{M}} \subseteq \bigcup_{\bar{a} \in M_0} B_{<\varepsilon}(\text{tp}(\bar{a}/\mathfrak{M}))$. M_0 must be non-empty, because $F(\bar{x})$ is consistent. Therefore $\mathfrak{M} \models \exists \bar{x}_0 \dots \bar{x}_{|M_0|-1} (\forall \bar{y} \in F) \bigvee_{i < |M_0|} d(\bar{x}_i, \bar{y}) < \varepsilon$, or, in other words, $\mathfrak{M} \models \neg \exists_{<\varepsilon}^{\text{cov} \geq |M_0|+1} \bar{x}F(\bar{x})$. Since M_0 is non-empty, $|M_0| + 1 > 1$. \square

An easy compactness argument shows that if $F(\bar{x}, \bar{a})$ is $(< \varepsilon)$ -algebraic for any choice of parameters \bar{a} , then there is a uniform n such that $\models \neg \exists_{<\varepsilon}^{\text{cov} \geq n} \bar{x}F(\bar{x}, \bar{a})$ for any \bar{a} . A similar statement is true for incomplete theories.

Corollary 2.3.60. *For any complete theory T , any set of parameters A , any $n \leq \omega$, and any $\varepsilon > 0$, the set of non- $(< \varepsilon)$ -algebraic types in $S_n(A)$ is topologically closed.*

Proof. Fix a model $\mathfrak{M} \supseteq A$. Consider the set $S_n(\mathfrak{M}) \setminus (\mathfrak{M}^n)^{<\varepsilon}$, which is closed since \mathfrak{M}^n as a subset of $S_n(\mathfrak{M})$ is a union of definable sets and is therefore locatable. Let

$f : S_n(\mathfrak{M}) \rightarrow S_n(A)$ be the natural restriction map. For any type $p \in S_n(A)$, p is $(< \varepsilon)$ -algebraic if and only if $f^{-1}(p) \subseteq (\mathfrak{M}^n)^{<\varepsilon}$. This implies that $f(S_n(\mathfrak{M}) \setminus (\mathfrak{M}^n)^{<\varepsilon})$ is precisely the set of non- $(< \varepsilon)$ -algebraic types in $S_n(A)$, which is closed since it is the continuous image of a compact set. \square

Corollary 2.3.61. *For any $\varepsilon > 0$ and any closed (type-set) formula $F(\bar{x}, \bar{y})$, there is an open (type-set) formula $U_\varepsilon(\bar{y})$ such that $\mathfrak{M} \models U_\varepsilon(\bar{a})$ if and only if $F(\bar{x}, \bar{a})$ is $(< \varepsilon)$ -algebraic over $\text{Th}(\mathfrak{M}_{\bar{a}})$. In particular, the set of types $p(\bar{y})$ in $S_{\bar{y}}(T)$ such that $\mathfrak{M} \models p(\bar{a})$, then $F(\bar{x}, \bar{a})$ is algebraic is G_δ .*

Proof. $U_\varepsilon(\bar{y}) = \bigvee_{n < \omega} \neg \exists_{<\varepsilon}^{\text{cov} \geq n} \bar{x} F(\bar{x}, \bar{y})$.

$F(\bar{x}, \bar{a})$ is algebraic if and only if $\mathfrak{M} \models \bigwedge_{k < \omega} U_{2^{-k}}(\bar{a})$. \square

Contrast Corollary 2.3.61 with the fact that the the analogous property in discrete logic is open.

The following lemma is occasionally useful for verifying that a formula is algebraic.

Lemma 2.3.62. *If F, G are closed formulas such that $\llbracket F \rrbracket \subseteq \text{int}[\llbracket G \rrbracket]$ and $G(\mathfrak{M})$ is metrically compact, then F is algebraic.*

Proof. Find an open formula $U(\bar{x})$ such that $\llbracket F(\bar{x}) \rrbracket \subseteq \llbracket U(\bar{x}) \rrbracket \subseteq \text{int}[\llbracket G(\bar{x}) \rrbracket]$. For each $\varepsilon > 0$, there is an $n_\varepsilon < \omega$ such that $\mathfrak{M} \models \neg \exists_{\leq \varepsilon}^{\text{cov} \geq n_\varepsilon} \bar{x} U(x)$ (i.e. $U(\mathfrak{M})$ requires fewer than n_ε closed ε -balls to cover). This implies that $\mathfrak{M} \models \neg \exists_{< 2\varepsilon}^{\text{cov} \geq n_\varepsilon} \bar{x} F(\bar{x})$ (i.e. $F(\mathfrak{M})$ requires fewer than n_ε open 2ε -balls to cover), so, by Proposition 2.3.57, F is algebraic. \square

Proposition 2.3.63. *A closed type-set formula $F(\bar{x})$ is algebraic over the theory T if and only if it is bounded (i.e. there is a fixed cardinal κ such that for any $\mathfrak{M} \models T$, $|F(\mathfrak{M})| \leq \kappa$).*

Proof. It is obvious that algebraic closed type-set formulas are bounded with $\kappa = 2^{\aleph_0}$. For the converse, assume that a closed type-set formula is not algebraic, then for some $\varepsilon > 0$, $T \models \exists_{<\varepsilon}^{\text{cov}} \geq n \bar{x} F(\bar{x})$ for every $n < \omega$. By Proposition 0.3.5, this implies that $T \models \exists_{\geq \varepsilon}^{\text{ent}} \geq n \bar{x} F(\bar{x})$ as well (by considering sufficiently saturated models of T). Therefore some model of T has an infinite ($\geq \varepsilon$)-separated set of realizations of F , and by compactness this implies that F can have arbitrarily many realizations. \square

In particular, an algebraic type always has at most 2^{\aleph_0} realizations.

Lemma 2.3.64. *If A is a set of parameters and $B \subseteq A^n$, for some $n \leq \omega$, is metrically compact, then the set $\llbracket F \rrbracket = \{\text{tp}(\bar{b})\}_{\bar{b} \in B} \subseteq S_{\bar{x}}(A)$ is definable.*

Proof. Singleton sets of parameters are obviously definable, finite unions of definable sets are definable, and $\llbracket F \rrbracket$ is a Hausdorff limit of finite sets of parameters. \square

Proposition 2.3.65. *Let T be a complete theory. Any algebraic closed type-set formula $F(\bar{x})$ over T is definable over T .*

Proof. By Lemma 2.3.64, $F(\mathfrak{C})$ is definable over the monster \mathfrak{C} , since it is a compact set of parameters. By Proposition B.1.7, $F(\bar{x})$ is definable without parameters, since it is invariant under the automorphisms of \mathfrak{C} . \square

Corollary 2.3.66. *If $F(\bar{x})$ is an algebraic closed formula over T , then for any structures $\mathfrak{M} \preceq \mathfrak{N}$ such that $\mathfrak{M} \models T$, $F(\mathfrak{M}) = F(\mathfrak{N})$.*

Proof. Suppose that $F(\mathfrak{M}) \subset F(\mathfrak{N})$. Since these sets are both compact, this implies that for some sufficiently small $\varepsilon > 0$, $\#_{>\varepsilon}^{\text{ent}} F(\mathfrak{M}) < \#_{>\varepsilon}^{\text{ent}} F(\mathfrak{N})$. But since $\#_{>\varepsilon}^{\text{ent}} F(\mathfrak{M})$ is finite, its value is a first-order property of \mathfrak{M} by Proposition 2.3.57 part (i), and so this contradicts the fact that \mathfrak{M} is an elementary sub-structure of \mathfrak{N} . \square

Note, however, that a pre-structure can omit an algebraic partial type.

Definition 2.3.67. Given a pre-structure \mathfrak{M} and a set $A \subseteq M$, the *algebraic closure* of A , written $\text{acl}_{\mathfrak{M}}(A)$, is the set $\{b \in M : \text{tp}(b/A) \text{ is algebraic}\}$. We will typically omit the subscript \mathfrak{M} when it is clear from context. \triangleleft

Note that by definition we have that if $\mathfrak{M} \preceq \mathfrak{N}$, then for any $A \subseteq \mathfrak{M}$, $\text{acl}_{\mathfrak{M}}(A) = \text{acl}_{\mathfrak{N}}(A)$.

Proposition 2.3.68. For any \mathfrak{M} , $A \mapsto \text{acl}_{\mathfrak{M}}(A)$ is a closure operator (i.e. $A \subseteq B \Rightarrow A \subseteq \text{acl}_{\mathfrak{M}}(A) \subseteq \text{acl}_{\mathfrak{M}}(B) = \text{acl}_{\mathfrak{M}}(\text{acl}_{\mathfrak{M}}(B))$) with countable character (i.e. $b \in \text{acl}_{\mathfrak{M}}(A)$ if and only if $b \in \text{acl}_{\mathfrak{M}}(A_0)$ for some countable $A_0 \subseteq A$).

Proof. Reflexivity ($A \subseteq \text{acl}_{\mathfrak{M}}(A)$): For any $a \in A$, $\text{tp}(a/A)$ is algebraic.

Monotonicity ($A \subseteq B \Rightarrow \text{acl}_{\mathfrak{M}}(A) \subseteq \text{acl}_{\mathfrak{M}}(B)$): $c \in \text{acl}_{\mathfrak{M}}(A)$ if and only if for any $\mathfrak{N} \succeq \mathfrak{M}$, there are at most continuum many e such that $c \equiv_A e$. This implies that there are at most continuum many e such that $c \equiv_B e$, so $c \in \text{acl}_{\mathfrak{M}}(B)$ as well.

Idempotence/Transitivity ($\text{acl}_{\mathfrak{M}}(A) = \text{acl}_{\mathfrak{M}}(\text{acl}_{\mathfrak{M}}(A))$): By monotonicity we have that $\text{acl}_{\mathfrak{M}}(A) \subseteq \text{acl}_{\mathfrak{M}}(\text{acl}_{\mathfrak{M}}(A))$, so we just need to show that $\text{acl}_{\mathfrak{M}}(A) \supseteq \text{acl}_{\mathfrak{M}}(\text{acl}_{\mathfrak{M}}(A))$. Assume that $b \in \text{acl}_{\mathfrak{M}}(\text{acl}_{\mathfrak{M}}(A))$. By countable character, there is some at most countable set $C \subseteq \text{acl}_{\mathfrak{M}}(A)$ such that $b \in \text{acl}_{\mathfrak{M}}(C)$. In the monster, there are at most continuum many C' satisfying $C' \equiv_A C$. Since there are also at most continuum many b' with $b' \equiv_C b$, this implies that there are at most continuum many $C'b'$ with $C'b' \equiv_A Cb$, so in particular there are at most continuum many b' such that $b' \equiv_A b$, and we have that $b \in \text{acl}_{\mathfrak{M}}(A)$, as required.

Countable Character: If $p = \text{tp}(b/A)$ is algebraic, then there is an A -definable distance predicate for the realizations of p . This means that that distance predicate axiomatizes the type p , but a distance predicate can only use at most countably many parameters, so the type of p is axiomatized by some restriction of it to some countable $A_0 \subseteq A$, and we have that $\text{tp}(b/A_0)$ is algebraic as well. \square

Proposition 2.3.69 (Approximate Finite Character). *If $b \in \text{acl}_{\text{M}}(A)$, then for every $\varepsilon > 0$, there is a finite $A_\varepsilon \subseteq A$ such that b is $(< \varepsilon)$ -algebraic over A_ε .*

Proof. Assume that for some $\varepsilon > 0$, b is not $(< \varepsilon)$ -algebraic over any finite subset of A . This implies that for any finite subset $A_0 \subseteq A$, there exists an infinite $(> \varepsilon)$ -separated set of realizations of $\text{tp}(b/A_0)$. By compactness, this implies that there exists an infinite $(\geq \varepsilon)$ -separated set of realizations of $\text{tp}(b/A)$, implying that b is not algebraic over A . \square

2.3.5 Definable Functions

A definable function is a special case of a uniformly definable set. We will also need the broader notion of a definable partial function.

Definition 2.3.70. For any real formula $\varphi(\bar{x}, y)$ and partial $\mathcal{L}(\bar{x})$ -type $\Sigma(\bar{x})$, we say that $\varphi(\bar{x}, y)$ defines a partial function over Σ if

$$\Sigma(\bar{x}) \models \exists y \varphi(\bar{x}, y) = 0 \wedge \forall z d(y, z) = \varphi(\bar{x}, z).$$

If T is an \mathcal{L} -theory, we say that $\varphi(\bar{x}, y)$ defines a function over T if

$$T \models \forall \bar{x} \exists y \varphi(\bar{x}, y) = 0 \wedge \forall z d(y, z) = \varphi(\bar{x}, z). \quad \triangleleft$$

Proposition 2.3.71. *If $\varphi(\bar{x}, y)$ defines a partial function over $\Sigma(\bar{x})$, then $\varphi(\bar{x}, y) = 0$ is \bar{x} -uniformly definable over $\Sigma(\bar{x})$.*

Furthermore, for any structure \mathfrak{M} and $\bar{a} \in M$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, the set $(\varphi(\bar{a}, \mathfrak{M}) = 0)$ is a singleton.

Proof. Let \mathfrak{M} be a sufficiently saturated structure, and let $\bar{a} \in M$ be such that $\mathfrak{M} \models \Sigma(\bar{a})$. Then we have that there is a $b \in M$ such that $\mathfrak{M} \models \varphi(\bar{a}, b) = 0 \wedge \forall z d(b, z) = \varphi(\bar{a}, z)$. Now we want to argue that $\mathfrak{M} \models \forall v \exists w (\varphi(\bar{a}, w) = 0 \wedge d(v, w) = \varphi(\bar{a}, v))$. For any $c \in M$, we have that $\mathfrak{M} \models \varphi(\bar{a}, b) = 0 \wedge d(b, c) = \varphi(\bar{a}, c)$, so we have that $\varphi(\bar{a}, y)$ is the distance predicate of its zeroset.

We also showed that this zeroset is a singleton in any sufficiently saturated structure \mathfrak{N} such that $\mathfrak{N} \models \Sigma(\bar{a}')$. This implies that in any model it can have at most one realization, but since $\varphi(\bar{x}, y)$ is \bar{x} -uniformly definable, this realization must exist in any structure \mathfrak{N} such that $\mathfrak{N} \models \Sigma(\bar{a}')$. \square

Corollary 2.3.72. *If $\varphi(\bar{x}, y)$ defines a partial function over $\Sigma(\bar{x})$, then for any real formula $\psi(\bar{x}, y, \bar{z})$, there exists a real formula $\chi(\bar{x}, \bar{z})$ such that if $\mathfrak{M} \models \Sigma(\bar{a})$ and $b \in M$ is the unique element such that $\mathfrak{M} \models \varphi(\bar{a}, b) = 0$, then for any $\bar{c} \in M$, $\mathfrak{M} \models \psi(\bar{a}, b, \bar{c}) = \chi(\bar{a}, \bar{c})$.*

Furthermore for any topological formula $X(\bar{x}, y, \bar{z})$ there is a topological formula $Y(\bar{x}, \bar{z})$ (of the same open-or-closed-ness) such that for any \mathfrak{M} and $\bar{a} \in M$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, $X(\bar{a}, b, \bar{z})$ is logically equivalent to $Y(\bar{a}, \bar{z})$. The same holds for type-set formulas.

Proof. $\chi(\bar{x}, \bar{z}) = (\exists y \in (\varphi(\bar{x}, \cdot) = 0))\psi(\bar{x}, y, \bar{c})$. The statement for topological and type-set formulas follows immediately. \square

Notation 2.3.73. For any real formula $\varphi(\bar{x}, y)$, we let

$$\begin{aligned} \text{FUN}_y \varphi(\bar{x}, y) &= \exists y (\varphi(\bar{x}, y) = 0 \wedge \forall z (dyz = \varphi(\bar{x}, z))) \text{ and} \\ \text{fun}_y \varphi(\bar{x}, y) &= \inf_y \max\{|\varphi(\bar{x}, y)|, \sup_z |dyz - \varphi(\bar{x}, z)|\}. \end{aligned} \quad \triangleleft$$

Just as with $\text{DEF}\bar{v}$, FUN_y is a quantifier that takes a real formula and produces a closed formula. Furthermore, $\text{FUN}_y \varphi(\bar{x}, y)$ is logically equivalent to $\text{fun}_y \varphi(\bar{x}, y) = 0$. FUN is somewhat similar to the quantifier $\exists!$ in discrete logic, but it is more strict in the sense that it requires $\varphi(\bar{x}, y)$ to be a very explicit witness of uniqueness.

Lemma 2.3.74. For any formulas $\varphi(\bar{x}, y), \psi(\bar{x}, y)$ and partial type $\Sigma(\bar{x})$,

$$\left\| \text{fun}_y \varphi(\bar{x}, y) - \text{fun}_y \psi(\bar{x}, y) \right\|_{\Sigma} \leq \|\varphi(\bar{x}, y) - \psi(\bar{x}, y)\|_{\Sigma}.$$

Proof. This follows from the same argument as Lemma 2.3.22. \square

Definition 2.3.75 (Definable (Partial) Functions). A *definable partial function* $f_{\varphi}(\bar{x})$ is an equivalence class of real formula $\varphi(\bar{x}, y)$ under an equivalence relation we will define here. For F , a closed formula, and U , an open formula, we define:

$$\begin{aligned} f_{\varphi}(\bar{x})\downarrow &= \text{FUN}_y \varphi(\bar{x}, y), \\ f_{\varphi}(\bar{x})\uparrow &= \neg f_{\varphi}(\bar{x})\downarrow, \\ (f_{\varphi}(\bar{x})\downarrow = y) &= f_{\varphi}(\bar{x})\downarrow \wedge (\varphi(\bar{x}, y) = 0), \\ f_{\varphi}(\bar{x})\downarrow \wedge G(\bar{x}, f_{\varphi}(\bar{x}), \bar{z}) &= f_{\varphi}(\bar{x})\downarrow \wedge (\exists y \in (\varphi(\bar{x}, \cdot) = 0))G(\bar{x}, y, \bar{z}), \text{ and} \\ f_{\varphi}(\bar{x})\uparrow \vee U(\bar{x}, f_{\varphi}(\bar{x}), \bar{z}) &= f_{\varphi}(\bar{x})\uparrow \vee (\exists y \in (\varphi(\bar{x}, \cdot) = 0))U(\bar{x}, y, \bar{z}). \end{aligned}$$

We may also write $f_\varphi(\bar{x})\downarrow \wedge F(\bar{x}, f_\varphi(\bar{x}), \bar{z})$ as $F(\bar{x}, f_\varphi(\bar{x})\downarrow, \bar{z})$, although never when it is in the scope of a negation. We may use this notation with multiple partial functions at once.

We say that $f_\psi(\bar{x})$ *extends* $f_\varphi(\bar{x})$, written $f_\psi(\bar{x}) \supseteq f_\varphi(\bar{x})$ or $f_\varphi(\bar{x}) \subseteq f_\psi(\bar{x})$, if

$$f_\varphi(\bar{x})\downarrow \models f_\varphi(\bar{x})\downarrow = f_\psi(\bar{x})\downarrow.$$

We regard two definable partial functions, $f_\varphi(\bar{x})$ and $f_\psi(\bar{x})$, as identical if $f_\varphi(\bar{x}) \subseteq f_\psi(\bar{x})$ and $f_\psi(\bar{x}) \subseteq f_\varphi(\bar{x})$. We will write equivalence classes directly with expressions such as $f(\bar{x})$ or $g(\bar{x})$, introducing them as definable partial functions. It is not hard to show that the notation defined here is invariant under this equivalence relation.

For a theory T , a definable partial function $f(\bar{x})$ is *total over* T if $T \models \forall \bar{x} f(\bar{x})\downarrow$. We also say that f is a *definable function over* T . \triangleleft

Note that $f(\bar{x})\downarrow$, $f(\bar{x})\downarrow = y$, and $f(\bar{x})\downarrow \wedge F(\bar{x}, f(\bar{x}), \bar{z})$ are closed formulas and $f(\bar{x})\uparrow$ and $f(\bar{x})\uparrow \vee U(\bar{x}, f(\bar{x}), \bar{z})$ are open formulas. Also note that we have not explicitly defined $U(\bar{x}, f(\bar{x})\downarrow, \bar{z})$, as this is not a topological formula, although it is sensible as a type-set formula.

We have engineered this notation so that the semantic meaning of each depends on the value of $\varphi(\bar{x}, y)$ on $S_{\bar{x}y}(F)$. For example, $\mathfrak{M} \models f(\bar{a})\downarrow \wedge F(\bar{a}, f(\bar{a}), \bar{c})$ if and only if $d(y, f(\bar{x}))$ defines a partial function over $\text{tp}(\bar{a})$, b is the unique element of $(d(\mathfrak{M}, f(\bar{a})) = 0)$, and $\mathfrak{M} \models F(\bar{a}, b, \bar{c})$. Although note that the last two are literal if we permit a short-circuiting interpretation of \wedge and \vee , as, for example, $F(\bar{x}, f(\bar{x}), \bar{z})$ isn't literally sensible unless $f(\bar{x})\downarrow$.

Proposition 2.3.76. *If $f(\bar{x})$ is a definable partial function, then for any closed formula $F(\bar{x})$, there is a definable partial function $g(\bar{x})$ such that $g(\bar{x}) \subseteq f(\bar{x})$ and $g(\bar{x})\downarrow$ is logically equivalent to $f(\bar{x})\downarrow \wedge F(\bar{x})$. By compactness there is a finite $F_0 \subseteq \llbracket F \rrbracket$ such that $\bigcup_{p \in F_0} \llbracket \chi_{\varepsilon, \delta(\varepsilon, p)} \rrbracket \supseteq \llbracket F \rrbracket$.*

Proof. Let $f(\bar{x}) = f_\varphi(\bar{x})$ and $F(\bar{x}) = (\psi(\bar{x}) = 0)$. Then $g(\bar{x}) = f_{|\varphi|+|\psi|}(\bar{x})$. \square

Notation 2.3.77. If $f(\bar{x})$ is a definable partial function and $F(\bar{x})$ is a closed formula, we write $(f \upharpoonright F)(\bar{x})$ for the definable partial function from Proposition 2.3.76. \triangleleft

The following is more of an observation than a proposition, but it is very useful.

Proposition 2.3.78 (Metric Closure of Image of Definable Function is Definable). *For any theory T , if $f(\bar{x})$ is a definable total function over T , then in any $\mathfrak{M} \models T$, the distance predicate of $\overline{f(M^{|\bar{x}|})}$ is $\varphi(\bar{x}) = \inf_{\bar{y}} d(x, f(\bar{y}))$. \square*

Proposition 2.3.79 (Pointwise Definable (Partial) Functions are Definable). *If $F(\bar{x})$ and $G(\bar{x}, y)$ are closed formulas such that for any structure \mathfrak{M} and $\bar{a} \in M$ such that $\mathfrak{M} \models F(\bar{a})$ there is a unique $b \in M$ such that $\mathfrak{M} \models G(\bar{a}, b)$, then there exists a definable partial function $f(\bar{x})$ such that $F(\bar{x}) \equiv f(\bar{x})\downarrow$ and $(f(\bar{x})\downarrow = y) \equiv F(\bar{x}) \wedge G(\bar{x}, y)$.*

Proof. Let $G(\bar{x}, y) \equiv (\varphi(\bar{x}, y) = 0)$. Assume without loss that $\models \varphi(\bar{x}, y) \geq 0$. Let

$$\chi_{\varepsilon, \delta}(\bar{x}) = \forall yz (\varphi(\bar{x}, \bar{y}) \leq \delta \wedge \varphi(\bar{x}, \bar{z}) \leq \delta \rightarrow d\bar{y}\bar{z} < \varepsilon).$$

For each type $p(\bar{x}) \in \llbracket F \rrbracket$ and for every $\varepsilon > 0$, there is a $\delta(\varepsilon, p) > 0$ such that for any $\bar{a} \in \mathfrak{M}$ with $\mathfrak{M} \models F(\bar{a})$, $\mathfrak{M} \models \chi_{\varepsilon, \delta(\varepsilon, p)}(\bar{a})$, by compactness. Also, clearly, $\mathfrak{M} \models \exists y \varphi(\bar{a}, y) = 0$.

For each $\varepsilon > 0$, we have that $\{\llbracket \chi_{\varepsilon, \delta(\varepsilon, p)} \rrbracket : p \in \llbracket F \rrbracket\}$ is an open cover of $\llbracket F \rrbracket$. Therefore,

by compactness there is finite $F_0(\varepsilon) \subseteq \llbracket F \rrbracket$ such that $\bigcup_{p \in F_0(\varepsilon)} \llbracket \chi_{\varepsilon, \delta(\varepsilon, p)} \rrbracket \supseteq \llbracket F \rrbracket$. Let

$$\gamma(\varepsilon) = \min\{\varepsilon, \min\{\delta(\varepsilon, p) : p \in F_0(\varepsilon)\}\}.$$

Now we have that $F(\bar{x}) \models \chi_{\varepsilon, \gamma(\varepsilon)}(\bar{x})$.

So now we have that for any $\varepsilon > 0$, for any $p(\bar{x}, y) \in \llbracket F \rrbracket_{\bar{x}y}$, if $p \in \llbracket \varphi < \gamma(\varepsilon) \rrbracket$, then $p \in \llbracket G \rrbracket^{d/\bar{x} < \varepsilon}$. Therefore $G(\bar{x}, y)$ is \bar{x} -uniformly definable over F , and there is a real formula $\eta(\bar{x}, y)$ such that for any $p(\bar{x}, y) \in \llbracket F \rrbracket_{\bar{x}y}$, $d/\bar{x}(p, \llbracket G \rrbracket) = \eta(p)$. By assumption we have $F(\bar{x}) \models \text{FUN}y\eta(\bar{x}, y)$, so G gives a definable partial function over F , as required. \square

What is remarkable about Proposition 2.3.79 is that the analogous fact fails for closed formulas $F(\bar{x}, \bar{y})$ such that $F(\bar{a}, \bar{y})$ is definable for every parameter tuple \bar{a} , even if we require that $F(\bar{a}, \bar{y})$ be algebraic. Note that it is not enough to check Proposition 2.3.79 in a single model (see Counterexample C.1.1).

Corollary 2.3.80 (Definable Function iff Graph is Type-Definable). *If $\Sigma(\bar{x}, y)$ is a partial type (or, equivalently, a closed type-set formula) and T is a theory, then $\Sigma(\mathfrak{M})$ is the graph of a function for any model $\mathfrak{M} \models T$ if and only if there is a real formula $\varphi(\bar{x}, y)$ defining a function over T such that $\Sigma(\bar{x}, y) \equiv_T (\varphi(\bar{x}, y) = 0)$.*

Proof. The \Leftarrow direction is obvious. For the \Rightarrow direction, we just need to show that there is a closed formula $F(\bar{x}, y)$ such that $\Sigma(\bar{x}, y) \models F(\bar{x}, y)$ and such that F has the same unique witness property that Σ does.

Assume that for some $\varepsilon > 0$, for every closed formula $G(\bar{x}, y)$ such that $\Sigma(\bar{x}, y) \models G(\bar{x}, y)$, there exists $\mathfrak{M} \models T$ and $\bar{a}bc \in M$ such that $\mathfrak{M} \models G(\bar{a}, b) \wedge G(\bar{a}, c) \wedge dbc \geq \varepsilon$. Then by compactness, there exists $\mathfrak{M} \models T$ and $\bar{a}bc \in M$ such that $\mathfrak{M} \models \Sigma(\bar{a}, b) \wedge$

$\Sigma(\bar{a}, c) \wedge dbc \geq \varepsilon$, which contradicts our assumption, so there is some $G_\varepsilon(\bar{x}, y)$ such that $T \models \forall yz(G_\varepsilon(\bar{x}, y) \wedge G_\varepsilon(\bar{x}, z)) \rightarrow dyz < \varepsilon$ and $\Sigma(\bar{x}, y) \models G_{<\varepsilon}(\bar{x}, y)$.

Let $F(\bar{x}, y) = \bigwedge_{n < \omega} G_{2^{-n}}(\bar{x}, y)$. Then by construction we have that $\Sigma(\bar{x}, y) \models F(\bar{x}, y)$ and that $F(\bar{x}, y)$ has the same unique witness property that $\Sigma(\bar{x}, y)$ does. Therefore by Proposition 2.3.79, there is a real formula $\varphi(\bar{x}, y)$ such that $F(\bar{x}, y) \equiv_T (\varphi(\bar{x}, y) = 0)$ and such that $\varphi(\bar{x}, y)$ defines a function over T . This implies that $\Sigma(\bar{x}, y) \equiv_T (\varphi(\bar{x}, y) = 0)$, as well. \square

Definition 2.3.81. Given a pre-structure \mathfrak{M} and a set $A \subseteq M$, the *definable closure* of A , written $\text{dcl}_{\mathfrak{M}}(A)$, is the set $\{b \in M : \text{tp}(b/A) \models \text{tp}(b/M)\}$. We will typically omit the subscript \mathfrak{M} when it is clear from context. \triangleleft

Given our definition we get this immediately.

Proposition 2.3.82. *For any signature \mathcal{L} , any \mathcal{L} -pre-structure, any set of parameters $A \subseteq \mathfrak{M}$, and any tuple of variables \bar{x} , the natural reduct map $S_{\bar{x}}(\text{dcl}(A)) \rightarrow S_{\bar{x}}(A)$ is a topological homeomorphism.*

Proof. By the definition of dcl the map is a bijection. By Fact A.2.11, continuous bijections between compact Hausdorff spaces are homeomorphisms. \square

Now we can give a roundabout characterization of structures among pre-structures.

Proposition 2.3.83. *Let \mathcal{L} be a signature. For any reduced \mathcal{L} -structure \mathfrak{M} , any set of parameters $A \subseteq \mathfrak{M}$, and \mathcal{L} -structure $\mathfrak{N} \succeq \mathfrak{M}$, $\bar{A} \subseteq \text{dcl}_{\mathfrak{N}}(A) \subseteq \overline{\mathfrak{M}}$. In particular, $\text{dcl}_{\mathfrak{N}}(\mathfrak{M}) = \overline{\mathfrak{M}}$.*

Proof. Let $\bar{a} \in A$ be a Cauchy sequence with limit a_ω . $\text{tp}(a_\omega/A)$ has a unique realization

in any model containing A , so $\text{tp}(a_\omega/A) \models \text{tp}(a_\omega/M)$, and we have that $a_\omega \in \text{dcl}_{\mathfrak{M}}(A)$.

□

Corollary 2.3.84. *A pre-structure \mathfrak{M} is a reduced pre-structure if and only if the map $a \mapsto \text{tp}(a/M)$ is an injection. A reduced pre-structure \mathfrak{M} is a structure if and only if for any reduced pre-structure $\mathfrak{N} \succeq \mathfrak{M}$, the reduct map $S_1(\mathfrak{N}) \rightarrow S_1(\mathfrak{M})$ is a proper surjection (i.e. a surjection that is not a bijection).*

□

Almost as in discrete logic, the definable closure is witnessed by definable partial functions.

Proposition 2.3.85. *For any structure \mathfrak{M} and $A \subseteq M$, $b \in \text{dcl}_{\mathfrak{M}}(A)$ if and only if there is a definable partial function f such that for some $\bar{a} \in A$, $\mathfrak{M} \models f(\bar{a}) \downarrow = b$.*

Proof. For the \Leftarrow direction, note that $(f(\bar{a}) \downarrow = y) \in \text{tp}(b/A)$ and that therefore $\text{tp}(b/A)$ can have at most one realization in any $\mathfrak{N} \succeq \mathfrak{M}$. This implies that $\text{tp}(b/A) \models \text{tp}(b/N) \models \text{tp}(b/M)$.

For the \Rightarrow direction, assume that $\text{tp}(b/A) \models \text{tp}(b/M)$. Since $\text{tp}(b/A) \models \text{tp}(b/M)$, by compactness for each $n < \omega$, there is a closed formula $F_n(\bar{a}_n, y) \in \text{tp}(b/A)$ such that $F_n(\bar{a}_n, y) \models dby < 2^{-n}$. Let $G(\bar{a}, y) = \bigwedge_{n < \omega} F_n(\bar{a}_n, y)$. Now we have that $G(\bar{a}, y) \models dby < 2^{-n}$ for every $n < \omega$, so $G(\bar{a}, y) \models b = y$, so since $\text{tp}(b/A) \models G(\bar{a}, y)$, by Proposition 2.3.79 there is a real formula $\varphi(\bar{x}, y)$ that defines a partial function over $\text{tp}(\bar{a})$ such that $G(\bar{a}, y) \equiv (\varphi(\bar{a}, y) = 0)$. □

Unlike in discrete logic, there is in general no guarantee that definable closure will be witnessed by a total function, because in general definable partial functions may fail to

be extendable to definable total functions, even on a definable set (see Counterexample C.1.2).⁷

Corollary 2.3.86. *For any structures $\mathfrak{M} \succeq \mathfrak{N}$ and $A \subseteq \mathfrak{M}$, $\text{dcl}_{\mathfrak{M}}(A) = \text{dcl}_{\mathfrak{N}}(A)$. \square*

Corollary 2.3.87. *For any structure \mathfrak{M} , $A \mapsto \text{dcl}_{\mathfrak{M}}(A)$ is an abstract closure operator with countable character and approximate finite character, i.e.*

(i) $A \subseteq \text{dcl}_{\mathfrak{M}}(A)$.

(ii) $\text{dcl}_{\mathfrak{M}}(\text{dcl}_{\mathfrak{M}}(A)) = \text{dcl}_{\mathfrak{M}}(A)$.

(iii) If $A \subseteq B$, then $\text{dcl}_{\mathfrak{M}}(A) \subseteq \text{dcl}_{\mathfrak{M}}(B)$.

(iv) (Countable Character) If $b \in \text{dcl}_{\mathfrak{M}}(A)$, then there is at most countable $A_0 \subseteq A$ such that $b \in \text{dcl}_{\mathfrak{M}}(A_0)$.

(v) (Approximate Finite Character) For any $\varepsilon > 0$, if $b \in \text{dcl}_{\mathfrak{M}}(A)$, then there is a finite $A_0 \subseteq A$ such that $\text{tp}(b/A_0) \models \text{dby} < \varepsilon$.

Furthermore, $\text{dcl}_{\mathfrak{M}}(A) \subseteq \text{acl}_{\mathfrak{M}}(A)$.

Proof. (i) and (iii) are straightforward.

(ii) If $c \in \text{dcl}_{\mathfrak{M}}(\text{dcl}_{\mathfrak{M}}(A))$, then there exists a definable partial function $f(\bar{x})$ and a tuple $\bar{b} \in \text{dcl}_{\mathfrak{M}}(A)$ such that $\mathfrak{M} \models f(\bar{b}) \downarrow = c$. By definition, for each $b_i \in \bar{b}$, there is a definable partial function $g_i(\bar{x}_i)$ and $\bar{a}_i \in A$ such that $\mathfrak{M} \models g_i(\bar{a}_i) \downarrow = b_i$. Therefore $\mathfrak{M} \models \bigwedge_{i < \omega} g(\bar{a}_i) \downarrow \wedge f(g_0(\bar{a}_0), g_1(\bar{a}_1), \dots) \downarrow = c$ and so $c \in \text{dcl}_{\mathfrak{M}}(A)$, as required.

(iv) This follows from compactness and the fact that a closed formula uses at most countably many parameters.

⁷This answers a question raised in [BY10a] immediately after Lemma 1.23.

(v) This follows from compactness, the logical completeness of restricted closed formulas, and the fact that a restricted closed formula uses finitely many parameters.

For the furthermore statement, note that if $b \in \text{dcl}_{\mathfrak{M}}(A)$, then $tp(b/A)$ has one realization in any elementary extension of \mathfrak{M} , which is bounded, so $b \in \text{acl}_{\mathfrak{M}}(A)$. \square

Definition 2.3.88. If \bar{x} is a $\Sigma(\bar{x})$ is a partial type and $f(\bar{x})$ and $g(\bar{x})$ are definable partial functions such that $\Sigma(\bar{x}) \models f(\bar{x})\downarrow \wedge g(\bar{x})\downarrow$, then the *logical distance between $f(\bar{x})$ and $g(\bar{x})$ over $\Sigma(\bar{x})$* , written $d_{\Sigma}(f, g)$, is given by

$$d_{\Sigma}(f, g) = \sup\{d^{\mathfrak{M}}(f^{\mathfrak{M}}(\bar{a}), g^{\mathfrak{M}}(\bar{a})) : \mathfrak{M} \models \Sigma(\bar{a})\}.$$

If Σ is empty, we write d_{\equiv} for d_{Σ} . \triangleleft

Proposition 2.3.89. *If $f(\bar{x})$ and $g(\bar{x})$ are definable partial functions such that $\Sigma(\bar{x}) \models f(\bar{x})\downarrow \wedge g(\bar{x})\downarrow$ and if $\varphi(\bar{x}, y)$ and $\psi(\bar{x}, y)$ are real formulas representing $f(\bar{x})$ and $g(\bar{x})$, then $d_{\Sigma}(f, g) = \|\varphi - \psi\|_{\Sigma}$.*

Proof. The proof of this is the same as the proof of Proposition 1.4.3. \square

Proposition 2.3.90. *For any $\Sigma(\bar{x})$ the set of definable partial functions $f(\bar{x})$ such that $\Sigma(\bar{x}) \models f(\bar{x})\downarrow$ is complete under the metric d_{Σ} . Furthermore, the metric density character of this set is at most $\aleph_0 + |\mathcal{L}| + |\bar{x}|$.*

Proof. Let $\{f_i(\bar{x})\}_{i < \omega}$ be a Cauchy sequence under d_{Σ} . Assume without loss that $d_{\Sigma}(f_i, f_{i+1}) \leq 2^{-i-1}$. Now consider the closed formula

$$F(\bar{x}, y) = \bigwedge_{i < \omega} f_i(\bar{x})\downarrow \wedge d(y, f_i(\bar{x})) \leq 2^{-i}.$$

By construction, for any structure \mathfrak{M} and $\bar{a} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, $\mathfrak{M} \models F(\bar{a}, b)$ for a unique $b \in \mathfrak{M}$. Therefore by Proposition 2.3.79, there is a definable partial function $g(\bar{x})$ such that $\Sigma(\bar{x}) \models g(\bar{x}) \downarrow$ and $d_\Sigma(g, f_i) \leq 2^{-i}$ for every $i < \omega$. \square

Now we can finally give an unexciting answer to a minor loose end that arose in Proposition 1.4.4. Namely, what is the metric completion of the space of $\mathcal{L}(V)$ -terms under d_{\equiv} ?

Corollary 2.3.91. *If \bar{x} is an at most countable tuple of variables, $\Sigma(\bar{x})$ is a partial $\mathcal{L}(\bar{x})$ -type, then d_Σ in the sense of Definition 1.4.1 agrees with d_Σ in the sense of Definition 2.3.88. Furthermore, if $\{t_i(\bar{x})\}_{i < \omega}$ is a sequence of $\mathcal{L}(V)$ -terms that is a Cauchy sequence under d_Σ , then there is a definable partial function $f(\bar{x})$ such that for any \mathfrak{M} and \bar{a} such that $\mathfrak{M} \models \Sigma(\bar{a})$, $f^{\mathfrak{M}}(\bar{a}) = \lim_{i \rightarrow \infty} t_i^{\mathfrak{M}}(\bar{a})$.*

Proof. The equivalence of the two definitions follows from Proposition 2.3.89. The furthermore statement follows from Proposition 2.3.90. \square

Although, note that the closure of the set of terms under d_Σ is not in general the whole collection of definable partial functions over Σ , just as it is in discrete logic where not every definable function is witnessed by a term.

2.4 Dictionaric Type Spaces and Theories

It's well known that definable sets are poorly behaved in continuous logic. As it will turn out, totally transcendental theories (and in particular ω -stable theories) have better behavior with regards to definable sets than arbitrary continuous first-order theories,

even relative to strictly superstable theories. The specific property they have seems to be important enough to have its own name.

Definition 2.4.1.

- Let X be a topometric space. We say that X is *dictionary* if for every $p \in X$ and closed $F \subseteq X$ with $p \notin F$, there is a definable set $D \subseteq X$ such that $p \in \text{int}_X D$ and $D \cap F = \emptyset$ (i.e. X has a basis of definable neighborhoods).
- A complete theory T is dictionary if for every $n < \omega$ and parameter set A , $S_n(A)$ is dictionary. \triangleleft

Clearly every discrete theory is dictionary, but, as we have already seen with Proposition 2.3.32, not every continuous theory is dictionary. We will now present a characterization of dictionary type spaces and definable sets.

Proposition 2.4.2. *Let X be a compact topometric space with an open metric. The following are equivalent:*

- (i) X is dictionary.
- (ii) For every $\varepsilon > 0$, X has a basis of open sets U satisfying $\text{cl}(U) \subseteq U^{<\varepsilon}$.
- (iii) Definable sets separate disjoint closed subsets of X .
- (iv) For every closed $F, G \subseteq X$ with $F \cap G = \emptyset$, there is a definable set D such that either $F \subseteq D$ and $D \cap G = \emptyset$ or $G \subseteq D$ and $D \cap F = \emptyset$.
- (v) X has a network of definable sets (i.e. for every $p \in U \subseteq X$, with U open, there is a definable set D such that $p \in D \subseteq U$).

(vi) For every $p \in U \subseteq X$, with U open, and every $\varepsilon > 0$, there is an open set $V \subseteq U$ and a closed set $F \subseteq X$ such that $p \in V$ and $d_H(V, F) < \varepsilon$ (where d_H is the Hausdorff metric on sets).

Proof. We will prove (i) \Rightarrow (v) \Rightarrow (ii) \Rightarrow (i), (i) \Rightarrow (iv) \Rightarrow (iii) \Rightarrow (i), and (ii) \Leftrightarrow (vi).

(i) \Rightarrow (v). This is immediate.

(v) \Rightarrow (ii). For any definable set D and $\varepsilon > 0$, the set $U = D^{<\varepsilon/2}$ satisfies $\text{cl } U \subseteq D^{\leq\varepsilon/2} \subseteq D^{<\varepsilon}$. So since given $x \in D \subseteq U$ we can always find $\varepsilon > 0$ small enough that $D^{\leq\varepsilon/2} \subseteq U$ (by compactness), we have that (ii) holds.

(ii) \Rightarrow (i). Let $x \in X$ be a point with open neighborhood $V \subseteq X$. Find O open such that $x \in O$ and $\text{cl } O \subseteq V$. Find $\varepsilon_0 > 0$ small enough that $(\text{cl } O)^{\leq 2\varepsilon_0} \subseteq V$. Using the assumption of (ii), find an open neighborhood U_0 of x such that $\text{cl } U_0 \subseteq O$ and such that $\text{cl } U_0 \subseteq U_0^{<\varepsilon_0}$. For each $n < \omega$, given U_n , find V_n open such that $\text{cl } U_n \subseteq V_n$ and $\text{cl } V_n \subseteq U_n^{<\varepsilon_n}$. Then find $\varepsilon_{n+1} > 0$ small enough that $(\text{cl } V_n)^{\leq\varepsilon_{n+1}} \subseteq U_n^{<\varepsilon_n}$ and $\varepsilon_{n+1} < \varepsilon_n \downarrow 2^{-n}$. Using the assumption of (ii), for each $y \in \text{cl } U_n$, find an open neighborhood $O_{n,y}$ such that $\text{cl } O_{n,y} \subseteq V_n$ and such that $\text{cl } O_{n,y} \subseteq O_{n,y}^{<\varepsilon_{n+1}}$. By compactness this cover has a finite subcover. Let U_{n+1} be its union. Note that we have that

$$\text{cl } U_n \subseteq \text{cl } U_{n+1} \subseteq U_{n+1}^{<\varepsilon_{n+1}} \subseteq (\text{cl } U_{n+1})^{\leq\varepsilon_{n+1}} \subseteq U_n^{<\varepsilon_n} \subseteq (\text{cl } O)^{\leq 2\varepsilon_0} \subseteq V.$$

Now finally let $D = \bigcap_{n < \omega} (\text{cl } U_n)^{\leq\varepsilon_n} = \bigcap_{n < \omega} U_n^{<\varepsilon_n}$. Clearly $U_0 \subseteq D \subseteq V$, so D is a neighborhood of x that is a sub-neighborhood of V . Furthermore D is clearly closed. To see that D is definable, note that for any $\delta > 0$ there is an n such that $\varepsilon_n < \delta$, so we have that $D^{<\delta} \supseteq U_n^{<\varepsilon_n} \supseteq D$. Therefore $D \subseteq \text{int } D^{<\delta}$ for every $\delta > 0$ and D is definable. Therefore X is dictionaric.

(i) \Rightarrow (iv). Let F and G be disjoint closed sets. For each point $x \in F$, find a definable neighborhood D_x disjoint from G . By compactness finitely many of these cover F . Since the union of finitely many definable sets is definable we have that their union D is a neighborhood of F disjoint from G .

(iv) \Rightarrow (iii). Given F and G , disjoint closed sets, let $U \supseteq F$ and $V \supseteq G$ be disjoint open neighborhoods, and let $L = X \setminus (U \cup V)$ be a closed separator between F and G . Now let D be the guaranteed definable set separating $F \cup G$ and L . If $D \supseteq F \cup G$, then $D \cap U$ and $D \cap V$ are the required definable sets (because if the intersection of a definable set and an open set is closed, then it is definable). If $D \not\supseteq F \cup G$, then consider the disjoint closed sets F and $D \cup V$. Find $U' \supseteq F$ and $V' \supseteq D \cup V$ disjoint open neighborhoods, and let $L' = X \setminus (U' \cup V')$ be a closed separator between F and $D \cup V$. Let E be the guaranteed definable set separating $F \cup D \cup V$ and L' . If $E \supseteq F \cup D \cup V$, then $E \cap U'$ is the required definable set for F . If $E \not\supseteq F \cup D \cup V$, then $E \cup U'$ is the required definable set for F . So in any case we have that there are definable, disjoint E and D such that $F \subseteq E$ and $G \subseteq D$.

(iii) \Rightarrow (i), If $x \in U \subseteq X$ with U open, we can find an open set V such that $x \in V$ and $\overline{V} \subseteq U$. A definable set separating $\text{cl } V$ and $X \setminus U$ is the required definable neighborhood of x .

(ii) \Rightarrow (vi). This is immediate.

(vi) \Rightarrow (ii). Given $x \in V \subseteq X$, with V open, find open V' such that $x \in V'$ and $\text{cl } V' \subseteq V$. Fix $\varepsilon > 0$, find $U \subseteq V'$ with $x \in U$ and such that for some closed set F , $d_H(U, F) < \frac{\varepsilon}{5}$. By the definition of the Hausdorff metric this implies that $U \subseteq \text{cl } U \subseteq F^{\leq \varepsilon/3} \subseteq U^{\leq 2\varepsilon/3} \subseteq U^{< \varepsilon}$, so we get the required basis of open sets. \square

A notable omission in Proposition 2.4.2 is the following plausible separation property:

For every $p, q \in X$, with $p \neq q$, there are definable sets D and E such that $p \in \text{int}_X D$, $q \in \text{int}_X E$, and $D \cap E = \emptyset$.

Since the intersection of definable sets is not in general definable, there is no reason a priori to suppose that this implies dictionaricness (as this is the typical avenue of the proof that compact Hausdorff spaces in which every point has a basis of clopen neighborhoods are zero-dimensional), and indeed it does not follow (Counterexample C.1.12). In fact, every theory is bi-interpretable with a theory with this property (Proposition 2.4.18).

We really only care about this result in the context of certain type spaces, as well as definable subsets of those type spaces.

Proposition 2.4.3 (Hereditariness to Definable Subsets). *If X is a dictionaric type space or definable set, and D is a definable subset of X , then D is dictionaric.*

Proof. Let D be a definable set in X . Let U satisfy $x \in U \subseteq D$ be an open neighborhood of x in D . Find V an open-in- D set such that $x \in V$ and $\text{cl } V \subseteq U$. Then let $\varepsilon > 0$ be small enough that $(\text{cl } V)^{<3\varepsilon} \cap (D \setminus U) = \emptyset$. Since a relatively open subset of a definable set is locatable, we have that $V^{<\varepsilon}$ is open as a subset of X , so by dictionaricness we can find a definable set E such that $x \in \text{int } E$ and $E \subseteq V^{<\varepsilon}$. Now consider the set $F = E^{\leq \varepsilon} \cap D$. By the triangle inequality this is disjoint from $D \setminus U$ and so is a subset of U . Furthermore it is a neighborhood of x in D . Also, since $E \subseteq V^{<\varepsilon}$ and $V \subseteq D$, we have that $E \subseteq F^{<2\varepsilon}$, so $F \subseteq \text{int}_D F^{<4\varepsilon}$. Let O be an open-in- D set such that $F \subseteq O$, $\text{cl } O \subseteq U$, and $\text{cl } O \subseteq \text{int}_D F^{<4\varepsilon}$. Then we have that O is an open neighborhood of x , smaller than U such that $\text{cl } O \subseteq O^{<4\varepsilon}$. Since we can do this for arbitrary x , U , and ε we have by Proposition 2.4.2 that X is dictionaric. \square

The analog of the previous result fails for arbitrary closed $F \subseteq X$, i.e. there is a dictionaric type space with a closed subset that fails to be dictionaric (see Counterexample 15). Compare this to the fact that any closed subspace of a totally disconnected compact Hausdorff space is totally disconnected.

Compare the following Proposition 2.4.4 to the fact that if X is a totally disconnected compact Hausdorff space, $F \subseteq X$ is a closed subset, and $Q \subseteq F$ is relatively clopen, then there is a clopen set $D \subseteq X$ such that $D \cap F = Q$.

Proposition 2.4.4 (Extension). *If X is a dictionaric type space or definable set and we have $Q \subseteq F \subseteq X$ with Q and F closed and Q definable-in- F , then there is a definable set $D \subseteq X$ such that $D \cap F = Q$.*

Proof. Use the following Lemma 2.4.5 with $F_i = F$ for all $i < \omega$ to get some definable set D and formula φ such that $D \cap F \cap \llbracket \varphi \leq 1 \rrbracket = Q$, but $\llbracket \varphi \leq 1 \rrbracket = X$, so $D \cap F = Q$, as required. \square

Although Proposition 2.4.4 is the more attractive statement, we will occasionally need this technical strengthening.

Lemma 2.4.5. *Let $\{F_i\}_{i < \omega}$ be a family of closed sets in X , a type space or definable set (not necessarily dictionaric), and let Q be a closed set such that $Q \subseteq F_i$ and Q is relatively definable in F_i for each $i < \omega$ (i.e. $Q \subseteq \text{int}_{F_i} Q^{<\varepsilon}$ for every $\varepsilon > 0$).*

(i) *There is a closed set $C \subseteq X$ and a $[0, 1]$ -valued formula ψ with $Q \subseteq \llbracket \psi = 0 \rrbracket$, such that*

- $Q \subseteq C$ and for each $\varepsilon > 0$, $Q \subseteq \text{int}_X C^{<\varepsilon}$, and
- for each $i < \omega$, $C \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket = Q$ (in particular, $C \cap F_0 = Q$).

(ii) (Strong Extension) *If X is dictionary, then C can be taken to be definable.*

Proof. In this proof, all instances of $\llbracket \cdot \rrbracket$ are understood to be $\llbracket \cdot \rrbracket_X$. Specifically, they are subsets of X . Also, without loss of generality, assume that $\text{diam}(X) \leq 1$. We may accomplish this by replacing d with $\min\{d, 1\}$. This is still a topometric, is still open, and has the same definable sets.

(i) Without loss of generality we may assume that $F_i \subseteq F_{i+1}$, since we have that if Q is relatively definable in two closed sets F and G , then Q is relatively definable in $F \cup G$ as well. So we can replace F_i by $\bigcup_{j \leq i} F_j$ if necessary. To see that such a Q is relatively definable in $F \cup G$, fix $\varepsilon > 0$, and find $U \subseteq F$ and $V \subseteq G$, each relatively open, such that $Q \subseteq U$ and $Q \subseteq V$ and such that if $a \in U$ or if $a \in V$, then $d(a, Q) < \varepsilon$. $U \cup V$ is a relatively open neighborhood of Q in $F \cup G$ with the same property, so Q is relatively definable in $F \cup G$.

Assume without loss of generality that the metric diameter of X is at most 1. For each $i < \omega$, let φ_i be a $[0, 1]$ -valued formula witnessing in F_i that Q is relatively definable, i.e. $Q \subseteq \llbracket \varphi_i = 0 \rrbracket$ and for every type $p \in F_i$, $d_{\text{inf}}(p, Q) \leq \varphi_i(p)$. These exist by the Tietze extension theorem.⁸ Let $\varphi = \sum_{i < \omega} 2^{-i} \varphi_i$. Note that $\llbracket \varphi = 0 \rrbracket \cap F_i = Q$ for every $i < \omega$.

For each $k < \omega$, let

$$G_k = F_k \cap \llbracket 4^{-k-1} \leq \varphi \leq 4^{-k} \rrbracket,$$

$$U_k = \llbracket \varphi < 4^{-k} \rrbracket, \text{ and}$$

$$V_k = \llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket \cap U_{k+3}^{< 2^{-k+1}}.$$

⁸There is a continuous function on F_i witnessing that Q is relatively definable by Proposition 2.3.15 part (vi). By the Tietze extension theorem (Fact A.2.8), this extends to a continuous function on the whole type space, which is equivalent to some formula.

Note that V_k is an open set. Also note that $G_k \subseteq \llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket$ and that if $p \in G_k$, then $\varphi(p) \leq 4^{-k}$. This implies in particular that $2^{-k}\varphi_k(p) \leq 4^{-k}$, and so $\varphi_k(p) \leq 2^{-k}$ and $d_{\text{inf}}(p, Q) \leq 2^{-k}$. This in turn implies that $p \in Q^{\leq 2^{-k}} \subseteq U_{k+3}^{< 2^{-k+1}}$, and thus $G_k \subseteq V_k$.

Let $A = X \setminus \bigcup_{k < \omega} V_k$. A is clearly closed. Note that $\llbracket \varphi = 0 \rrbracket \cap V_k = \emptyset$ for each $k < \omega$, so we have that $Q \subseteq \llbracket \varphi = 0 \rrbracket \subseteq A$. We want to show that for any $\varepsilon > 0$, $Q \subseteq \text{int}_X A^{< \varepsilon}$.

Assume that $q \notin A$. It must be in V_k for some $k < \omega$. Let $q_0 = q$ and $k(0) = k$. Assume we're given q_ℓ and $k(\ell)$ such that $q_\ell \in V_{k(\ell)}$. By construction there exists $q_{\ell+1} \in U_{k(\ell)+3}$ such that $d(q_\ell, q_{\ell+1}) < 2^{-k(\ell)+1}$. Stop if $q_{\ell+1} \in A$, otherwise $q_{\ell+1} \in V_{k(\ell+1)}$ for some $k(\ell+1)$ strictly larger than $k(\ell)$.

After the construction in either case the total distance traversed along the sequence q_0, q_1, \dots is $\leq \sum_{k \leq m < \omega} 2^{-m+1} = 2^{-k+2}$. If we stopped, then the final point is in A , so $d_{\text{inf}}(q, A) < 2^{-k+3}$. If the sequence never stopped, then it is a Cauchy sequence whose limit, q_ω , by continuity has $\varphi(q_\omega) = 0$, so $q_\omega \in A$ as well. So we have that $\llbracket \varphi = 0 \rrbracket \subseteq \llbracket \varphi < 2^{-k-1} \rrbracket \subseteq A^{< 2^{-k+3}}$, whence $Q \subseteq \llbracket \varphi = 0 \rrbracket \subseteq \text{int}_X A^{< \varepsilon}$ for every $\varepsilon > 0$. Hence we have verified the first bullet point of the lemma for the set A .

Let $\psi = \min\{4\sqrt{\varphi}, 1\}$. This is clearly a $[0, 1]$ -valued formula.

We need to adjust A in order to satisfy the second bullet point in the $i = 0$ case; specifically, we need to ensure that $C \subseteq \llbracket \varphi < 2^{-0} \rrbracket = \llbracket \varphi < 1 \rrbracket$. It is not hard to show that if $O \supseteq Q$ is any open neighborhood such that $\overline{O} \subseteq \llbracket \varphi < 1 \rrbracket$, then $B = A \cap \overline{O}$ still satisfies the property stated in the first bullet point of the lemma and clearly satisfies $B \subseteq \llbracket \varphi < 1 \rrbracket$.

Claim 1. C still satisfies the first bullet point of the lemma.

Proof of claim 1. We will actually show that $A \cap O$ satisfies the first bullet point of the lemma, from which the same follows easily for C .

Fix $x \in Q$. Fix open $Y \ni x$ such that $(\text{cl} Y) \subseteq O$. Find $\varepsilon > 0$ small enough that $(\text{cl} Y)^{\leq \varepsilon} \subseteq O$ (this must exist by compactness). By assumption we have that $x \in \text{int} A^{< \varepsilon}$. Let $Z = Y \cap \text{int} A^{< \varepsilon}$. Note that this is an open neighborhood of x . By construction, $Z^{< \varepsilon} \subseteq (\text{cl} Z)^{\leq \varepsilon} \subseteq (\text{cl} Y)^{\leq \varepsilon} \subseteq O$. Therefore $Z \subseteq (A \cap O)^{< \varepsilon}$, and we have that $x \in \text{int}(A \cap O)^{< \varepsilon}$ for any sufficiently small ε , which implies it for every $\varepsilon > 0$.

Since we can do this for any $x \in Q$, we have that $Q \subseteq \text{int}(A \cap O)^{< \varepsilon}$ for every $\varepsilon > 0$, and likewise we have that C satisfies the first bullet point of the lemma. $\square_{\text{claim 1}}$

To verify the second bullet point, consider $C \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket$. If $i = 0$, then since $C \subseteq \llbracket \psi < 2^{-0} \rrbracket$, we have that for any $p \in C$ that $\psi(p) < 1$, and so $4\sqrt{\varphi(p)} < 1$ (this is the particular point that we needed to shrink A for, given the definition of ψ). From this it follows that $\sqrt{\varphi(p)} < \frac{1}{4}$ and $\varphi(p) < 4^{-2}$. If $i > 0$, then for any $p \in \llbracket \psi \leq 2^{-i} \rrbracket$, we have $4\sqrt{\varphi} \leq 2^{-i}$, so $\sqrt{\varphi(p)} \leq 2^{-i-2}$ and $\varphi(p) \leq 4^{-i-2}$. So for any i , if $p \in C \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket$, then $\varphi(p) \leq 4^{-i-2}$.

Claim 2. For any i , $\llbracket \varphi < 4^{-i-1} \rrbracket \cap C \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$.

Proof of claim 2. Assume that $p \in \llbracket \varphi < 4^{-i-1} \rrbracket \cap A \cap F_i$ but $\varphi(p) > 0$. Then $p \in \llbracket 0 < \varphi < 4^{-i-1} \rrbracket$ and for all $k \geq i$, $p \in F_k$ (since $F_k \supseteq F_i$). So p must be in G_k for some $k > i$, but the G_k are all disjoint from A by construction and therefore also disjoint from C . Thus we have a contradiction. $\square_{\text{claim 2}}$

This implies that $\llbracket \varphi \leq 4^{-i-2} \rrbracket \cap C \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$, as well, since it's a smaller set. Therefore $C \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket \subseteq \llbracket \varphi = 0 \rrbracket$, but $F_i \cap \llbracket \varphi = 0 \rrbracket = Q$, so we have $C \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket = Q$, as required.

(ii) Assume that X is dictionaric. Continuing from just after the definition of ψ in

the proof of part (i), we need to cover $A \setminus \llbracket \varphi = 0 \rrbracket$ by definable sets so that the overall union will be closed (this will be enough to imply that the union is definable) without spoiling the property stated in the second bullet point of the lemma.

For each $k < \omega$, let $H_k = A \cap \llbracket 4^{-k-1} \leq \varphi \leq 4^{-k} \rrbracket$. Note that by construction $H_k \cap F_k = \emptyset$. Let $W_k = \llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket \setminus F_k$, which is an open neighborhood of H_k . Let D_k be a definable set such that $H_k \subseteq D_k \subseteq W_k$. Finally let $E = A \cup \bigcup_{k < \omega} D_k$.

First to see that E is closed, note that any convergent net $\{q_i\}_{i \in I}$ in E either eventually stays within some W_k or has $\lim_{i \in I} \varphi(q_i) = 0$. In the first case, $E \cap (W_{k-1} \cup W_k \cup W_{k+1})$ is relatively closed in $W_{k-1} \cup W_k \cup W_{k+1}$, since in that set it is a finite union of closed sets, so the net converges to a point in E . In the second case, the net must be converging to a point in $\llbracket \varphi = 0 \rrbracket \subseteq A \subseteq E$ by continuity.

To see that E is definable, note that for any $\varepsilon > 0$,

$$\begin{aligned} \text{int}_X E^{<\varepsilon} &= \text{int}_X \left[A \cup \bigcup_{k < \omega} D_k \right]^{<\varepsilon} \supseteq \text{int}_X A^{<\varepsilon} \cup \bigcup_{k < \omega} \text{int}_X D_k^{<\varepsilon} \\ &\supseteq \llbracket \varphi = 0 \rrbracket \cup \bigcup_{k < \omega} D_k = E. \end{aligned}$$

Since E is a definable set and X is dictionaric, we have that E is dictionaric as well. Let $D \subseteq E$ be a definable set such that $Q \subseteq D \subseteq \llbracket \psi < 1 \rrbracket$. (This is only necessary to handle the $i = 0$ case.) Consider $D \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket$. If $i = 0$, then since $D \subseteq \llbracket \psi < 2^{-0} \rrbracket$ we have that for any $p \in D$, $\psi(p) < 1$, so $4\sqrt{\varphi(p)} < 1$, $\sqrt{\varphi(p)} < \frac{1}{4}$, and $\varphi(p) < 4^{-2}$. If $i > 0$, then $2^{-i} < 1$, so $4\sqrt{\varphi(p)} \leq 2^{-i}$ and we have $\sqrt{\varphi(p)} \leq 2^{-i-2}$ and $\varphi(p) \leq 4^{-i-2}$. So for any $i < \omega$ we have that if $p \in D \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket$, then $\varphi(p) \leq 4^{-i-2}$.

Claim 3. For each $i < \omega$, $\llbracket \varphi < 4^{-i-1} \rrbracket \cap D \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$.

Proof of claim 3. Note that by construction

$$\llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket \cap E = \llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket \cap \bigcup_{k-2 \leq \ell \leq k+2} D_\ell.$$

so we have that $\llbracket 4^{-k-2} < \varphi < 4^{-k+1} \rrbracket \cap E \cap F_{k-2} = \emptyset$, since D_ℓ , for $\ell \geq k-2$, are all disjoint from F_{k-2} (because $F_m \subseteq F_{m+1}$ for all $m < \omega$). This implies that $\llbracket 0 < \varphi < 4^{-k+1} \rrbracket \cap E \cap F_{k-2} = \emptyset$ as well. Therefore $\llbracket \varphi < 4^{-k+1} \rrbracket \cap E \cap F_{k-2} \subseteq \llbracket \varphi = 0 \rrbracket$. Shifting by 2 gives $\llbracket \varphi < 4^{-i-1} \rrbracket \cap D \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$. $\square_{\text{claim 3}}$

Given claim 3, we also have $\llbracket \varphi \leq 4^{-i-2} \rrbracket \cap D \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$ (since this is a smaller set). This implies that $\llbracket \psi \leq 2^{-i} \rrbracket \cap D \cap F_i \subseteq \llbracket \varphi = 0 \rrbracket$, but $F_i \cap \llbracket \varphi = 0 \rrbracket = Q$, so we have $D \cap F_i \cap \llbracket \psi \leq 2^{-i} \rrbracket = Q$, as required. So we can satisfy part (ii) by setting $C = D$. \square

Lemma 2.4.6. *If X is a dictionary type space or definable set, $D \subseteq X$ is a definable set, $F \subseteq X$ is a closed set, $U \subseteq X$ is a (relatively) open set such that $F \subseteq U$, and $E \subseteq D$ is a definable set such that $F \cap D \subseteq E$, then there exists a definable set E' such that $F \subseteq E' \subseteq U$ and $D \cap E' = E$.*

Proof. We need to argue that $E \cup F$ is relatively definable in $D \cup F \cup (X \setminus U)$, then the result follows from Proposition 2.4.4.

To see that $E \cup F$ is relatively definable in $D \cup F \cup (X \setminus U)$, first note that since E is definable in D , E is definable in X . Now notice that any net $\{x_i\}_{i \in I}$ converging to a point in $E \cup F$ must either be limiting to a point in G or be eventually contained in $E \cup F$, so in either case we have that $\lim_{i \in I} d_{\text{inf}}(x, G \cup F) = 0$, so $E \cup F$ is relatively definable. Now let E' be an extension of $E \cup F$ to all of X . We have that E' is the required definable set. \square

Proposition 2.4.7 (Approximate Intersection). *If X is a dictionary type space or definable set, D is a definable subset of X , $F \subseteq X$ is a closed set, and $U \supseteq F$ is an open-in- X set, then there is a definable set E such that $F \subseteq E \subseteq U$, and such that $D \cap E$ is definable.*

Proof. First since D is definable it is itself dictionary by Proposition 2.4.3. Let $G \subseteq D$ be a definable set such that $G \subseteq U$ and $\text{int}_X G \supseteq F \cap D$. Then we can apply Lemma 2.4.6. \square

Corollary 2.4.8 (Approximate Intersection of Definable Sets). *If D and E are definable sets in a dictionary type space, then for any $\varepsilon > 0$, there is a definable set $E' \supseteq E$ with $d_H(E, E') < \varepsilon$ such that $D \cap E'$ is also definable.*

Proof. If $E \subseteq E' \subseteq E^{<\varepsilon/2}$, then $d_H(E, E') \leq \frac{\varepsilon}{2} < \varepsilon$. \square

Although Corollary 2.4.8 is an easy consequence of Proposition 2.4.7, for definable sets D and E , $d_H(D, E) < \varepsilon$ can be a very useful condition. It implies that quantifying over elements of D is ‘approximately the same’ as quantifying over elements of E to within an accuracy of ε . If D is a uniformly discrete set, then there is an $\varepsilon > 0$ such that $d_H(D, E) < \varepsilon$ implies that there is a definable equivalence relation \sim on E such that E/\sim has a canonical bijection with D .

Proposition 2.4.9. *If X is a dictionary type space or definable set and ρ is a definable pseudo-metric on X , then X/ρ is dictionary.*

Proof. Let $q \in U \subseteq X/\rho$ be a type with open neighborhood U , consider $\pi^{-1}(q) \subseteq \pi^{-1}(U)$ where $\pi : X \rightarrow X/\rho$ is the natural projection map. By continuity $\pi^{-1}(q)$ is closed and $\pi^{-1}(U)$ is open. Let D be a definable set such that $\pi^{-1}(q) \subseteq D$ and $D \cap U = \emptyset$. The

projection $\pi(D)$ is a definable set (with distance function $\inf_{y \in D} \rho(x, y)$). It's clearly contained in U and contains q , so X/ρ has a network of definable sets, and by Proposition 2.4.2, X/ρ is *dictionary*. \square

Proposition 2.4.10. *If $S_n(T)$ is dictionary for every $n < \omega$, then $S_\omega(T)$ is dictionary as well.*

Proof. This follows from the fact that given any type $p \in S_\omega(T)$ and an open neighborhood U , there is a restricted formula $\varphi(\bar{x})$ such that $p(\bar{x}, \dots) \in \llbracket \varphi(\bar{x}) < \frac{1}{3} \rrbracket$ and $\llbracket \varphi(\bar{x}) \leq \frac{2}{3} \rrbracket \subseteq U$. By dictionaryness of $S_{\bar{x}}(T)$ we can find a definable set D such that $\llbracket \varphi(\bar{x}) < \frac{1}{3} \rrbracket \subseteq D \subseteq \llbracket \varphi(\bar{x}) \leq \frac{2}{3} \rrbracket$ and this set is still definable in the sort of ω -tuples under the metric given in Definition 1.2.1, so $S_\omega(T)$ is *dictionary* as well. \square

It is natural to think about strengthenings of dictionaryness, for instance one might hope for a version that is uniform in parameters, such as in the following definition.

Definition 2.4.11. A (possibly incomplete) theory T is *uniformly parametrically dictionary* if for any finite tuples of variables \bar{v} and \bar{x} and any pair of inconsistent closed formulas $F(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$, there is a \bar{x} -uniformly definable set $D(\bar{v}, \bar{x})$ (over T) such that $F(\bar{v}, \bar{x}) \models D(\bar{v}, \bar{x})$ and $D(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$ are inconsistent. \triangleleft

This is equivalent to requiring that each type space $S_{\bar{v}\bar{x}}(T)$ be *dictionary* with regards to the metric $d_{/\bar{x}}$. As it turns out this is an extremely strong condition.

Proposition 2.4.12. *A theory T is uniformly parametrically dictionary if and only if $S_{\bar{x}}(T)$ is totally disconnected for every tuple of variables \bar{x} .*

Proof. If T has totally disconnected type spaces, then by compactness, if $F(\bar{v}, \bar{x})$ and

$G(\bar{v}, \bar{x})$ are inconsistent closed formulas, we can find a clopen formula $Q(\bar{v}, \bar{x})$ such that $F(\bar{v}, \bar{x}) \models Q(\bar{v}, \bar{x})$ and $Q(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$ are inconsistent.

Conversely, assume that T is uniformly parametrically dictionaric. Fix a finite tuple of variables \bar{x} , and let v be a distinct variable. For any pair of inconsistent closed $\mathcal{L}(\bar{x})$ -formulas, $F(\bar{x})$ and $G(\bar{x})$, by assumption we can find a \bar{x} -uniformly definable set $D(v, \bar{x})$ such that $F(\bar{x}) \models D(v, \bar{x})$ and $D(v, \bar{x})$ and $G(\bar{x})$ are inconsistent. By Proposition 2.3.18, the formula $Q(\bar{x}) = (\exists v \in D(\cdot, \bar{x}))(v = v)$ is clopen and we have necessarily that $F(\bar{x}) \models Q(\bar{x})$ and that $Q(\bar{x})$ and $G(\bar{x})$ are inconsistent. Therefore $S_{\bar{x}}(T)$ has a basis of clopen sets and is totally disconnected. This implies that $S_{\bar{x}}(T)$ is totally disconnected for every finite tuple of variables \bar{x} , which in turn implies the same for arbitrary tuples of variables. \square

We can weaken this, as in the following definition. This definition uses concepts from Section 3.5, but we have included it here for the sake of organization.

Definition 2.4.13. For any (possibly incomplete) theory T , any at most countable tuple \bar{v} of variables, and arbitrary tuple \bar{x} of variables, a closed type-set formula $F(\bar{v}, \bar{x})$ is *\bar{x} -pointwise definable over T* if for any tuple of parameters $\bar{a} \in \mathfrak{M} \models T$ assigned to \bar{x} , $F(\bar{v}, \bar{a})$ is a definable set over $\text{Th}(\mathfrak{M})$.

A (possibly incomplete) theory T is *pointwise parametrically dictionaric* if for finite tuples \bar{v} and \bar{x} and any pair of inconsistent closed formulas $F(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$, there is a \bar{x} -pointwise definable set $D(\bar{v}, \bar{x})$ over T such that $T, F(\bar{v}, \bar{x}) \models D(\bar{v}, \bar{x})$ and $D(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$ are inconsistent.

A (possibly incomplete) theory T is *almost uniformly parametrically dictionaric* if for finite tuples \bar{v} and \bar{x} and any pair of inconsistent closed formulas $F(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$,

there is a formula $\varphi(\bar{v}, \bar{x}, r)$ that gives an almost \bar{x} -uniformly definable set over T such that $T, F(\bar{v}, \bar{x}) \models \varphi(\bar{v}, \bar{x}, r) = 0$ and $\varphi(\bar{v}, \bar{x}, r) = 0$ and $G(\bar{v}, \bar{x})$ are inconsistent. \triangleleft

It turns out that these weakened definitions are equivalent to dictionaricness and to each other, and it allows us to characterize dictionaricness of a theory in terms of the properties of the type spaces $S_n(T)$ for $n < \omega$, rather than those of $S_n(A)$ for arbitrary sets of parameters A .⁹

Theorem 2.4.14. *For a (possibly incomplete) theory T , the following are equivalent:*

- (i) *Every completion of T is dictionaric.*
- (ii) *T is pointwise parametrically dictionaric.*
- (iii) *T is almost uniformly parametrically dictionaric.*
- (iv) *For every restricted formula $\varphi(\bar{v}, \bar{x})$ and every rational $\varepsilon > 0$, there is a finite list $\delta_0(\bar{v}, \bar{x}), \dots, \delta_{n-1}(\bar{v}, \bar{x})$ of restricted formulas and a finite list $\gamma_0, \dots, \gamma_{n-1}$ of rational numbers such that $0 < \gamma_i < \varepsilon$ for every $i < n$ and such that*

$$T \models \forall \bar{x} \bigvee_{i < n} U_{\varphi, \delta_i, \gamma_i}(\bar{x}),$$

where

$$\begin{aligned} U_{\varphi, \delta, \gamma}(\bar{x}) = & \forall \bar{v} \left[(\varphi(\bar{v}, \bar{x}) = 0 \rightarrow \delta(\bar{v}, \bar{x}) < 2\gamma) \right. \\ & \wedge (\varphi(\bar{v}, \bar{x}) = 1 \rightarrow \delta(\bar{v}, \bar{x}) > 3\gamma) \\ & \left. \wedge \left(\text{def}_{\bar{v}} \delta(\bar{v}, \bar{x}) < \gamma \right) \right]. \end{aligned}$$

⁹Although an easy compactness argument shows that it is sufficient to check for dictionaricness of a theory over finite sets of parameters.

Proof. (ii) \Rightarrow (i) is obvious.

(iii) \Rightarrow (ii): For any inconsistent closed formulas $F(\bar{v}, \bar{x})$ and $G(\bar{v}, \bar{x})$ let $\varphi(\bar{v}, \bar{x}, \bar{r})$ be a formula that gives an almost \bar{x} -uniformly definable set over T and for which $T, F(\bar{v}, \bar{x}) \models \varphi(\bar{v}, \bar{x}, \bar{r}) = 0$ and $\{\varphi(\bar{v}, \bar{x}, \bar{r}) = 0, G(\bar{v}, \bar{x})\}$ is inconsistent. Then it follows that $T, F(\bar{v}, \bar{x}) \models \exists r \varphi(\bar{v}, \bar{x}, r) = 0 \wedge \text{DEF} \bar{u} \varphi(\bar{u}, \bar{x}, r)$. To see this, let $\bar{a}\bar{b} \in \mathfrak{M} \models T$ be such that $\mathfrak{M} \models F(\bar{a}, \bar{b})$. Then by assumption, $\mathfrak{M} \models \varphi(\bar{a}, \bar{b}, c) = 0$ for any c in the relevant compact sort. Also by assumption there is a c in the sort such that $\mathfrak{M} \models \text{DEF} \bar{u} \varphi(\bar{u}, \bar{b}, c)$, so we have that $F(\bar{a}, \bar{b}, c) \models \varphi(\bar{a}, \bar{b}, c) = 0 \wedge \text{DEF} \bar{u} \varphi(\bar{u}, \bar{b}, c)$. Since we can do this for any model and any tuples satisfying F , we have the required implication. By Lemma 3.5.5, $H(\bar{v}, \bar{x}) = \exists r \varphi(\bar{v}, \bar{x}, r) = 0 \wedge \text{DEF} \bar{u} \varphi(\bar{u}, \bar{x}, r)$ is \bar{x} -pointwise definable over T , and we have that $T, F(\bar{v}, \bar{x}) \models H(\bar{V}, \bar{x})$. Furthermore, $H(\bar{v}, \bar{x})$ is necessarily inconsistent with $G(\bar{v}, \bar{x})$, so we have that T is pointwise parametrically dictionaric.

(iv) \Rightarrow (iii): Fix inconsistent closed formulas $G_0(\bar{v}, \bar{x})$ and $G_1(\bar{v}, \bar{x})$. By compactness there is a restricted formula $\varphi(\bar{v}, \bar{x})$ such that $\llbracket G_0(\bar{v}, \bar{x}) \rrbracket_T \subseteq \llbracket \varphi(\bar{v}, \bar{x}) = 0 \rrbracket_T$ and $\llbracket G_1(\bar{v}, \bar{x}) \rrbracket_T \subseteq \llbracket \varphi(\bar{v}, \bar{x}) = 1 \rrbracket_T$.

Now we will build a finitely branching tree $T \subseteq \omega^{<\omega}$. Each node $\sigma \in T$ will be labeled with a triple $\langle \eta_\sigma(\bar{v}, \bar{x}), F_\sigma(\bar{x}), \varepsilon_\sigma \rangle$, where η_σ is a restricted real formula, F_σ is a closed formula, and ε_σ is a positive rational number. For the root node set $\eta_\emptyset = [\varphi]_0^1$, $F_\emptyset(\bar{x}) = \top$, and $\varepsilon_\emptyset = \frac{1}{4}$.

At each node $\sigma \in T$, given $\langle \eta_\sigma, F_\sigma, \varepsilon_\sigma \rangle$, find a restricted formula $\psi_\sigma(\bar{x}, \bar{y})$ such that $\llbracket \eta_\sigma \leq 2\varepsilon_\sigma \rrbracket \subseteq \llbracket \psi_\sigma = 0 \rrbracket$ and $\llbracket \eta_\sigma \geq 3\varepsilon_\sigma \rrbracket \subseteq \llbracket \psi_\sigma = 1 \rrbracket$. By assumption, we can find restricted formulas $\delta_0^\sigma(\bar{v}, \bar{x}), \dots, \delta_{n_\sigma-1}^\sigma(\bar{v}, \bar{x})$ and rational numbers $\gamma_0^\sigma, \dots, \gamma_{n_\sigma-1}^\sigma$, satisfying $0 < \gamma_i^\sigma < \varepsilon_\sigma$ for each $i < n_\sigma$, such that $F_\sigma(\bar{x}) \models \bigvee_{i < n_\sigma} U_{\psi_\sigma, \delta_i^\sigma, \gamma_i^\sigma}(\bar{x})$. By Fact A.2.16, we can find a sequence of closed formulas $F_{\sigma \smallfrown i}(\bar{x})$ for $i < n_\sigma$ such that $F_{\sigma \smallfrown i} \models U_{\psi_\sigma, \delta_i^\sigma, \gamma_i^\sigma}$

and $\llbracket F_\sigma \rrbracket = \bigcup_{i < n_\sigma} \llbracket F_{\sigma \smallfrown i} \rrbracket$. For each $i < n_\sigma$ such that $F_{\sigma \smallfrown i}$ is consistent, add the node $\sigma \smallfrown i$ to T , labeled by $\langle \delta_i^\sigma, F_{\sigma \smallfrown i}, \min\{\gamma_i^\sigma, 2^{-|\sigma|}\} \rangle$.

Let X be the compact Hausdorff space of paths through T , for any $\sigma \in \omega^{<\omega}$, let $[\sigma]$ be the clopen set of extensions of σ in X . For each $i < \omega$, let $H_i \subseteq S_{\bar{x}}(T) \times X$ be the union of all sets of the form $\llbracket F_\sigma \rrbracket \times [\sigma]$ for $\sigma \in T$ with $|\sigma| = i$. By construction, we have that $H_{i+1} \subseteq H_i$, so let $H_\omega = \bigcap_{i < \omega} H_i$. For any $p \in S_{\bar{x}}(T)$, we have that $(\{p\} \times X) \cap H_i$ is non-empty for every $i < \omega$, so $(\{p\} \times X) \cap H_\omega$ is non-empty as well, by compactness. Let

$$H = \{ \langle p, \alpha \rangle \in S_{\bar{v}\bar{x}}(T) \times X : \langle p \upharpoonright \bar{x}, \alpha \rangle \in H_\omega \}.$$

(Note that this is closed, since it is the pre-image of H_ω under a continuous function.)

For each $i < \omega$, let $f_i : H \rightarrow \mathbb{R}$ be the function defined by $f_i(p, \alpha) = \eta_{\alpha \upharpoonright i}(p)$. Note that, by construction, each f_i is a continuous function. We need to argue that the sequence $\{f_i\}_{i < \omega}$ is converging in the uniform norm.

Fix $\langle p, \alpha \rangle \in H$. For any $i > 0$, we have that $p(\bar{v}, \bar{x}) \models \text{def}_{\bar{v}} \eta_{\alpha \upharpoonright i}(\bar{v}, \bar{x}) < \varepsilon_{\alpha \upharpoonright i} \leq 2^{-i}$, $\llbracket \eta_{\alpha \upharpoonright i} \leq 2\varepsilon_{\alpha \upharpoonright i} \rrbracket \subseteq \llbracket \eta_{\alpha \upharpoonright i+1} < 2\varepsilon_{\alpha \upharpoonright i+1} \rrbracket$, and $\llbracket \eta_{\alpha \upharpoonright i} \geq 3\varepsilon_{\alpha \upharpoonright i} \rrbracket \subseteq \llbracket \eta_{\alpha \upharpoonright i+1} > 3\varepsilon_{\alpha \upharpoonright i+1} \rrbracket$. Fix \mathfrak{M} and $\bar{a}\bar{b} \in \mathfrak{M}$ such that $\mathfrak{M} \models p(\bar{a}, \bar{b})$. For any $i < \omega$ such that $5 \cdot 2^{-i} < \text{db}(\mathcal{L})$, let $\eta_{\alpha \upharpoonright i}(p) = r$. Since $p(\bar{v}, \bar{x}) \models \text{def}_{\bar{v}} \eta_{\alpha \upharpoonright i}(\bar{v}, \bar{x}) < 2^{-i}$, there are two cases.

Case 1 ($\eta_{\alpha \upharpoonright i}$ is approximately the distance predicate of the empty set): $p(\bar{v}, \bar{x}) \models \sup_{\bar{x}} |\eta_{\alpha \upharpoonright i}(\bar{v}, \bar{x}) - \text{db}(\mathcal{L})| < 2^{-i}$. Since $4 \cdot 2^{-i} < \text{db}(\mathcal{L})$, this implies that $\eta_{\alpha \upharpoonright i}(\bar{a}, \bar{c}) > \frac{4}{5} \text{db}(\mathcal{L}) > 4 \cdot 2^{-i}$ for all $\bar{v} \in \mathfrak{M}$, therefore $\llbracket \eta_{\alpha \upharpoonright i}(\bar{a}, \bar{x}) \geq 3\varepsilon_{\alpha \upharpoonright i} \rrbracket$ is all of $S_{\bar{x}}(\bar{a})$, and we must have that $\llbracket \eta_{\alpha \upharpoonright i+1}(\bar{a}, \bar{x}) < 2\varepsilon_{\alpha \upharpoonright i+1} \rrbracket$ is empty. This implies that $p(\bar{v}, \bar{x}) \models \sup_{\bar{x}} |\eta_{\alpha \upharpoonright i+1}(\bar{v}, \bar{x}) - \text{db}(\mathcal{L})| < 2^{-i-1}$ (i.e. case 1 holds for $i+1$ as well), so we have that $|\eta_{\alpha \upharpoonright i}(p) - \eta_{\alpha \upharpoonright i+1}(p)| < 2^{-i} + 2^{-i-1} < 2^{-i+1}$.

Case 2 ($\eta_{\alpha|i}$ is approximately the distance predicate of a non-empty set):

$$p(\bar{v}, \bar{x}) \models \sup_{\bar{v}} \max\left\{ \inf_{\bar{w}} \max\{|\eta_{\alpha|i}(\bar{w}, \bar{x})|, |d(\bar{v}, \bar{w}) - \eta_{\alpha|i}(\bar{v}, \bar{x})|\}\right\} \\ \sup_{\bar{w}} \eta_{\alpha|i}(\bar{v}, \bar{x}) \div (\eta_{\alpha|i}(\bar{w}, \bar{x}) + d(\bar{v}, \bar{w})) \} < \varepsilon_{\alpha|i}.$$

This implies that $p(\bar{v}, \bar{x}) \models \inf_{\bar{w}} \max\{|\eta_{\alpha|i}(\bar{w}, \bar{x})|, |d(\bar{v}, \bar{w}) - \eta_{\alpha|i}(\bar{v}, \bar{x})|\} < \varepsilon_{\alpha|i}$, so we can find $\bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models \max\{|\eta_{\alpha|i}(\bar{c}, \bar{b})|, |d(\bar{a}, \bar{c}) - \eta_{\alpha|i}(\bar{a}, \bar{b})|\} < \varepsilon_{\alpha|i}$, so in particular $\text{tp}(\bar{c}\bar{b}) \in \llbracket \eta_{\alpha|i} \leq 2\varepsilon_{\alpha|i} \rrbracket \subseteq \llbracket \eta_{\alpha|i+1} < 2\varepsilon_{\alpha|i+1} \rrbracket$, or in other words $\mathfrak{M} \models \eta_{\alpha|i+1}(\bar{c}, \bar{b}) < 2\varepsilon_{\alpha|i+1}$. This implies that case 2 holds for $i + 1$ by the contrapositive of part of the argument in case 1. Therefore $\mathfrak{M} \models \inf_{\bar{w}} \max\{|\eta_{\alpha|i+1}(\bar{w}, \bar{b})|, |d(\bar{c}, \bar{w}) - \eta_{\alpha|i+1}(\bar{c}, \bar{b})|\} < \varepsilon_{\alpha|i+1}$, and we can find $\bar{e} \in \mathfrak{M}$ such that $\mathfrak{M} \models \max\{|\eta_{\alpha|i+1}(\bar{e}, \bar{b})|, |d(\bar{c}, \bar{e}) - \eta_{\alpha|i+1}(\bar{c}, \bar{b})|\} < \varepsilon_{\alpha|i}$, so in particular $d(\bar{c}, \bar{e}) < 2\varepsilon_{\alpha|i+1} + \varepsilon_{\alpha|i+1} = 3\varepsilon_{\alpha|i+1}$.

Finally, we also have that $p(\bar{v}, \bar{x}) \models \sup_{\bar{v}} \sup_{\bar{w}} \eta_{\alpha|i}(\bar{v}, \bar{x}) \div (\eta_{\alpha|i}(\bar{w}, \bar{x}) + d(\bar{v}, \bar{w})) < \varepsilon_{\alpha|i}$, as well as the same for $i + 1$. This is equivalent to saying that $\eta_{\alpha|i}(\bar{v}, \bar{a})$ and $\eta_{\alpha|i}(\bar{v}, \bar{a})$ are 1-Lipschitz whenever \bar{a} satisfies $p \upharpoonright \bar{x}$.

Putting this all together we get

$$\begin{aligned} \eta_{\alpha|i+1}(\bar{a}, \bar{b}) &< d(\bar{a}, \bar{e}) + \varepsilon_{\alpha|i+1} \\ &< d(\bar{a}, \bar{c}) + 3\varepsilon_{\alpha|i+1} + \varepsilon_{\alpha|i+1} \\ &< \eta_{\alpha|i} + \varepsilon_{\alpha|i} + 3\varepsilon_{\alpha|i+1} + \varepsilon_{\alpha|i+1} \\ &< \eta_{\alpha|i} + 5 \cdot 2^{-i}. \end{aligned}$$

For the other direction, since case 2 holds for $i + 1$, we have that there is a $\bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models \eta_{\alpha|i+1} < \varepsilon_{\alpha|i+1}$ and $\mathfrak{M} \models |d(\bar{a}, \bar{c}) - \eta_{\alpha|i+1}(\bar{a}, \bar{b})| < \varepsilon_{\alpha|i+1}$. This implies

that $\text{tp}(\bar{c}, \bar{b}) \in \llbracket \eta_{\alpha|i+1}(\bar{x}, \bar{b}) \leq 3\varepsilon_{\alpha|i+1} \rrbracket \subseteq \llbracket \eta_{\alpha|i}(\bar{x}, \bar{b}) < 3\varepsilon_{\alpha|i} \rrbracket$, or in other words, $\mathfrak{M} \models \eta_{\alpha|i}(\bar{c}, \bar{b}) < 3\varepsilon_{\alpha|i}$. There exists $\bar{e} \in \mathfrak{M}$ such that $\eta_{\alpha|i}(\bar{e}, \bar{b}) < \varepsilon_{\alpha|i}$ and $|d(\bar{c}, \bar{e}) - \eta_{\alpha|i}(\bar{c}, \bar{a})| < \varepsilon_{\alpha|i}$, implying that $d(\bar{c}, \bar{e}) < \varepsilon_{\alpha|i} + 3\varepsilon_{\alpha|i} = 4\varepsilon_{\alpha|i}$.

Putting this all together we get

$$\begin{aligned}
\eta_{\alpha|i}(\bar{a}, \bar{b}) &< \eta_{\alpha|i}(\bar{e}, \bar{b}) + d(\bar{e}, \bar{a}) \\
&< \varepsilon_{\alpha|i} + d(\bar{e}, \bar{c}) + d(\bar{c}, \bar{a}) \\
&< \varepsilon_{\alpha|i} + 4\varepsilon_{\alpha|i} + d(\bar{c}, \bar{a}) \\
&< \varepsilon_{\alpha|i} + 4\varepsilon_{\alpha|i} + \eta_{\alpha|i+1}(\bar{a}, \bar{b}) + \varepsilon_{\alpha|i+1} \\
&\eta_{\alpha|i+1}(\bar{a}, \bar{b}) + 6 \cdot 2^{-i}.
\end{aligned}$$

Therefore we have that $|\eta_{\alpha|i}(\bar{a}, \bar{b}) - \eta_{\alpha|i+1}(\bar{a}, \bar{b})| < 6 \cdot 2^{-i}$.

So in either case, $|\eta_{\alpha|i}(p) - \eta_{\alpha|i+1}(p)| < 6 \cdot 2^{-i}$. Since we can do this for any $\langle p, \alpha \rangle$, we have that f_i converges uniformly on H . Let f_ω be its limit.

By construction, we have that $(\llbracket G_0 \rrbracket \times X) \cap H \subseteq \{f_\omega = 0\}$ and $(\llbracket G_1 \rrbracket \times X) \cap H \cap \{f_\omega = 0\} = \emptyset$. Therefore, by compactness there is an $r > 0$ such that $(\llbracket G_1 \rrbracket \times X) \cap H \subseteq \{f_\omega \geq r\}$. This implies that we can extend f_ω to a continuous function g on $H \cup (\llbracket G_0 \rrbracket \times X) \cup (\llbracket G_1 \rrbracket \times X)$ such that for all $\langle p, \alpha \rangle \in \llbracket G_0 \rrbracket \times X$, $g(p, \alpha) = 0$, and for all $\langle p, \alpha \rangle \in \llbracket G_1 \rrbracket \times X$, $g(p, \alpha) \geq r$. Finally, we can extend g to a continuous function h on all of $S_{\bar{v}\bar{x}}(T) \times X$. We can then find a real formula $\chi(\bar{v}, \bar{x}, r)$, with r a variable on the compact parameter sort K_X , such that χ agrees with h as a function on $S_{\bar{v}\bar{x}}(T) \times X$ (identifying X with $S_r(X)$).

Now we have that $G_0(\bar{v}, \bar{x}) \models \chi(\bar{v}, \bar{x}, r) = 0$, $G_1(\bar{v}, \bar{x})$ and $\chi(\bar{v}, \bar{x}, r) = 0$ are inconsistent, and $\chi(\bar{v}, \bar{x}, r)$ gives an almost \bar{x} -uniformly definable set over T .

So, since we can do this for any inconsistent $G_0(\bar{v}, \bar{x})$ and $G_1(\bar{v}, \bar{x})$, we have that T is almost uniformly parametrically dictionaric, as required.

(i) \Rightarrow (iv): Fix a restricted formula $\varphi(\bar{v}, \bar{x})$ and a rational $\varepsilon > 0$. For each type $p(\bar{x}) \in S_{\bar{x}}(T)$, find \mathfrak{M} and \bar{a} such that $\mathfrak{M} \models p(\bar{a})$. Let T_p be the set of closed sentences contained in p . By dictionaricness, there is a definable set $D(\bar{v}, \bar{a})$ such that $\llbracket \varphi(\bar{v}, \bar{a}) = 0 \rrbracket_{T_p} \subseteq \llbracket D(\bar{v}, \bar{a}) \rrbracket_{T_p}$ and such that $D(\bar{v}, \bar{a})$ and $\varphi(\bar{v}, \bar{a}) = 1$ are inconsistent. By compactness, we can find a restricted formula $\delta_p(\bar{v}, \bar{x})$ and a rational $\gamma_p > 0$ with $\gamma_p < \varepsilon$ such that

$$\begin{aligned} \llbracket D(\bar{v}, \bar{a}) \rrbracket_{T_p} &\subseteq \llbracket \delta_p(\bar{v}, \bar{a}) < 2\gamma_p \rrbracket_{T_p} \subseteq \llbracket \delta_p(\bar{v}, \bar{a}) \leq 2\gamma_p \rrbracket_{T_p} \subseteq \llbracket D(\bar{v}, \bar{a}) < \varepsilon \rrbracket_{T_p} \text{ and} \\ \llbracket \varphi(\bar{v}, \bar{x}) = 1 \rrbracket_{T_p} &\subseteq \llbracket \delta(\bar{v}, \bar{x}) > 3\gamma_p \rrbracket_{T_p}. \end{aligned}$$

So in particular, $p(\bar{x}) \models U_{\varphi, \delta_p, \gamma_p}(\bar{x})$. This is an open formula, so by compactness, there is a finite list p_0, \dots, p_{n-1} of types such that $\bigcup_{i < n} \llbracket U_{\varphi, \delta_{p_i}, \gamma_{p_i}}(\bar{x}) \rrbracket_T = S_{\bar{x}}(T)$. If we set $\delta_i = \delta_{p_i}$ and $\gamma_i = \gamma_{p_i}$, then we have that $T \models \forall \bar{x} \bigvee_{i < n} U_{\varphi, \delta_i, \gamma_i}(\bar{x})$, which is precisely the required condition. \square

Corollary 2.4.15. *In a countable signature \mathcal{L} , for any theory T , the set of completions $T' \in S_0(T)$ which are dictionaric is G_δ (i.e. a countable intersection of open sets).*

Much later, in Theorem 5.1.1, we will show that a theory for which $S_n(T)$ is totally disconnected for every $n < \omega$ has $S_n(A)$ totally disconnected for every n and parameter set A (actually, we really only need that $S_2(T)$ is totally disconnected). Dictionaricness is a weakening of total disconnectedness, and so a natural question arises.

Question 2.4.16. *If $S_n(T)$ is dictionaric for every $n < \omega$, does it follow that T is dictionaric?*

The characterization of theories with totally disconnected type spaces takes advantage of the fact that there is a canonical obstruction to total disconnectedness, namely a non-trivial continuum as a subspace. This raises the question as to whether or not there is a canonical obstruction or class of obstructions to dictionariness.

Question 2.4.17. *Is there a simple class of (ideally compact) topometric spaces \mathcal{O} such that a type space $S_n(T)$ fails to be dictionaric if and only if there is some $X \in \mathcal{O}$ which embeds into $S_n(T)$?*

We will not need the following proposition for anything. We have just included it for completeness.

Proposition 2.4.18. *Every complete theory T is bi-interpretable with a theory T° with the property that for any $n < \omega$, any set of parameters A , and any distinct $p, q \in S_n(A)$, there are disjoint definable sets $D, E \subseteq S_n(A)$ such that $p \in D$ and $q \in E$.*

Proof. If models of T have a single element, then this statement is trivial, so assume that models of T have more than one element.

Let \mathfrak{M} be a model of T . Consider the structure $(\mathfrak{M} \times S^1, P, Q)$, where S^1 is the circle (with its typical metric) and P and Q are unary $[-1, 1]$ -valued predicates chosen so that for any $(a, b) \in \mathfrak{M} \times S^1$, $P((a, b))$ is the x coordinate of b and $Q((a, b))$ is the y coordinate of b (as a point in \mathbb{R}^2). We will regard P and Q as a single S^1 -valued predicate R . We give $\mathfrak{M} \times S^1$ the max metric (but leave the original metric on \mathfrak{M} as a predicate symbol).

It is not hard to show that this structure is bi-interpretable with \mathfrak{M} , and moreover that this interpretation is uniform in models of T . Let T° be the theory of this structure. It is also not hard to show that if $\varphi(\bar{x})$ is any $[0, 1]$ -valued formula in the original language

of \mathfrak{M} and $f : [0, 1] \rightarrow S^1$ is any continuous function, then $\llbracket R(x_0) = \varphi(\bar{x}) \rrbracket$ is a definable subset of $S_{\bar{x}}(T)$.

Let A be a set of parameters, and let $p, q \in S_n(A)$ be a pair of distinct types. If for some $i < n$, p and q entail distinct values r and s , respectively, of $R(x_i)$, then $\llbracket R(x_i) = r \rrbracket$ and $\llbracket R(x_i) = s \rrbracket$ are disjoint definable sets separating p and q , so assume that for each $i < n$, p and q entail the same value for $R(x_i)$.

Find a $[0, 1]$ -valued A -formula $\varphi(\bar{x})$ such that $\varphi(p) = 0$ and $\varphi(q) = 1$. Let r be the value of $R(x_0)$ entailed by both p and q . We have that $\llbracket R(x_0) = r + \varphi(\bar{x}) \rrbracket$ and $\llbracket R(x_0) = r + \varphi(\bar{x}) - 1 \rrbracket$ are disjoint definable sets containing p and q , respectively, where we understand addition to be in the sense of radians parameterizing S^1 and we understand $\varphi(\bar{x})$ as a formula on $\mathfrak{M} \times S^1$ in the obvious way. \square

2.5 Small Type Spaces and ω -stability

Here, first of all, we will complete some discussion of approximate ω -saturation from Section 1.8. Recall that a pre-structure \mathfrak{M} is κ -saturated over a dense sub-pre-structure if there is a dense sub-pre-structure $\mathfrak{M}_0 \subseteq \mathfrak{M}$ such that for any set $A \subseteq M_0$ with $|A| < \kappa$, for every type $p(x, A) \in S_x(A)$, there is $b \in M$ such that $\mathfrak{M} \models p(b, A)$. Also recall that a pre-structure \mathfrak{M} is *approximately ω -saturated* if for any finite $\bar{a} \in \mathfrak{M}$, any type $p(x, \bar{a}) \in S_x(\bar{a})$, and any $\varepsilon > 0$, there is $\bar{b}, c \in \mathfrak{M}$ such that $d(\bar{a}, \bar{b}) < \varepsilon$, $\bar{a} \equiv \bar{b}$, and $\mathfrak{M} \models p(c, \bar{b})$.

Proposition 2.5.1. *Let \mathfrak{M} be a structure.*

- (i) *If \mathfrak{M} is ω -saturated over a dense sub-pre-structure, then it is approximately ω -saturated.*

(ii) If \mathfrak{M} is approximately ω -saturated, then for any finite or countable $\bar{a} \in \mathfrak{M}$, any $n \leq \omega$, any $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$, any $\varepsilon > 0$, there is a $\bar{b} \in \mathfrak{M}$ with $d(\bar{a}, \bar{b}) < \varepsilon$ and $\bar{a} \equiv \bar{b}$ such that $p(\bar{x}, \bar{b})$ is realized in \mathfrak{M} .

Proof. (i) Let $\mathfrak{M}_0 \subseteq \mathfrak{M}$ be a dense sub-pre-structure such that for any finite $\bar{a} \in \mathfrak{M}_0$, every type in $S_1(\bar{a})$ is realized in \mathfrak{M} .

Claim. For any $\bar{a} \in \mathfrak{M}_0$ and any type $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$ and any $\varepsilon > 0$, there exists $\bar{b} \in \mathfrak{M}_0$ such that $d(tp(\bar{b}\bar{a}), p) < \varepsilon$.

Proof of claim. For each $i \leq n$, let $p_i(x_0, \dots, x_{i-1}, \bar{a})$ be the restriction of p to the variables x_0, \dots, x_{i-1} . First, find $c_0 \in \mathfrak{M}$ such that $\mathfrak{M} \models p_1(c_0, \bar{a})$. Find $b_0 \in \mathfrak{M}_0$ such that $d(c_0, b_0) < \varepsilon$. Note that $d(tp(b_0, \bar{a}), p_1) < \varepsilon$. Now, for each positive $i < n$, given $b_0, \dots, b_{i-1} \in \mathfrak{M}_0$, satisfying $d(tp(b_0, \dots, b_{i-1}, \bar{a}), p_i) < \varepsilon$, find an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ such that for some $\bar{e} \in \mathfrak{N}$, $\mathfrak{N} \models p(\bar{e}, \bar{a})$ and $d^{\mathfrak{N}}(b_0, \dots, b_{i-1}; e_0, \dots, e_{i-1}) < \varepsilon$. Let $q_i = tp(b_0, \dots, b_{i-1}, e_i, \bar{a})$, and note that $d(p_{i+1}, q_i) < \varepsilon$. By approximate ω -saturation, there is $c_i \in \mathfrak{M}$ such that $\mathfrak{M} \models q_i(b_0, \dots, b_{i-1}, c_i, \bar{a})$. Find $b_i \in \mathfrak{M}_0$ close enough to c_i such that $d(tp(b_0, \dots, b_i, \bar{a}), p_{i+1}) < \varepsilon$.

By induction, $\bar{b} = b_0, \dots, b_{n-1}$ has the required properties. □_{claim}

Fix $\bar{a} \in \mathfrak{M}$, and let $p(x, \bar{a}) \in S_\omega(\bar{a})$ be an 1-type over \bar{a} . Fix $\varepsilon > 0$. Let $\bar{c}_0 \in \mathfrak{M}_0$ satisfy $d(\bar{a}, \bar{c}_0) < 2^{-2}\varepsilon$ (and so also $< 2^{-1}\varepsilon$). Let $\mathfrak{N}^0 \succeq \mathfrak{M}$ be an elementary extension such that for some $f_0 \in \mathfrak{N}^0$, $\mathfrak{N}^0 \models p(f, \bar{c})$. Let $q_0 = tp(f_0, \bar{c}_0)$. By approximate ω -saturation, there is $e_0 \in \mathfrak{M}$ such that $\mathfrak{M} \models q_0(e_0, \bar{c}_0)$. Note that we have $d(q_0, p) < 2^{-2}\varepsilon$.

For any $i < \omega$, given $e_i, \bar{c}_i \in \mathfrak{M}_0$ and $q_i = tp(e_i, \bar{c}_i)$ such that $d(q_i, p) < 2^{-i-2}\varepsilon$, find an elementary extension $\mathfrak{N}^{i+1} \succeq \mathfrak{M}$ such that for some $f_{i+1}, \bar{a}_{i+1} \in \mathfrak{N}^{i+1}$, $\mathfrak{N}^{i+1} \models p(f_{i+1}, \bar{a}_{i+1})$ and $d(e_i, \bar{c}_i; f_{i+1}, \bar{a}_{i+1}) < 2^{-i-2}\varepsilon$. Let $q_{i+1}(y, \bar{x}, z, \bar{w}) = tp(e_i, \bar{c}_i, f_{i+1}, \bar{a}_{i+1})$. By the claim, there exists $e_{i+1}, \bar{c}_{i+1} \in \mathfrak{M}_0$ such that $d(tp(e_i, \bar{c}_i, e_{i+1}, \bar{c}_{i+1}), q_{i+1}) < 2^{-i-3}\varepsilon$.

Note that this implies that

$$d(e_i, \bar{c}_i; e_{i+1}, \bar{c}_{i+1}) < 2^{-i-3}\varepsilon + d(e_i, \bar{c}_i; f_{i+1}, \bar{a}_{i+1}) + 2^{-i-3}\varepsilon < 2^{-i-1}\varepsilon.$$

Therefore, $\{e_i \bar{c}_i\}_{i < \omega}$ forms a Cauchy sequence. Let $e_\omega \bar{c}_\omega \in \mathfrak{M}$ be its limit. We have that $d(p, \text{tp}(e_\omega, \bar{c}_\omega)) < \delta$ for every $\delta > 0$, so $p = \text{tp}(e_\omega, \bar{c}_\omega)$. Finally, we have that $d(\bar{a}, \bar{c}_\omega) < \sum_{i < \omega} 2^{-i-1}\varepsilon = \varepsilon$, so setting $\bar{b} = \bar{c}_\omega$ we get the required conclusion.

(ii) It is sufficient to prove this for $n = \omega$. First assume that \bar{a} is finite. Let $p(\bar{x}, \bar{a})$ be an ω -type over \bar{a} , and fix $\varepsilon > 0$. By approximate ω -saturation, we can find $c_0^1, \bar{a}^1 \in \mathfrak{M}$ such that c_0^1, \bar{a}^1 satisfies the restriction of $p(\bar{x}, \bar{y})$ to x_0, \bar{y} and such that $d(\bar{a}^1, \bar{a}) < 2^{-1}\varepsilon$. For each positive $i < \omega$, given $c_0^i, \dots, c_{i-1}^i, \bar{a}^i$ satisfying $p(\bar{x}, \bar{y})$ restricted to $x_0, \dots, x_{i-1}, \bar{y}$, by approximate ω -saturation, we can find $c_0^{i+1}, \dots, c_i^{i+1}, \bar{a}^{i+1} \in \mathfrak{M}$ such that $c_0^{i+1}, \dots, c_i^{i+1}, \bar{a}^{i+1}$ satisfies the restriction of $p(\bar{x}, \bar{y})$ to x_0, \dots, x_i, \bar{y} and such that $d(c_0^i, \dots, c_{i-1}^i, \bar{a}^i; c_0^{i+1}, \dots, c_i^{i+1}, \bar{a}^{i+1}) < 2^{-i-1}\varepsilon$.

By construction, $\{\bar{a}^i\}_{i < \omega}$ is a Cauchy sequence and for each $j < \omega$, $\{c_i^j\}_{j < i < \omega}$ is a Cauchy sequence. Let \bar{a}^ω and c_i^ω be the limits of these Cauchy sequences. By construction, $\mathfrak{M} \models p(\bar{c}^\omega, \bar{a}^\omega)$ and $d(\bar{a}, \bar{a}^\omega) < \sum_{i < \omega} 2^{-i}\varepsilon = \varepsilon$, so setting $\bar{b} = \bar{a}^\omega$ we get the required conclusion.

Now assume that \bar{a} is countable. Given $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$ and $\delta > 0$, we can find $k < \omega$ such that for any ω -tuple $\bar{c} \in \mathfrak{M}$, $d(\bar{a}; a_0, \dots, a_{k-1}, \bar{c}) < \frac{1}{2}\delta$. Then we can apply the finite case to the type $p(\bar{x}, a_0, \dots, a_{k-1}, \bar{y})$ with $\varepsilon = \frac{1}{2}\delta$, and we get the required conclusion. \square

Remark 2.5.2. Propositions 1.8.3 and 2.5.1 part (i) are evidence that approximate ω -saturation is the ‘correct’ generalization of ω -saturation to continuous logic. This contention is further evidenced by the facts that ω -categoricity and ω -stability in general only guarantee the existence of approximately ω -saturated separable models, as well as the fact that an approximately ω -saturated separable model (of a theory in a countable language) is always unique [BYU07, Fact 1.5].

In my opinion, the strongest evidence is the fact that every continuous theory T is bi-interpretable with a continuous theory S with the property that a model of S is ω -saturated if and only if it is \aleph_1 -saturated. Specifically, S is the theory of structures of the form \mathfrak{M}^ω for $\mathfrak{M} \models T$.

The following proposition is the topometric analog of the fact that in a countable Baire space (i.e. a space in which the Baire category theorem holds, such as a locally compact Hausdorff space or a completely metrizable space) isolated points are dense. This result was originally proven in greater generality in [BY08c], but this specific case can be proven more directly using the Baire category theorem.

Proposition 2.5.3. *Let (X, τ, d) be a topometric space. If (X, τ) is locally compact or completely metrizable and $\#^{\text{dc}}(X, d) \leq \aleph_0$, then d -atomic points are dense in X .*

Proof. Since $\#^{\text{dc}}X \leq \aleph_0$, for any $\varepsilon > 0$ there is a countable set $X_\varepsilon \subseteq X$ such that $X \subseteq \bigcup_{x \in X_\varepsilon} B_{\leq \varepsilon}(x)$. By the Baire category theorem and since each $B_{\leq \varepsilon}(x)$ is closed, for any non-empty open set U there must be $x \in X_\varepsilon$ such that $B_{\leq \varepsilon}(x) \cap U$ has non-empty interior. This implies that for every $\varepsilon > 0$, the set $Y_\varepsilon = \bigcup_{x \in X} \text{int } B_{\leq \varepsilon}(x)$ is open and dense. Therefore by the Baire category theorem, $Z = \bigcap_{k < \omega} Y_{2^{-k}}$ is non-empty and

dense. For any $z \in Z$, for each $\varepsilon > 0$, there is $k < \omega$ such that $2^{-k} < \frac{\varepsilon}{2}$ and there is x such that $z \in \text{int } B_{\leq 2^{-k}}(x) \subseteq B_{\leq \varepsilon}(z)$. Therefore $z \in \text{int } B_{\leq \varepsilon}(z)$ for every $\varepsilon > 0$ and z is d -atomic. \square

For a generalization of Proposition 2.5.3, see Proposition B.2.12.

Definition 2.5.4. A type space $S_V(\Sigma)$ is ρ -small if $\#^{\text{dc}}(S_V(\Sigma), \rho) \leq \aleph_0$. We will just say small if $\rho = d$, the induced metric on type space. A complete theory T is *small* if each $S_n(T)$ for $n < \omega$ is small. \triangleleft

Proposition 2.5.5. *A countable complete theory T is small if and only if it has an approximately ω -saturated separable model.*

Proof. (\Rightarrow) Assume that T is small, and assume without loss that \mathcal{L} is relational. For every type $p(\bar{y}) \in S_n(T)$ and every $j < \omega$ let $\{q_{p,j}^i(\bar{x}, \bar{y})\}_{i < \omega}$, where \bar{x} is a j -tuple, be a metrically tail-dense sequence in $S_{\bar{x}\bar{y}}(p(\bar{y}))$ (i.e. every final segment is metrically dense) with the metric d (i.e. the metric on $S_{\bar{x}\bar{y}}(T)$ restricted to $S_{\bar{x}\bar{y}}(p(\bar{y}))$), not the metric $d_{/\bar{y}}$, note that these are all metrically separable as they are subspaces of metrically separable spaces).

It is not hard to build a countable pre-structure \mathfrak{M}_0 such that for every finite tuple $\bar{a} \in \mathfrak{M}_0$, if $p(\bar{y}) = \text{tp}(\bar{a})$, then every type $q_{p,j}^i(x, \bar{a})$ is realized in \mathfrak{M}_0 for every i and j . Let \mathfrak{M} be the completion of \mathfrak{M}_0 . We need to argue that \mathfrak{M} is approximately ω -saturated.

Let $\bar{a} \in \mathfrak{M}$ be a finite tuple, and let $p(x, \bar{a}) \in S_x(\bar{a})$ be a type over \bar{a} . Fix $\varepsilon > 0$. Let \bar{b}^0 be a tuple of elements of \mathfrak{M}_0 such that $d(\bar{a}, \bar{b}^0) < 2^{-1}\varepsilon$, so in particular $d^{\text{tp}}(\bar{a}, \bar{b}^0) < 2^{-1}\varepsilon$. Let $r_0(\bar{y}) = \text{tp}(\bar{b}^0)$. There is a type $s(x, \bar{y})$ extending $r_0(\bar{y})$ such that $d(p, s_0) < 2^{-1}\varepsilon$. This implies that we can find some $q_{r_0,1}^{i_0}$ such that $d(p, q_{r_0,1}^{i_0}) < 2^{-1}\varepsilon$ (by the triangle inequality). Therefore by construction there is $c^0 \in \mathfrak{M}_0$ such that $\mathfrak{M} \models q_{r_0,1}^{i_0}(c^0, \bar{b}^0)$.

Now for each $k < \omega$, given $\bar{b}^k, c^k \in \mathfrak{M}_0$ such that $d(p, \text{tp}(c^k \bar{b}^k)) < 2^{-k-1}\varepsilon$. Let $u_k(x, \bar{y}) = \text{tp}(c^k \bar{b}^k)$, and let $t_k(x, \bar{y}, z, \bar{w})$ be a type such that

$$\begin{aligned} t_k(x, \bar{y}, z, \bar{w}) &\models u_k(x, \bar{y}) \\ &\models p(z, \bar{w}) \\ &\models d(x\bar{y}, z\bar{w}) \leq \delta_k < 2^{-k-1}\varepsilon \end{aligned}$$

for some $\delta_i > 0$. Now by construction we can find a type $q_{u_k, |z\bar{w}|}^{\ell_k}(x, \bar{y}, z, \bar{w})$ such that

$$d\left(t_k, q_{u_k, |z\bar{w}|}^{\ell_k}\right) < \min\left\{\frac{2^{-k-1}\varepsilon - \delta_k}{2}, 2^{-k-2}\right\},$$

and by construction there are $c^{k+1}, \bar{b}^{k+1} \in \mathfrak{M}_0$ such that $\mathfrak{M} \models q_{u_k, |z\bar{w}|}^{\ell_k}(c^k, \bar{b}^k, c^{k+1}, \bar{b}^{k+1})$.

This implies that

$$d(c^k \bar{b}^k, c^{k+1} \bar{b}^{k+1}) < \delta_k + 2 \min\left\{\frac{2^{-k-1}\varepsilon - \delta_k}{2}, 2^{-k-2}\right\} \leq \delta_k + 2^{-k-1}\varepsilon - \delta_k = 2^{-k-1}\varepsilon,$$

and

$$d(p, \text{tp}(c^{k+1} \bar{b}^{k+1})) < \min\left\{\frac{2^{-k-1}\varepsilon - \delta_k}{2}, 2^{-k-2}\right\} \leq 2^{-k-2}.$$

So we have that $\{c^k \bar{b}^k\}_{k < \omega}$ is a Cauchy sequence. Let its limit be $c^\omega \bar{b}^\omega$, and note that

$$d(\bar{a}, \bar{b}^\omega) < \sum_{k < \omega} 2^{-k-1}\varepsilon = \varepsilon.$$

We have by construction that $d(p, \text{tp}(c^\omega \bar{b}^\omega)) < \gamma$ for every $\gamma > 0$, so $\mathfrak{M} \models p(c^\omega, \bar{b}^\omega)$, which in particular implies that $\bar{a} \equiv \bar{b}^\omega$, so we have that \bar{b}^ω is the required tuple for

which $p(x, \bar{b}^\omega)$ is realized in \mathfrak{M} .

Since we can do this for any $\bar{a} \in \mathfrak{M}$, $p(x, \bar{a}) \in S_x(\bar{a})$, and $\varepsilon > 0$, we have that \mathfrak{M} is approximately ω -saturated.

(\Leftarrow) Assume that T has an approximately ω -saturated separable model. This implies that it realizes every type in $S_n(T)$ for every $n < \omega$. Since the map from tuples to their type is metrically 1-Lipschitz, this implies that each $S_n(T)$ is metrically separable, or in other words small. \square

The following Lemma 2.5.6 is equivalent to Proposition 12.8 in [BYBHU08], and Proposition 2.5.7 and Corollary 2.5.8 are direct¹⁰ consequences of Corollary 12.9 in [BYBHU08], but we have included proofs and in particular a more detailed proof of Proposition 2.5.7 for completeness. These facts and their proofs are also very similar to Fact 1.5 in [BYU07] (the uniqueness of approximately ω -saturated separable models).

Lemma 2.5.6. *If $\text{tp}(\bar{a}\bar{b})$ is atomic and \mathfrak{M} is any model containing \bar{b} , then for every $\varepsilon > 0$ there exists $a'\bar{b}' \in \mathfrak{M}$ such that $\bar{a}\bar{b} \equiv a'\bar{b}'$ and $d(\bar{b}, \bar{b}') < \varepsilon$.*

Proof. Let $D(x, \bar{y})$ be a distance predicate for the definable set $\text{tp}(\bar{a}\bar{b})$. We have that $\mathfrak{M} \models \inf_x D(x, \bar{b})$, since this is part of the type of \bar{b} . Let a_0 be such that $D(a_0, \bar{b}) < \varepsilon$. Since D is a distance predicate this implies that there is $a'\bar{b}' \in D(\mathfrak{M})$ with $d(a_0\bar{b}, a'\bar{b}') < \varepsilon$. $a'\bar{b}'$ is the required sequence of elements. \square

Proposition 2.5.7. *If a countable theory T has a prime model, it has a unique prime model. Moreover if $\bar{a} \equiv \bar{b}$ with $\bar{a} \in \mathfrak{M}$ and $\bar{b} \in \mathfrak{N}$, with \mathfrak{M} and \mathfrak{N} prime, then for every $\varepsilon > 0$ there is an isomorphism $f : \mathfrak{M} \rightarrow \mathfrak{N}$ such that $d(f(\bar{a}), \bar{b}) < \varepsilon$.*

¹⁰[BYBHU08] does not state the conclusion that in countable theories a model is prime if and only if it is separable and atomic, but this is immediate from what is stated there.

Proof. Let \mathfrak{M} and \mathfrak{N} be prime. Let $\bar{a} \in \mathfrak{M}$ and $\bar{b} \in \mathfrak{N}$ with $\bar{a} \equiv \bar{b}$. Fix $\varepsilon > 0$. By the omitting types theorem, every type realized in \mathfrak{M} or \mathfrak{N} is atomic. Let $\{a_i\}_{i < \omega}$ be an enumeration of a tail dense sequence in \mathfrak{M} (i.e. every final segment is dense), starting with \bar{a} , and let $\{b_i\}_{i < \omega}$ be an enumeration of a tail dense sequence in \mathfrak{N} , starting with \bar{b} . We are going to build an array $\{c_j^i\}_{i < \omega, j \leq i} \subseteq \mathfrak{M}$ and an array $\{e_j^i\}_{i < \omega, j \leq i} \subseteq \mathfrak{N}$ with the following properties:

- For every $i < \omega$ and $j \leq i$, $d(c_j^i, c_j^{i+1}) < 2^{-i-2}\varepsilon$ and $d(e_j^i, e_j^{i+1}) < 2^{-i-2}\varepsilon$.
- At each stage i , $c_0^i c_1^i \dots c_i^i \equiv e_0^i e_1^i \dots e_i^i$.
- For every $k < \omega$, $c_{2k}^{2k} = a_k$ and $e_{2k+1}^{2k+1} = b_k$.

For the initial step, let $c_0^0 = a_0$. Since $\text{tp}(a_0)$ is atomic, we can find $e_0^0 \in \mathfrak{M}$ such that $c_0^0 \equiv e_0^0$.

On odd step $2k + 1$, given $c_{\leq 2k}^{2k}$ and $e_{\leq 2k}^{2k}$ satisfying $c_{\leq 2k}^{2k} \equiv e_{\leq 2k}^{2k}$, let $e_j^{2k+1} = e_j^{2k}$ for $j \leq 2k$, and let $e_{2k+1}^{2k+1} = b_k$. In particular note that $d(e_j^{2k}, e_j^{2k+1}) = 0 < 2^{-2k-2}\varepsilon$ for each $j \leq 2k$. Now by the lemma we can find $c_{\leq 2k+1}^{2k+1} \in \mathfrak{M}$ such that $d(c_{\leq 2k}^{2k}, c_{\leq 2k}^{2k+1}) < 2^{-2k-3}\varepsilon$ and such that $c_{\leq 2k+1}^{2k+1} \equiv e_{\leq 2k+1}^{2k+1}$.

On even step $2k+2$, given $c_{\leq 2k+1}^{2k+1}$ and $e_{\leq 2k+1}^{2k+1}$ satisfying $c_{\leq 2k+1}^{2k+1} \equiv e_{\leq 2k+1}^{2k+1}$, let $c_j^{2k+2} = c_j^{2k+1}$ for $j \leq 2k + 1$, and let $c_{2k+2}^{2k+2} = a_k$. In particular note that $d(c_j^{2k+1}, c_j^{2k+2}) = 0 < 2^{-2k-3}\varepsilon$ for each $j \leq 2k + 1$. Now by the lemma we can find $e_{\leq 2k+2}^{2k+2} \in \mathfrak{N}$ such that $d(e_{\leq 2k+1}^{2k+1}, e_{\leq 2k+1}^{2k+2}) < 2^{-2k-4}\varepsilon$ and such that $e_{\leq 2k+2}^{2k+2} \equiv c_{\leq 2k+2}^{2k+2}$.

For each $j < \omega$, let $c_j = \lim_{i \rightarrow \omega} c_j^i$ and $e_j = \lim_{i \rightarrow \omega} e_j^i$. By construction these are all Cauchy sequences and thus have limits in \mathfrak{M} and \mathfrak{N} respectively. Also by uniform continuity of formulas we have that $c_0 c_1 \dots \equiv e_0 e_1 \dots$.

We need to show that the sequences $c_{<\omega}$ and $e_{<\omega}$ are metrically dense in \mathfrak{M} and \mathfrak{N} respectively. This follows by the fact that the sequences $a_{<\omega}$ and $b_{<\omega}$ were tail dense. For any $x \in \mathfrak{M}$ and any $\delta > 0$ there is an a_i such that $d(x, a_i) < \frac{1}{2}\delta$ and $d(a_i, c_j) < \frac{1}{2}\delta$ for some c_j , implying that $d(x, c_j) < \delta$. Therefore $c_{<\omega}$ is dense in \mathfrak{M} . Likewise for $e_{<\omega}$ in \mathfrak{N} , so we have that the elementary map $f_0 : c_{<\omega} \rightarrow e_{<\omega}$ extends to an isomorphism $f : \mathfrak{M} \rightarrow \mathfrak{N}$.

Now note that if \bar{c} is the initial segment of $c_{<\omega}$ of the same length as \bar{a} , then we have by construction that $d(\bar{c}, \bar{a}) < \sum_{i < \omega} 2^{-i-2}\varepsilon = \frac{1}{2}\varepsilon$. Likewise if \bar{e} is the initial segment of $e_{<\omega}$ of the same length as \bar{b} , then by construction we have that $d(\bar{e}, \bar{b}) < \frac{1}{2}\varepsilon$. Therefore we have that $d(f(\bar{a}), \bar{b}) \leq d(f(\bar{a}), f(\bar{c})) + d(\bar{e}, \bar{b}) < \frac{1}{2}\varepsilon + \frac{1}{2}\varepsilon$, as required. \square

We will finish the characterization of countable theories with prime models in Proposition B.1.5.

Corollary 2.5.8. *If \mathfrak{M} is a prime model of a countable theory, then it is approximately \aleph_0 -homogeneous (i.e. for any finite tuples $\bar{a} \equiv \bar{b}$ and for any $\varepsilon > 0$ there is an automorphism σ of \mathfrak{M} such that $d(\sigma(\bar{a}), \bar{b}) < \varepsilon$).*

For the following lemma, note that there is no requirement that f be a continuous function. See Proposition 4.5.9 for an example of using this lemma with a discontinuous function.

Lemma 2.5.9. *If (X, d) is a separable metric space and $f : X \rightarrow \mathbb{R}$ is a (not necessarily continuous) function, then for all but countably many $r \in \mathbb{R}$,*

$$\{f = r\} \subseteq \overline{\{f < r\}}.$$

Proof. Assume that for an uncountable set $R \subseteq \mathbb{R}$ and for each $r \in R$, $\{f \leq r\} \not\subseteq \overline{\{f < r\}}$. By definition this means that for each $r \in R$ there is $x_r \in X$ with $f(x_r) = r$ such that $d_{\text{inf}}(x_r, \{f < r\}) > \varepsilon_r$. Since there are uncountably many such r , by the cofinality of \mathbb{R} , there is an $R_0 \subseteq R$ and an $\varepsilon > 0$ such that for all $r \in R_0$, $\varepsilon_r > \varepsilon$. This implies that for any distinct $r, s \in R_0$ with $r > s$, $d(x_r, x_s) \geq d_{\text{inf}}(x_r, \{f < r\}) > \varepsilon_r > \varepsilon$, so R_0 is an uncountable ($> \varepsilon$)-separated set, contradicting that (X, d) is separable. \square

Proposition 2.5.10. *If $S_n(T)$ is small, then for any real formula φ , for all but countably many r , $(\varphi \leq r)$ is a definable set.*

Proof. By Lemma 2.5.9, for all but countably many r , $\llbracket \varphi \leq r \rrbracket \subseteq \overline{\llbracket \varphi < r \rrbracket}$. $\llbracket \varphi < r \rrbracket$ is open and therefore locatable, so $\llbracket \varphi \leq r \rrbracket$ is locatable and therefore definable. \square

Note, however, that we cannot guarantee that $(\varphi = r)$ will be definable for any $r \in I(\varphi)$ (see Counterexample C.1.7).

Definition 2.5.11. A complete theory T is ω -stable if $\#^{\text{dc}} S_1(A) \leq \aleph_0$ for every countable set of parameters A . \triangleleft

Corollary 2.5.12. *If a countable theory T has the property that for every $n < \omega$ and every \bar{a} , $S_n(\bar{a})$ is small (in particular if T is ω -stable, T has an exactly ω -saturated separable model, or $T_{\bar{a}}$ is ω -categorical for every \bar{a}), then T is dictionaric.*

Proof. If F, G are inconsistent closed A -formulas, then there is a restricted formula $\varphi(\bar{x}, \bar{a})$ such that $F(\bar{x}) \models \varphi(\bar{x}, \bar{a}) = 0$ and $G(\bar{x}) \models \varphi(\bar{x}, \bar{a}) = 1$. For some $r, s \in (0, 1)$, with $r < s$, $(\varphi(\bar{x}, \bar{a}) \leq r)$ and $(\varphi(\bar{x}, \bar{a}) \geq s)$ are definable. Since we can do this for any pair of inconsistent closed formulas, T is dictionaric. \square

See Definition B.2.1 for the definition of *totally transcendental*.

Corollary 2.5.13. *If T is totally transcendental, then for any real formula φ , for all but countably many r , $(\varphi \leq r)$ is a definable set.*

Proof. Pass to a countable reduct $T \upharpoonright \mathcal{L}_0$ such that φ is an \mathcal{L}_0 -formula. □

Corollary 2.5.14 (Strong Intersection Property). *If $S_n(T)$ is a small type space and $\{D_i\}_{i < \kappa}$ is a family of definable sets with $\kappa < 2^{\aleph_0}$, then for all but $\aleph_0 + \kappa$ many r , $(\varphi \leq r)$, $(\varphi \geq r)$, $(\varphi \leq r) \cap D_i$, and $(\varphi \geq r) \cap D_i$ are definable for every $i < \kappa$.*

Chapter 3

Imaginarities and Interpretations

3.1 Many-Sorted Signatures

Just as in discrete logic, the passage from single-sorted to many-sorted signatures is straightforward but tedious. Unsurprisingly, in continuous logic it is even more fraught with tedium than it is in discrete logic. We will not bother proving any many-sorted analogs of many of the earlier results in this thesis, as the generalizations should be obvious, but we do need to set conventions and notation.

Definition 3.1.1. A *many-sorted signature* \mathcal{L} contains the following data:

- A non-empty set of *sorts*, $\mathcal{S}(\mathcal{L})$. For each sort $O \in \mathcal{S}(\mathcal{L})$, there is a *diameter bound* $\text{db}(O)$, a positive real number.
- A set of *predicate symbols*, $\mathcal{P}(\mathcal{L})$. For each $P \in \mathcal{P}$, we have a compact interval $I(P) \subseteq \mathbb{R}$, an arity $a(P) \in \mathcal{S}(\mathcal{L})^{\leq \omega}$, and if $a(P) \neq \emptyset$, a modulus ω_P . For each sort $O \in \mathcal{S}(\mathcal{L})$, there is a special designated predicate symbol d_O with $I(d_O) = [0, \text{db}(O)]$, $a(d_O) = OO$, and $\omega_{d_O}(r) = 2r$.
- A set of *function symbols* $\mathcal{F}(\mathcal{L})$. For each $f \in \mathcal{F}$, we have a *codomain sort* $S(f) \in \mathcal{S}(\mathcal{L})$, an arity $a(f) \in \mathcal{S}(\mathcal{L})^{\leq \omega}$, and if $a(f) \neq \emptyset$, a modulus ω_f . Function symbols f

with $a(f) = \emptyset$ are called *constant symbols*. The set of constant symbols is written $\mathcal{C}(\mathcal{L})$.

For any many-sorted signature \mathcal{L} , the *cardinality of \mathcal{L}* , written $|\mathcal{L}|$, is $|\mathcal{S}(\mathcal{L})| + |\mathcal{P}(\mathcal{L})| + |\mathcal{F}(\mathcal{L})|$.

A many sorted signature \mathcal{L}' is an *expansion of \mathcal{L}* , written $\mathcal{L}' \supseteq \mathcal{L}$ or $\mathcal{L} \subseteq \mathcal{L}'$, if $\mathcal{S}(\mathcal{L}') \supseteq \mathcal{S}(\mathcal{L})$, $\mathcal{P}(\mathcal{L}') \supseteq \mathcal{P}(\mathcal{L})$, and $\mathcal{F}(\mathcal{L}') \supseteq \mathcal{F}(\mathcal{L})$, with $a(P)$, $a(f)$, ω_P , ω_f , $I(P)$, and $S(f)$ the same in \mathcal{L} and \mathcal{L}' for all $P, f \in \mathcal{L}$. We also say that \mathcal{L} is a *reduct of \mathcal{L}'* . \triangleleft

Note in particular that we allow expansions to include additional sorts.

Notation 3.1.2. For any many-sorted signature \mathcal{L} and any non-empty $\mathcal{S}_0 \subseteq \mathcal{S}$, we write $\mathcal{L} \upharpoonright \mathcal{S}_0$ for the reduct of \mathcal{L} containing only the sorts \mathcal{S}_0 and any predicate or function symbols whose arity sorts and, in the case of function symbols, codomain sorts are contained in \mathcal{S}_0 . If Σ is an \mathcal{L} -type (in particular, an \mathcal{L} -theory), then we write $\Sigma \upharpoonright \mathcal{S}_0$ for the reduct of Σ to $\mathcal{L} \upharpoonright \mathcal{S}_0$. \triangleleft

For each sort O in some signature \mathcal{L} , we allow arbitrary collections of *variable symbols of sort O* . There is also a special set of variable symbols of sort O , the *variable symbols for binding of sort O* , written $\mathcal{V}_b(O) = \{\dot{v}_i^O\}_{i < \omega_1}$, indexed by ω_1 . The important thing is that if two signatures have a sort O in common, then variable symbols of sort O are the same objects for both sorts. If we need to emphasize the sort of a variable symbol v , we will write expressions such as $v:O$. This will typically be in the context of quantifiers, such as in expressions like $(\forall v:O)F(v)$. We extend the notation for codomain sorts from functions to terms in general. For any variable symbol v , $S(v)$ is the sort of v and for any term $f\bar{t}$, $S(f\bar{t}) = S(f)$.

The definitions of $\mathcal{L}(V)$ -terms and $\mathcal{L}(V)$ -formulas and most other concepts for many-sorted \mathcal{L} are obvious. Although note that in quantification with infinite sets or tuples of variables, we allow arbitrary collections of variables, possibly of infinitely many distinct sorts. There are only two tricky things: the metric on ω -products and the notion of a restricted terms and formulas.

Definition 3.1.3. If \mathcal{L} is a many-sorted signature and \bar{O} is an ω -tuple of sorts in \mathcal{L} , then we take the metric on tuples \bar{a}, \bar{b} of the same sorts as \bar{O} to be

$$d(\bar{a}, \bar{b}) = \sup_{i < \omega} \frac{\text{db}(O_0)}{2^i \text{db}(O_i)} d_{O_i}(a_i, b_i).$$

We will use this freely in formulas. ◁

The factor of $\text{db}(O_0)$ is only there so that this will agree with the single-sorted definition of the metric on ω -products. We require uniform continuity with regards to this metric in the definition of an \mathcal{L} -structure for many-sorted \mathcal{L} . The metric on finite product sorts is still the max metric.

The notions of restricted term and formula do not generalized in a direct way to many-sorted signatures, so we need to do something new. Generally the problem is that we may have arities involving infinitely many distinct sorts. The solution to this is to require for any given term $f\bar{t}$ that for all but finitely many i and all but finitely many sorts O , if $t_i = O$, then t_i is some particular fixed dummy variable, and then to quantify over these dummy variables in front of each atomic formula.

Remark 3.1.4. This shaded section deals with a precise definition of a notion of *restricted formula* for many-sorted continuous logic. The combination of having many sorts and allowing ω -ary terms and predicates creates the need for some very baroque bookkeeping in defining such a concept. This section can be safely skipped if one buys the following statement: There is a systematic way of constructing a class of restricted \mathcal{L} -formula, given any many-sorted signature \mathcal{L} and any set of variable symbols V , $R\mathcal{L}$, such that

- there are precisely $\aleph_0 + |\mathcal{L}| + |V|$ restricted formulas,
- $R\mathcal{L}$ is dense in \mathcal{L} under $\|\cdot\|_{\equiv}$,
- each restricted formula involves only finitely many symbols in \mathcal{L} and has finitely many free variables, and
- if \mathcal{L} is single-sorted, then $R\mathcal{L}(V)$ agrees with Definition 1.3.2.

Definition 3.1.5. If X is a predicate or function symbol, \bar{t} is a tuple of terms of the same sorts as $a(X)$, and $k < \omega$, then we say that \bar{t} is *restricted for X with index k* if either $a(X)$ is finite or $a(X)$ is infinite and there exists $N < \omega$ such that for all $i, j > N$,

- if $S(t_i) = S(t_j)$, then $t_i = t_j$, and
- for all but finitely many sorts $O \in \mathcal{S}(\mathcal{L})$, if $S(t_i) = O$, then $t_i = \dot{v}_k^O$.

We say that \bar{t} is *restricted for X* if it is restricted for X with index k for some $k < \omega$. \triangleleft

The only reason why we need dummy variables of different indices is that we might be using some variable symbols for binding as free variables. If we were content to never use \dot{v}_0^O as a free variable, we could require that $k = 0$.

Definition 3.1.6. Let \mathcal{L} be a many-sorted signature. The class of *restricted $\mathcal{L}(V)$ -terms with index k* is defined inductively:

- Each variable symbol $v \in V$ is a restricted $\mathcal{L}(V)$ -term.
- For each function symbol f and tuple \bar{t} of restricted $\mathcal{L}(V)$ -terms with index k of the same sorts as $a(f)$, if \bar{t} is restricted for f with index k , then $f\bar{t}$ is a restricted $\mathcal{L}(V)$ -term with index k .

A *restricted $\mathcal{L}(V)$ -term* is a restricted $\mathcal{L}(V)$ -term with index k for some $k < \omega$. \triangleleft

Definition 3.1.7. Let \mathcal{L} be a many-sorted signature. For any $k < \omega$ and any set of sorts \mathcal{S} , let $\mathcal{V}_b(\mathcal{S}, k)$ be the set of variables of the form v_k^O with $O \in \mathcal{S}$. The class of *restricted $\mathcal{L}(V)$ -formula* are defined as follows:

(i) If P is a predicate symbol such that $a(P)$ contains only finitely many sorts, then for any tuple of restricted $\mathcal{L}(V)$ -terms \bar{t} which is restricted for P , then $P\bar{t}$ is a restricted $\mathcal{L}(V)$ -formula.

(ii) If

- P is a predicate symbol such that $a(P)$ contains infinitely many sorts,
- k is an ordinal less than ω ,
- \bar{t} is a tuple of restricted $\mathcal{L}(V \cup \mathcal{V}_b(\mathcal{S}, k))$ -terms which is restricted for P with index k , and
- \mathcal{S}_0 is the set of all sorts of variable symbols in \bar{t} (which, note, is always at most countable),

then $\inf_{\mathcal{V}_b(\mathcal{S}_0, k)} P\bar{t}$ and $\sup_{\mathcal{V}_b(\mathcal{S}_0, k)} P\bar{t}$ are restricted $\mathcal{L}(V)$ -formula.

- (iii) 1 is a restricted $\mathcal{L}(V)$ -formula.
- (iv) If φ is a restricted $\mathcal{L}(V)$ -formula, then for any $r \in \mathbb{Q}$, $r \cdot \varphi$ is a restricted $\mathcal{L}(V)$ -formula.
- (v) If φ and ψ are restricted $\mathcal{L}(V)$ -formulas, then $\varphi + \psi$, $\max\{\varphi, \psi\}$, and $\min\{\varphi, \psi\}$ are restricted $\mathcal{L}(V)$ -formulas.
- (vi) If $i < \omega$, $O \in \mathcal{S}(\mathcal{L})$, and φ is a restricted $\mathcal{L}(V\dot{v}_i^O)$ -formula, then $\sup_{v_i^O} \varphi$ and $\inf_{v_i^O} \varphi$ are restricted $\mathcal{L}(V)$ -formulas. ◁

We don't actually need both sup and inf in part (ii) of Definition 3.1.7. Also, parts (i) and (ii) are only separated to ensure that these definitions are literally generalizations of the definitions of restricted terms and formulas for single-sorted signatures given in Section 1.3. The issue being that technically $\sup_{xy} \varphi$ and $\sup_x \sup_y \varphi$ are different formulas.

Proposition 3.1.8. *Let \mathcal{L} be a many-sorted signature and V a set of variable symbols with sorts in \mathcal{L} .*

- (i) *The collection of restricted $\mathcal{L}(V)$ -terms has cardinality at most $\aleph_0 + |\mathcal{L}| + |V|$.*
- (ii) *The collection of restricted $\mathcal{L}(V)$ -formulas has cardinality $\aleph_0 + |\mathcal{L}| + |V|$.*
- (iii) *Every restricted formula contains at most finitely many predicate and function symbols and has at most finitely many free variables.*

- (iv) If for every $P \in \mathcal{P}(\mathcal{L})$, $a(P)$ contains finitely many sorts, then the set of restricted $\mathcal{L}(V)$ -terms is dense in the set of $\mathcal{L}(V)$ -terms under d_{\equiv} .
- (v) For any $k < \omega$, every $\mathcal{L}(V)$ -term is a limit of $\mathcal{L}(V \cup \mathcal{V}_b(\mathcal{S}, k))$ -terms under d_{\equiv} .
- (vi) The set of restricted $\mathcal{L}(V)$ -formulas is dense in the set of $\mathcal{L}(V)$ -formulas under $\|\cdot\|_{\equiv}$.
- (vii) Every restricted real formula is logically equivalent to one of the form

$$Q_0 V_0 Q_1 V_1 \dots Q_{\ell-1} V_{\ell-1} \max_{n < N} \min_{m < M_n} a_{nm} + \sum_{k < K_{nm}} b_{nmk} \cdot P_{nmk} \bar{t}_{nmk}$$

with $\ell, N, M_n, K_{nm} < \omega$, with $a_{nm}, b_{nmk} \in \mathbb{Q}$, with \bar{t}_{nmk} a tuple of restricted terms restricted for P_{nmk} , with each Q_i a quantifier, and with each V_i either a singleton or a set of the form $\mathcal{V}_b(\mathcal{S}_0, i)$ for some $\mathcal{S}_0 \subseteq \mathcal{S}(\mathcal{L})$ and $i < \omega$. And likewise with \max and \min switched.

The proofs of the parts of Proposition 3.1.8 are either straightforward or very similar to corresponding proofs in Section 1.3. The only vaguely notable differences are parts (iv) and (v).

3.2 Imaginary Sorts and T^{eq}

In discrete logic, we typically take imaginaries to be formed by two operations: passing to finite product sorts and passing to definable quotients. There is another operation that would be admissible, but is not really needed, which is passing to definable subsets.

This is because whenever $\varphi(x)$ is a formula (and there is a definable element $a \in \mathfrak{M}$ such that $\mathfrak{M} \models \varphi(a)$), the set $\varphi(\mathfrak{M})$ can canonically be identified with a quotient by a certain definable equivalence relation, namely

$$xEy = (x = y) \vee (\neg\varphi(x) \wedge y = a) \vee (x = a \wedge \neg\varphi(y)) \vee (\neg\varphi(x) \wedge \neg\varphi(y)).$$

In continuous logic, we can already guess that we will need to alter this picture. We will probably need to consider ω -product sorts, in addition to finitary product sorts, and we will need to consider definable pseudo-metrics, rather than definable equivalence relations (although see Lemma 3.4.4), but more than this we need to allow the operation of passing to definable subsets, because of examples like this:

Counterexample 3.2.1. *A structure with a definable set that is not metrically homeomorphic to any quotient of a Cartesian power of it.*

Description. Let \mathcal{L} be the signature with a single 1-Lipschitz unary predicate P with $I(P) = [0, 1]$ and with $\text{db}(\mathcal{L}) = 1$. Let \mathfrak{J} be the \mathcal{L} -structure with $[0, 1]$ as its universe, with $P^{\mathfrak{J}}(x) = x$, and with $d^{\mathfrak{J}}(x, y) = |x - y|$.

Let $D(x) = (P(x) = 0) \vee (P(x) = 1)$. This is a definable set which is topologically disconnected, but any continuous quotient of any Cartesian power of \mathfrak{J} will be connected, so $D(\mathfrak{J})$ is not metrically homeomorphic to any quotient of \mathfrak{J}^n for any $n \leq \omega$. \square

This is clearly a wholly continuous obstacle.

Remark 3.2.2. That said, a similar situation to Counterexample 3.2.1 can occur in discrete logic in the absence of parameters. Consider the discrete theory of a single infinite

and co-infinite predicate, P . The set of realizations of P feels like it ought to be an imaginary of this theory, but no \emptyset -definable quotient of the sort of n -tuples is in definable bijection with P , for any $n < \omega$. The analog of Theorem 3.2.18 in discrete logic states that any harmless expansion is equivalent to a definable subset of a quotient of some power of the home sort, which lends further credence to the idea that passing to definable sets should be allowed as an imaginary forming operation. Admittedly, however, this issue is rarely important in discrete logic.

Definition 3.2.3. Given a (possibly many-sorted) signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T with a finite sequence of sorts $\{O_i\}_{i < n}$, a (*finitary*) *product sort imaginary expansion of T by $\{O_i\}_{i < n}$* , also called a (*finitary*) *product sort*, is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ containing a single new sort P with diameter bound $\max_{i < n} \{\text{db}(O_i)\}$ and in which the only new predicate symbol is d_P and the only new function symbols are the 1-Lipschitz function symbols $\pi_i : P \rightarrow O_i$ and $\langle \cdot, \dots, \cdot \rangle : O_0 \times \dots \times O_{n-1} \rightarrow P$. Finally, T' is T with the following new axioms added:

- $(\forall x_0 : O_0) \dots (\forall x_{n-1} : O_{n-1}) \pi_i(\langle x_0, \dots, x_{n-1} \rangle) = x_i$ for each $i < n$ and
- $(\forall xy : P) d_P(x, y) = \max_{i < n} d_{O_i}(\pi_i(x), \pi_i(y))$. \triangleleft

Definition 3.2.4. Given a (possibly many-sorted) signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T with a countable sequence of sorts $\{O_i\}_{i < \omega}$, an (*infinitary*) *product sort imaginary expansion of T by $\{O_i\}_{i < \omega}$* , also called an (*infinitary*) *product sort*, is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ containing a single new sort P with diameter bound $\text{db}(O_0)$ and in which the only new predicate symbol is d_P and

the only new function symbols are the function symbols $\pi_i : P \rightarrow O_i$, which are each $2^i \frac{\text{db}(O_i)}{\text{db}(O_0)}$ -Lipschitz, and $\langle \cdot, \dots \rangle : O_0 \times O_1 \times \dots \rightarrow P$, which is 1-Lipschitz. T' is T with the following new axioms added:

- $(\forall x_0:O_0)(\forall x_1:O_1) \dots \pi_i(\langle x_0, x_1, \dots \rangle) = x_i$ for all $i < \omega$ and
- $(\forall xy:P)d_P(x, y) = \sup_{i < \omega} \frac{\text{db}(O_0)}{2^i \text{db}(O_i)} d_{O_i}(\pi_i(x), \pi_i(y))$. \triangleleft

Definition 3.2.5. A *product sort* is either an infinitary or a finitary product sort. A product sort of the sequence of sorts $\{O_i\}_{i < n}$ for some $n \leq \omega$ is written $\prod_{i < n} O_i$. The maps π_i are called *projection maps*, and the map $\langle \cdot, \dots \rangle$ is called the *product map*. \triangleleft

Definition 3.2.6. Given a (possibly many-sorted) signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T with a sort O and an \mathcal{L} -formula $\rho(x, y)$ with two free variables in O such that $T \models$

- $\forall x(\rho(x, x) = 0)$,
- $\forall xyz(\rho(x, z) \leq \rho(x, y) + \rho(y, z))$,
- $\forall xy(\rho(x, y) \geq 0)$, and
- $\forall xy(\rho(x, y) = \rho(y, x))$.¹

A *quotient sort imaginary expansion of T by ρ* , also called a *quotient sort*, is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ containing a single new sort Q with diameter bound $\max\{\sup I(\rho), 1\}$ and in which the only new relation symbol is d_Q (which by an abuse of notation we will typically write ρ) and the only new function symbol is $q : O \rightarrow Q$, which is ω_ρ -uniformly continuous. T' is T with the following new axioms added:

¹Or, equivalently as a single axiom, $T \models \forall xyz(|\rho(x, z) - \rho(z, y)| + 2|\rho(z, z)| \leq \rho(x, y))$.

- $(\forall xy:O)d_Q(x, y) = \rho(q(x), q(y))$ and
- $(\forall y:Q)(\exists x:O)y = q(x)$.

Q is typically written as O/ρ . The map q is called the *quotient map*. ◁

Note that in a quotient sort O/ρ , it is not generally true that for every $b \in (O/\rho)(\mathfrak{M})$ there is an $a \in O(\mathfrak{M})$ such that $q^{\mathfrak{M}}(a) = b$. Also note that just as in discrete logic, for any real formula $\varphi(x, \bar{y})$, the formula

$$\rho_\varphi(x, z) = \sup_{\bar{y}} |\varphi(x, \bar{y}) - \varphi(z, \bar{y})|$$

is a definable pseudo-metric, and furthermore if ρ is a definable pseudo-metric, then $\rho_\rho \equiv \rho$, so every definable pseudo-metric is of this form. Furthermore, we will see in Lemma 3.4.4 that pseudo-metrics are closely related to type-definable equivalence relations.

Definition 3.2.7. Given a (possibly many-sorted) signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T with a sort O and a closed \mathcal{L} -formula $D(x)$ with a single free variable in O such that $D(x)$ is definable over T , a *definable set sort imaginary expansion of T by D* , also called a *definable set sort*, is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ containing a single new sort E with diameter bound $\text{db}(O)$ and in which the only new relation symbol is d_E and the only new function symbol is $i : E \rightarrow O$, which is 1-Lipschitz. T' is T with the following new axioms added:

- $(\forall xy:E)d_E(x, y) = d_O(i(x), i(y))$ and
- $(\forall y \in D)(\exists x:E)y = i(x)$.

E is typically written $D(O)$. The map i is called the *inclusion map*. ◁

Note that in a definable set sort $D(O)$, it is always true that for every $b \in D(\mathfrak{M})$

there is an $a \in (D(O))(\mathfrak{M})$ such that $i^{\mathfrak{M}}(a) = b$. It should also be noted that unlike product and quotient sorts, definable set sorts depend very strongly on T . Product sorts do not depend on any way in the particular theory T and quotient sorts can be restricted to pseudo-metrics of the form ρ_φ without weakening the concept. In contrast to this there is no simple syntactic transformation, like $\varphi \mapsto \rho_\varphi$, which can take an arbitrary closed formula $F(x)$ and produces a closed formula $D_F(x)$ such that

- $D_F(x)$ is definable over the empty theory and
- for any theory T , if $F(x)$ is definable over T , then $D_F(x)$ is logically equivalent to $F(x)$ modulo T .

This is impossible for the simple reason that one can construct a closed formula $F(x)$ such that for every complete theory T , $F(x)$ is definable over T , but which is not itself definable over the empty theory.

Finally, in order to make certain results more uniform we will allow ourselves one more kind of imaginary.

Definition 3.2.8. Given a (possibly many-sorted) signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T , the *imaginary expansion of T by $\mathbf{2}$* , also called the *imaginary sort $\mathbf{2}$* , is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ containing a single new sort $\mathbf{2}$ with diameter bound 1 and in which the only new non-metric symbols are two constant symbols $\mathbf{0}$ and $\mathbf{1}$. T' is T with the following new axioms added:

- $d(\mathbf{0}, \mathbf{1}) = 1$ and
- $(\forall x:\mathbf{2})d(x, \mathbf{0}) = 0 \vee d(x, \mathbf{1}) = 0.$ \triangleleft

Remark 3.2.9. Of course, the sort $\mathbf{2}$ is very trivial by itself. We are really after the sort $\mathbf{2}^\omega$ and quotients of it, which can be any compact metric space (Fact A.2.14). In this way we can freely treat as imaginary sorts compact metric spaces in which each point is uniquely definable, which treatment we will explore more in Section 3.5. Also, as we will see in Section 3.5, formulas with variables in these sorts of sorts are equivalent to \mathbb{R}^K -valued formulas, where K is some compact metric space and we give \mathbb{R}^K the topology of uniform convergence. This kind of formula can also be thought of as a family of formulas continuously indexed by K .

Explicitly allowing the imaginary sort $\mathbf{2}$ is only necessary in theories in which it is possible for every sort to simultaneously have cardinality 0 or 1. If there is a sort O such that $T \models (\exists xy:O)d_Oxy \geq r$, for some $r > 0$, then we can realize $\mathbf{2}$ by the following operations: Take the quotient of O^2 by the pseudo-metric

$$\rho(xy, zw) = \frac{1}{r} |\min\{d_Oxy, r\} - \min\{d_Ozw, r\}|.$$

This sort is necessarily a compact metric space homeomorphic to some closed subset of $[0, 1]$ with a homeomorphism witnessed by a $[0, 1]$ -valued formula φ . By assumption we have that there are unique $a, b \in O^2/\rho$ (in any model) with $\varphi(a) = 0$ and $\varphi(b) = 1$. In particular these form a definable set D , and we can take the sort $D(O^2/\rho)$. This sort is now equivalent to $\mathbf{2}$. A similar procedure can be performed in theories T with a pair of sorts O and R such that $T \models (\exists xy:O)d_Oxy \geq r \vee (\exists xy:R)d_Rxy \geq r$, and other such situations, but no such procedure can be completely uniform in a theory in which all sorts can simultaneously be trivial. This occurs most frequently with empty

theories or universal theories. Given that it would be tedious to continually state the assumption “ T cannot consistently have all sorts empty or singleton” and given that the addition of $\mathbf{2}$ is useful, harmless (both informally and in the sense of Definition 3.2.14), and reducible to \mathbb{R}^K -valued formulas, I felt that this slightly more permissive definition of imaginary was appropriate.

The following proposition almost goes without saying, although note that it only partially extends to pre-structures.

Proposition 3.2.10. *For any \mathcal{L} -theory T , any model $\mathfrak{M} \models T$, and any imaginary expansion $T' \supseteq T$, \mathfrak{M} has a unique extension to a model of T' .*

Proof. This obviously follows for product, quotient, and definable set sorts, and the sort $\mathbf{2}$. The full statement follows by induction. \square

For reduced pre-models, this only works with product sorts and the sort $\mathbf{2}$. Reduced pre-models have a canonical extension by quotient sorts, but it is not the unique such extension. Pre-models in general never have unique extensions to imaginary expansions, unless the new sort is empty.

Definition 3.2.11. Given a (possibly incomplete) \mathcal{L} -theory T , an *imaginary expansion* of T is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ such that there exists chains $\{\mathcal{L}_i\}_{i < \alpha}$ and $\{T_i\}_{i < \alpha}$ such that

- $\mathcal{L}_0 = \mathcal{L}$ and $T_0 = T$,
- for each $i < \alpha$, T_{i+1} is either a product, quotient, or definable set sort expansion of T_i in the expanded language \mathcal{L}_{i+1} ,

- for each limit ordinal $\beta < \alpha$, $\mathcal{L}_\beta = \bigcup_{i < \beta} \mathcal{L}_i$ and $T_\beta = \bigcup_{i < \beta} T_i$, and
- $\mathcal{L}' = \bigcup_{i < \alpha} \mathcal{L}_i$ and $T' = \bigcup_{i < \alpha} T_i$. \triangleleft

Definition 3.2.12. Given a (possibly incomplete) \mathcal{L} -theory T , the *complete imaginary expansion* of T , written T^{eq} , is defined iteratively. Let $T_0 = T$ and $\mathcal{L}_0 = \mathcal{L}$. For each $i < \omega_1$, let $T_O \supseteq T_i$ be a single sort imaginary expansion in the language $\mathcal{L}_O \supseteq \mathcal{L}_i$. For distinct imaginary expansions with sorts O and O' , we take $\mathcal{L}_O \cap \mathcal{L}_{O'} = \mathcal{L}_i$. We let $T_{i+1} = \bigcup T_O$ and $\mathcal{L}_{i+1} = \bigcup \mathcal{L}_O$, where the unions are taken over all single sort imaginary expansions of \mathcal{L}_i . For each limit ordinal $\alpha < \omega_1$, let $T_\alpha = \bigcup_{i < \alpha} T_i$ and $\mathcal{L}_\alpha = \bigcup_{i < \alpha} \mathcal{L}_i$. Finally, let $T^{eq} = \bigcup_{i < \omega_1} T_i$ and $\mathcal{L}^{eq}(T) = \bigcup_{i < \omega_1} \mathcal{L}_i$. \triangleleft

In general, $\mathcal{L}^{eq}(T)$ depends on T and not just \mathcal{L} , but when no confusion can arise we will write \mathcal{L}^{eq} for $\mathcal{L}^{eq}(T)$.

Note that if T_1 is an extension of T_0 , then T_1^{eq} is not necessarily an extension of T_0^{eq} , although T_0^{eq} is interpretable in T_1^{eq} (in the sense of Definition 3.6.1). Usually we will only care about T^{eq} for complete T , so this won't be an issue. As in discrete logic, we will often freely pass to imaginary expansions, and we will often conflate T and T^{eq} .

Here we will see that although our construction of T^{eq} was in ω_1 stages, we could have stopped at 4.²

Proposition 3.2.13 (Imaginary Normal Form). *Let T be a (possibly incomplete) \mathcal{L} -theory.*

- (i) *For any sort $O \in \mathcal{L}^{eq}$, there are sorts $P, Q, E \in \mathcal{L}^{eq}$ and definable maps $f : O \rightarrow E$ and $g : E \rightarrow O$ over T such that P is a product of sorts in \mathcal{L} and the sort $\mathbf{2}$,*

²Technically we also need to know that imaginary expansions are harmless in the sense of Definition 3.2.14 to know that all of the relevant formulas in the proof of Proposition 3.2.13 are actually definable over a 4-step version of T^{eq} , but this isn't a very critical point.

$Q = P/\rho$ is a quotient of P , and $E = D(Q)$ is a definable subset of Q and such that f and g are isometries and inverses over T^{eq} .

(ii) If T is of the form T'_A for some complete \mathcal{L}' -theory T' and some set of parameters A , then P and Q can be taken to be imaginary sorts over T' .

Proof. (i) The statement is trivial for sorts of \mathcal{L} . Assume that we've shown the result for all sorts constructed before stage $\alpha < \omega_1$, and let O be an imaginary sort constructed at stage α .

• If $O = \prod_{i < n} O_i$ is a product sort (for some $n \leq \omega$), then for each $i < n$, we can find $P_i, \rho_i, Q_i, D_i, E_i, f_i, g_i$ with P_i a product of sorts in \mathcal{L} , $Q_i = P_i/\rho_i$, and $E_i = D_i(Q_i)$, with $f_i : E_i \rightarrow O_i$ and $g_i : O_i \rightarrow E_i$ definable maps that are isometries and inverses over T^{eq} . Assume that each P_i is $\prod_{k < n_i} R_{i,k}$, with $R_{i,k}$ either sorts in \mathcal{L} or $\mathbf{2}$. Let $\{R_{i(\ell),k(\ell)}\}_{\ell < m}$ (for some $m \leq \omega$) be an enumeration of all $R_{i,k}$, and let $P = \prod_{\ell < m} R_{i(\ell),k(\ell)}$. For each i, k , let $x_{i,k}$ and $y_{i,k}$ be variables of sort $R_{i,k}$, let \bar{x} and \bar{y} be listings of all of these variables, and for each $i < n$, let $\varphi_i(\bar{x}, \bar{y}) = \rho_i(\langle x_{i,0}, x_{i,1}, \dots \rangle, \langle y_{i,0}, y_{i,1}, \dots \rangle)$. Now we have that

$$\rho(\bar{x}, \bar{y}) = \sup_{i < \omega} \frac{\max\{\sup I(\delta_0), 1\}}{2^i \max\{\sup I(\delta_i), 1\}} \rho_i(\bar{x}, \bar{y})$$

is a pseudo-metric on P and that there is a definable bijection with definable inverse between $Q = P/\rho$ and $\prod_{i < n} Q_i$. Let $h : Q \rightarrow \prod_{i < n} Q_i$ be this map. We necessarily have that $\rho(z, w)$ and $d_{\prod_{i < n} Q_i}(h(z), h(w))$ are uniformly equivalent. Furthermore, the closed formula $F(x) = \bigwedge_{i < n} D_i(\pi_i(x))$ is a definable set on $\prod_{i < n} Q_i$, so we have that the image $h^{-1}(F)$ is a definable subset of Q as well. Let this definable set be $D(z)$, and let $E = D(Q)$, and now we have that h restricts to a bijection between $\prod_{i < n} D_i$ and E . If we let $c_i : E_i \rightarrow Q_i$ be the inclusion maps, then we have that $g(x) =$

$h^{-1}(\langle c_0 \circ \pi_0(x), c_1 \circ \pi_1(x), \dots \rangle)$ is the required definable bijection between O and E . Since it is a bijection in every model, by Proposition 2.3.79 we have that the inverse is a definable function as well, so let the inverse be f . By unpacking definitions, we get that the metric ρ we put on Q and therefore by extension on E agrees with the metric on O under the maps f and g , so we have the full required statement.

- If $O = R/\delta$ is a quotient sort with quotient map $u : R \rightarrow O$, then by the induction hypothesis we have that there are sorts P, Q, E with P a product of sorts in \mathcal{L} and $\mathbf{2}$, $Q = P/\rho$ a quotient of P with quotient map $q : P \rightarrow Q$, and $E = D(Q)$, a definable set sort with inclusion map $i : E \rightarrow Q$, and we have that there are definable isometries that are inverses $f : R \rightarrow E$ and $g : E \rightarrow R$. Let

$$\varphi(x:E, y:E) = \delta(g(x), g(y)).$$

By construction this is a pseudo-metric on E . Let

$$\psi(x:Q, y:Q) = \inf_{z, w \in D} \varphi(x, y) + \omega_\varphi(d_Q(xy, zw)).$$

We want to argue that for $a, b \in E$, $\psi(i(a), i(b)) = \varphi(a, b)$. Clearly by construction we have that $\psi(i(a), i(b)) \geq \varphi(a, b)$, furthermore by the definition of ω_φ , we have that for any other $c, e \in D$, that $\varphi(c, e) \leq \varphi(i(a), i(b)) + \omega_\varphi(d_Q(i(a)i(b), ce))$, implying that $\psi(i(a), i(b)) \leq \varphi(a, b)$ as well, so we have that they are equal. Let

$$\chi(x:Q, y:Q) = \sup_{z \in D} |\psi(x, z) - \psi(y, z)|.$$

By construction, $\chi(x, y)$ is a pseudo-metric on Q and furthermore we have that for

$a, b \in D$, $\psi(a, b) = \chi(a, b)$. Now if we let $\eta(x:P, y:P) = \chi(q(x), q(y))$ we get that η is a pseudo-metric on P as well. Let $q_0 : Q \rightarrow Q/\chi$ and $q_1 : P \rightarrow P/\eta$ be the quotient maps. We have that $F(x, y) = (\exists z:P)x = q_0(z) \wedge y = q_1(q(z))$ is the graph of a bijection between Q/χ and P/η in every model of T , so by Proposition 2.3.79, we get a definable bijection map $h : Q/\chi \rightarrow P/\eta$.

Now we have that the (metric closure of the) image $D' = \overline{h(q_0(i(E)))}$ is a definable subset of P/η . Let $E' = D'(P/\eta)$, with inclusion map $c : E' \rightarrow P/\eta$. Now by unpacking what we have done we get that the formula

$$G(x, y) = (\exists z:R)x = u(z) \wedge c(y) = h(q_0(i(f(z))))$$

is the graph of a bijection between O and E' in every model of T and that furthermore it is an isometry, so $P, P/\eta, E'$ is the required sequence of sorts for O .

- If $O = D(R)$ is a definable set sort with inclusion map $i : O \rightarrow R$, then by the induction hypothesis we have that there are sorts P, Q, E' with P a product of sorts in \mathcal{L} and $\mathbf{2}$, Q a quotient of P , and $E' = D'(P)$ a definable set sort, such that there are definable isometries that are inverses $f : R \rightarrow E'$ and $g : E' \rightarrow R$. Definable subsets of definable sets are themselves definable, so we have that the formula $D''(x) = D(i'(x))$, where $i' : E' \rightarrow Q$ is the inclusion map, is a definable subset of Q . Taking $E'' = D''(Q)$, with inclusion map $i'' : E'' \rightarrow Q$, we have that the formula $i'(f(i(x))) = i''(y)$ is the graph of a bijection between O and E'' in every model, so by Proposition 2.3.79, we get the required bijections. These are isometries by construction.

- If O is a copy of the sort $\mathbf{2}$, then the statement is trivial.

Therefore by induction we have that the result holds for all sorts in \mathcal{L}^{eq} .

(ii) By part (i), we can find for any sort O over T^{eq} , sorts P, Q, E with P a product of sorts in $\mathcal{L} = \mathcal{L}'_A$ (which, note, are the same as the sorts in \mathcal{L}'), $Q = P/\rho$, a quotient sort with quotient map $q : P \rightarrow Q$, and $E = D(Q)$, a definable set sort with inclusion map $i : E \rightarrow Q$. Write ρ as $\rho(x, y, \bar{a})$, where $\bar{a} \in A$ is an at most countable tuple of parameters sufficient to define ρ . \square

Now, Definition 3.2.11 is all well and good, but how can we say that it is *the* correct notion of imaginary expansion, rather than some admissible forms of imaginary expansion? As it happens we can give a tight characterization of this notion of imaginary expansion.

Definition 3.2.14. Given a signature \mathcal{L} and a (possibly incomplete) \mathcal{L} -theory T , a (*single sort*) *harmless expansion of T* is an expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ that contains a single new sort O' such that

- T' is conservative over T , i.e. for any closed \mathcal{L} -sentence F , $T' \models F$ if and only if $T \models F$,
- for any \mathcal{L}' -formula $\varphi(\bar{x})$ such that the sort of each x_i is in \mathcal{L} , there is an \mathcal{L} -formula $\psi(\bar{x})$ such that $T' \models \forall \bar{x}(\varphi(\bar{x}) = \psi(\bar{x}))$, and
- for any $\mathfrak{M} \models T'$, $O'(\mathfrak{M}) \subseteq \text{dcl}(M \upharpoonright \mathcal{L})$, where $M \upharpoonright \mathcal{L}$ is the set of elements of \mathfrak{M} contained in the sorts of \mathcal{L} .

An expansion $T' \supseteq T$ in an expanded language $\mathcal{L}' \supseteq \mathcal{L}$ is a *harmless expansion of T* , or just *harmless*, if there is an enumeration $\{O_i\}_{i < \alpha}$ of the new sorts in \mathcal{L}' such that for each $i < \alpha$, $T' \upharpoonright O_{\leq i}$ is a harmless expansion of $T' \upharpoonright O_{< i}$. \triangleleft

It is straightforward but tedious to verify the following:

Proposition 3.2.15. *Any imaginary expansion is harmless.* \square

The point of the concept of a harmless expansion is to validate our definition of imaginary expansion. Every harmless expansion is equivalent to an imaginary expansion in a strong sense. The proof of the following theorem uses concepts from a later section (Section 3.5), but we have chosen to put it here for the sake of organization.

Lemma 3.2.16. *Fix a theory T . Suppose that $\{H_i\}_{i \in I}$ and O are sorts such that for any $\mathfrak{M} \models T$, $H_i(\mathfrak{M})$ and $O(\mathfrak{M})$ are non-empty³ and $O(\mathfrak{M}) \subseteq \text{dcl} \bigcup_{i \in I} H_i(\mathfrak{M})$, then there exists a definable partial function $f(\bar{x}, r)$ with $\mathcal{S}(r) \in \mathbf{2}^\omega$, $\mathcal{S}(x_k) \in \{H_i\}_{i \in I}$, and $\mathcal{S}(r) = K$, such that for any $\mathfrak{M} \models T$ and $a \in O(\mathfrak{M})$, there is $\bar{b} \in \bigcup_{i \in I} H_i(\mathfrak{M})$ and $c \in K(\mathfrak{M})$ such that $f(\bar{b}, c) \downarrow = a$.*

Proof. Let y be a variable of sort O . Fix a type $p(y) \in S_y(T)$. For every $k < \omega$, by compactness there must be a finite list $\varphi_p^0(\bar{x}_p^0, y), \dots, \varphi_p^{n_p-1}(\bar{x}_p^{n_p-1}, y)$ of real formulas (where the sorts of each tuple \bar{x}_p^i come from $\{H_i\}_{i \in I}$) such that for any $\mathfrak{M} \models T$ and $a \in \mathfrak{M}$ with $\mathfrak{M} \models p(a)$, for some $i < n$, $\mathfrak{M} \models \exists \bar{x}_p^i \varphi_p^i(\bar{x}_p^i, a) < 2^{-k} \wedge \forall z (\varphi_p^i(\bar{x}_p^i, z) \leq 2^{-k} \rightarrow d(a, z) < 2^{-k+1})$. Or in other words,

$$p(y) \models \bigvee_{i < n_p} \exists \bar{x}_p^i \varphi_p^i(\bar{x}_p^i, y) < 2^{-k} \wedge \forall z (\varphi_p^i(\bar{x}_p^i, z) \leq 2^{-k} \rightarrow d(y, z) < 2^{-k+1}),$$

but since $p(y)$ is a complete type, there must be some $i < n_p$ such that $p(y) \models \exists \bar{x}_p^i \varphi_p^i(\bar{x}_p^i, y) < 2^{-k} \wedge \forall z (\varphi_p^i(\bar{x}_p^i, z) \leq 2^{-k} \rightarrow d(y, z) < 2^{-k+1})$. Call this formula $U_p^k(y)$, and let i_p^k be the corresponding i .

³This is necessary to avoid situations like this: Let \mathcal{L} be a signature with three sorts A, B, C and function symbols $f : A \rightarrow C$ and $g : B \rightarrow C$. Let T be an \mathcal{L} -theory that says that all three metrics are discrete and that f and g are injections with disjoint images that cover C . There is no definable function $f(x:A, y:B, i:2)$ witnessing that C is always contained in $\text{dcl}(A \cup B)$ because one of A or B could be empty.

By compactness, there is a finite list $p_0^k, \dots, p_{\ell_k-1}^k$ of types in $S_y(T)$ such that $\bigcup_{j < \ell_k} \llbracket U_{p_j^k}^k \rrbracket = S_y(T)$. By Fact A.2.16, we can find a family of closed formula F_j^k such that for each $j < \ell_k$, $\llbracket F_j^k \rrbracket \subseteq \llbracket U_{p_j^k}^k \rrbracket$ and such that $\bigcup_{j < \ell_k} \llbracket F_j^k \rrbracket = S_y(T)$.

Let $K = \prod_{k < \omega} \ell_k$, and let \bar{x} be a concatenation of the tuples $\bar{x}_{p_j^k}^{i_j^k}$ for each $k < \omega$ and $j < \ell_k$. For sanity's sake let $i_j^k = i_{p_j^k}^k$, $\bar{x}_j^k = \bar{x}_{p_j^k}^{i_j^k}$, $\varphi_j^k = \varphi_{p_j^k}^{i_j^k}$, and $U_j^k = U_{p_j^k}^k$. Let r be a variable of sort K . For each $k < \omega$, let

$$\begin{aligned} G_k(\bar{x}, r, y) &= \bigvee_{j < \ell_k} r(k) = j \\ &\quad \wedge F_j^k(y) \\ &\quad \wedge \varphi_j^k(\bar{x}_j^k, y) \leq 2^{-k} \\ &\quad \wedge \forall z (\varphi_j^k(\bar{x}_j^k, y) < 2^{-k} \rightarrow d(y, z) \leq 2^{-k+1}). \end{aligned}$$

(Note that the last two lines are weaker than the corresponding sub-formula of U_j^k .) By construction, for any $\mathfrak{M} \models T$ and $\bar{a}, c \in \mathfrak{M}$, if $G_\omega(\bar{a}, c, y)$ is consistent, there is a unique $b \in \mathfrak{M}$ such that $\mathfrak{M} \models G_\omega(\bar{a}, c, b)$, so by Proposition 2.3.79 there is a definable partial function $\chi(\bar{x}, r, y)$ such that $G_\omega(\bar{x}, r, y) \equiv_T (\chi(\bar{x}, r, y) = 0)$.

For any type $p(y) \in S_y(T)$, we have by construction that $\llbracket p(y) \rrbracket_{\bar{x}ry} \cap \bigcap_{m < k} \llbracket G_m \rrbracket_{\bar{x}ry}$ is non-empty for every $k < \omega$. Therefore by compactness if we set $G_\omega(\bar{x}, r, y) = \bigwedge_{k < \omega} G_k(\bar{x}, r, y)$ we have that $\llbracket p(y) \rrbracket_{\bar{x}ry} \cap \llbracket G_\omega \rrbracket_{\bar{x}ry}$ is non-empty. This implies that $T \models \forall y \exists \bar{x} r G_\omega(\bar{x}, r, y)$. We need to show actual existence, though.

Fix $\mathfrak{M} \models T$ and $a \in O(\mathfrak{M})$. For each $k < \omega$, we have that $\mathfrak{M} \models F_j^k(a)$ for some $j < \ell_k$. Let c be the sequence of these j 's. We have that $\mathfrak{M} \models U_j^k(a)$. This implies that there exists \bar{b}^k , of the same sorts as \bar{x}_j^k , such that $\mathfrak{M} \models \varphi_j^k(\bar{b}^k, a) < 2^{-k} \wedge \forall z \varphi_j^k(\bar{b}^k, z) \leq$

$2^{-k} \rightarrow d(a, z) < 2^{-k+1}$. Let \bar{b} be a tuple of elements of the same sorts as \bar{x} , where the spots corresponding to each \bar{x}_j^k are filled in with \bar{b}^k and the others are filled in arbitrarily (here we use that each H_i is always non-empty). Now we clearly have for each $k < \omega$ that

$$\mathfrak{M} \models F_{c^{(k)}}^k(a) \wedge \varphi_{c^{(k)}}^k(\bar{b}^k, a) \leq 2^{-k} \wedge \forall z(\varphi_j^k(\bar{b}^k, a) < 2^{-k} \rightarrow d(a, z) \leq 2^{-k+1}),$$

and therefore $\mathfrak{M} \models G_\omega(\bar{b}, c, a)$, as required.

By embedding K in 2^ω , we may change the sort of r to $\mathbf{2}^\omega$. Extending χ with the Tietze extension theorem (Fact A.2.8) gives the required $f(\bar{x}, r)$. \square

The following lemma is essentially the same as Lemma 1.24 in [BY10a].

Lemma 3.2.17. *Fix a theory T . If $f(\bar{x})$ is a partial definable function with codomain O (and such that $\mathcal{S}(x_i) = H_i$) such that $T \models (\forall y:O)\exists \bar{x}f(\bar{x})\downarrow = y$, then there exists a definable pseudo-metric $\rho(\bar{x}, \bar{y})$ on $\prod_{i < |\bar{x}|} H_i$ such that the image of $\llbracket f(\bar{x})\downarrow \rrbracket$ under the natural quotient map $q : \prod_{i < |\bar{x}|} H_i \rightarrow \prod_{i < |\bar{x}|} H_i/\rho$ is a definable set, D . Furthermore, there is a definable isometry $g : D \rightarrow O$ such that whenever $f(\bar{a})\downarrow$, then $f(\bar{a}) = g(q(\bar{a}))$.*

Proof. Let $\varphi(\bar{x}, y)$ be a real formula defining $f(\bar{x})$. Let $\rho(\bar{x}, \bar{y}) = \sup_{z:O} |\varphi(\bar{x}, z) - \varphi(\bar{y}, z)|$. This is a pseudo-metric by construction. Also by construction there is a real formula $\psi(w, z)$ with $\mathcal{S}(w) = \prod_{i < |\bar{x}|} H_i/\rho$ such that for any \bar{x} , $\varphi(\bar{x}, z) = \psi(q(\bar{x}), z)$. Let $\chi(w)$ be the formula $\inf_{z:O} \sup_{u:O} |\psi(w, z) - d_O(u, z)|$.

We want to argue that $\chi(w)$ is the distance predicate of the required definable set. Let \bar{a} be such that $f(\bar{a})\downarrow = b$. Then we have that $\chi(q(\bar{a})) = 0$, so $\llbracket \chi(w) = 0 \rrbracket$ contains the image of $\llbracket f(\bar{x})\downarrow \rrbracket$ under q . Suppose that $\chi(\bar{a}) = r$. Then in the monster there

is some b in sort O such that $\sup_{u:O} |\psi(w, b) - d_O(u, b)| = r$. By assumption there is also \bar{c} such that $f(\bar{c})\downarrow = b$. This implies that $\rho(\bar{a}, \bar{c}) = r$. Therefore we have that $d_{\text{inf}}(p, \llbracket \chi(w) = 0 \rrbracket) \leq \chi(p)$ for all types p in the appropriate variables. Conversely, assume that $d_{\text{inf}}(p, \llbracket \chi(w) = 0 \rrbracket) = r$, then there exists a type $q \in \llbracket \chi(w) = 0 \rrbracket$ such that $d(p, q) = r$. If \bar{a} is a realization of p and \bar{b} is a realization of q such that $d(\bar{a}, \bar{b}) = r$, then in the monster there exists c in sort O such that $\varphi(\bar{b}, e) = d(c, e)$ for all e in sort O , but this is the same as saying that $f(\bar{b})\downarrow = c$, therefore we have that $\sup_{u:O} |\psi(q(\bar{a}), c) - d_O(u, c)| = r$ and we have that $\chi(q(\bar{a})) \leq r$, so $d_{\text{inf}}(p, \llbracket \chi(w) = 0 \rrbracket) \geq \chi(p)$, and hence $d_{\text{inf}}(p, \llbracket \chi(w) = 0 \rrbracket) = \chi(p)$ for all types p in the appropriate variables.

The required function g exists by Proposition 2.3.79. \square

Theorem 3.2.18. *Let T be an \mathcal{L} -theory, and let T' be a single sort harmless expansion of T in the language $\mathcal{L}' \supseteq \mathcal{L}$ with new sort O' . There is an imaginary expansion $T'' \supseteq T$ in a language $\mathcal{L}'' \supseteq \mathcal{L}$ (which we take such that $\mathcal{L}' \cap \mathcal{L}'' = \mathcal{L}$) with an imaginary sort O'' such that there are definable functions $f : O' \rightarrow O''$ and $g : O'' \rightarrow O'$ over $T' \cup T''$ ⁴ such that*

$$T' \cup T'' \models \forall x g(f(x)) = x \wedge \forall x f(g(x)) = x \wedge \forall xy d_{O''}(f(x), f(y)) = d_{O'}(x, y).$$

Moreover, O'' can be taken to be a definable subset of a quotient of a product of sorts of the form $H \sqcup \mathbf{1}$ (for $H \in \mathcal{L}$) and the imaginary sort $\mathbf{2}$.

Proof. Let T^* be an imaginary expansion of T where we add $\mathbf{2}^\omega$ and $H \sqcup \mathbf{1}$ for each sort $H \in \mathcal{L}$. $T^* \cup T'$ has the property that for any $\mathfrak{M} \models T^* \cup T'$, $O'(\mathfrak{M})$ is contained in the definable closure of the sorts of the form $H \sqcup \mathbf{1}$ (with $H \in \mathcal{L}$), so by Lemma 3.2.16 for

⁴Which is consistent by the Craig interpolation theorem (Proposition B.1.8).

some $n \leq \omega$ we can find a definable partial function $h : \mathbf{2}^\omega \times \prod_{i < n} (H_i \sqcup \mathbf{1}) \rightarrow O$ such that for any $\mathfrak{M} \models T^* \cup T'$, $O(\mathfrak{M})$ is contained in the image of h .

By Lemma 3.2.17, we can find a definable pseudo-metric ρ on $\mathbf{2}^\omega \times \prod_{i < n} (H_i \sqcup \mathbf{1})$ and a definable subset $D \subseteq \mathbf{2}^\omega \times \prod_{i < n} (H_i \sqcup \mathbf{1}) / \rho$ with a definable isometry $g : D \rightarrow O'$. Let f be its inverse.

Now we just need to argue that ρ and D are definable over T^* . For each $\alpha, \beta \in \mathbf{2}^\omega$, let $A_{\alpha, \beta}$ be the definable subset of $\mathbf{2}^\omega \times \prod_{i < n} (H_i \sqcup \mathbf{1})$ of elements of the form $\langle \alpha, \bar{b} \rangle$ where $b_i \in H_i$ if and only if $\beta(i) = 1$ (note that this is definable over \mathcal{L}^*). For each $\alpha, \beta \in \mathbf{2}^\omega$, the restrictions $\rho \upharpoonright A_{\alpha, \beta} \times A_{\alpha, \beta}$ and $D \upharpoonright A_{\alpha, \beta}$ (where we're thinking of D as its distance predicate) correspond to some formulas $\varphi_{\alpha, \beta}(\bar{x}, \bar{y})$ and $\psi_{\alpha, \beta}(\bar{x})$, where \bar{x} and \bar{y} are sequences of variables corresponding to H_i for i such that $\beta(i) = 1$. These are \mathcal{L}' -formulas whose only free variables are in \mathcal{L} , therefore, since T' is a harmless expansion, there are \mathcal{L} -formulas $\varphi_{\alpha, \beta}^\dagger(\bar{x}, \bar{y})$ and $\psi_{\alpha, \beta}^\dagger(\bar{x})$ which are logically equivalent modulo T' .

Since we can do this for any α and β , this implies that there are \mathcal{L}^* -formulas ρ' and D' equivalent to ρ and D over $T^* \cup T'$. Let T'' be the imaginary expansion of T^* where we add a sort O'' for D . We have that O'' satisfies the required properties. \square

3.3 T^{eq} for Dictionaric T

Now that we have the machinery in place, we can discuss the properties of definable sets in T^{eq} .

Proposition 3.3.1. *If T is dictionaric, then T^{eq} is dictionaric over parameters from the home sort (i.e. if A is a set of parameters from the home sort and I is some imaginary, then the type space $S_I(A)$ is dictionaric). In particular T^{eq} is dictionaric over models.*

Furthermore, if T is dictionaric over models, then T^{eq} is as well.

Proof. By Proposition 3.2.13 every imaginary is equivalent to a definable subset of a definable quotient of the sort of ω -tuples, so the result follows from Propositions 2.4.3, 2.4.9, and 2.4.10. \square

Unfortunately in general we cannot conclude that T^{eq} is itself dictionaric. There are even discrete counterexamples, such as RCF, where if we consider the hyperimaginary given by quotienting by the ‘infinitesimally close’ equivalence relation (which can be seen as a continuous imaginary sort) and if we let a be any infinite element of this hyperimaginary, then $S_1(a)$ fails to be dictionaric in the connected component of $\text{tp}(a)$. This phenomenon is related to failure of elimination of hyperimaginaries. If T is a discrete theory that eliminates hyperimaginaries, then the continuous T^{eq} is dictionaric, so this will hold for any stable or supersimple discrete theory [PP87, BPW00]. It would be nice to know if this extends to continuous dictionaric theories.

Question 3.3.2. *If T is a continuous dictionaric theory and T is also stable or supersimple, does it follow that T^{eq} is dictionaric?*

This issue, fortunately, is not relevant in the main results of this thesis, as every totally transcendental (and in particular ω -stable) theory is dictionaric (and so also has dictionaric T^{eq}). This is a topometric analog of the fact that scattered compact Hausdorff spaces are automatically totally disconnected.

3.4 Additional Imaginary Constructions

Given the primitive imaginary-forming operations, many more are ‘constructible.’

Proposition 3.4.1 (Compact Metric Space Sorts). *For any compact metric space X , there is an imaginary sort K_X (over the empty theory) such that for each $x \in X$, there is a fix formula $\varphi_x(y)$ such that in any structure \mathfrak{M} , $K_X(\mathfrak{M})$ is isometric to X with isometry $f : X \rightarrow K_X(\mathfrak{M})$.*

Proof. If X is empty, then we can pass to an empty definable set to get K_X . Otherwise, if X is non-empty, by Fact A.2.14 there is a pseudo-metric on 2^ω , ρ such that $2^\omega/\rho$ is isometric to X . Since every element of 2^ω is defined by a term in every model, $S_{xy:2^\omega}(\mathcal{L})$ is homeomorphic to $(2^\omega)^2 \times S_0(\mathcal{L})$, so ρ is actually a definable pseudo-metric on 2^ω , and we can take $K_X = 2^\omega/\rho$. \square

Notation 3.4.2 (Finite Ordinal Sorts). For any $n < \omega$, we let \mathbf{n} denote K_n , where n is taken to be a compact metric space with n points with pairwise distance 1. We take the elements of \mathbf{n} to be labeled by constants $\mathbf{0}, \mathbf{1}, \dots, \mathbf{n} - \mathbf{1}$ (technically 0-ary functions definable over the empty theory). Furthermore, we may write quantifiers annotated with one of these sorts with expressions such as $(\forall x < \mathbf{n})$, rather than $(\forall x:\mathbf{n})$. \triangleleft

Just as with ordinary ordinals, the severe overloading of boldface numerals should not cause any confusion. Note though, that there is in general no sort \mathbf{n} with $n \geq \omega$.

Proposition 3.4.3 (Finite Disjoint Union Sorts). *For any pair of sorts O and O' , for any $r \geq \frac{1}{2} \max\{\text{db}(O), \text{db}(O')\}$, there is an imaginary sort $O \sqcup_r O'$ (over the empty theory) with definable functions $f : O \rightarrow O \sqcup_r O'$ and $g : O' \rightarrow O \sqcup_r O'$ and a $[0, 1]$ -valued formula $\varphi(x)$ such that the following hold:*

- $(\forall xy:O)d(x, y) = d(f(x), f(y))$,
- $(\forall xy:O')d(x, y) = d(g(x), g(y))$,

- $(\forall x:O)(\forall y:O')d(f(x), g(y)) = r,$
- $(\forall x:O)\varphi(f(x)) = 0,$
- $(\forall x:O \sqcup_r O')(\exists y:O)x=f(y)$
- $(\forall x:O')\varphi(g(x)) = 1.$
- $\vee(\exists z:O')x=g(z),$ and

In particular, in any structure \mathfrak{M} , $(O \sqcup_r O')(\mathfrak{M})$ is isometric to $O(\mathfrak{M}) \sqcup_r O'(\mathfrak{M})$, as witnessed by $f^{\mathfrak{M}}$ and $g^{\mathfrak{M}}$.

Proof. Consider the imaginary sort $O \times O' \times \mathbf{2}$ and the definable pseudo-metric

$$\rho((x, y, i), (z, w, j)) = d(i, \mathbf{1})d(j, \mathbf{1})d(x, z) + d(i, \mathbf{0})d(j, \mathbf{1})d(y, w) + d(i, j)r.$$

$O \times O' \times \mathbf{2}/\rho$ is the required imaginary sort. If we let $q : O \times O' \times \mathbf{2} \rightarrow O \times O' \times \mathbf{2}/\rho$ be the natural quotient map, then

$$d(f(x), y) = \inf_{z \in O'} \rho(q(x, z, \mathbf{0}), y) \text{ and}$$

$$d(g(x), y) = \inf_{z \in O} \rho(q(z, x, \mathbf{1}), y). \quad \square$$

We could have taken finite disjoint unions as a primitive imaginary forming operation and allowed empty products (where an empty product contains a single element $\langle \rangle$, the empty tuple), and this would also have allowed us to construct $\mathbf{2}$ in a uniform way.

The following proposition was originally shown by Ben Yaacov in the context of CATs [BY05, Thm. 2.20].

Lemma 3.4.4. *Let T be a theory, and suppose that $E(\bar{x}, \bar{y})$ is a closed type-set formula with \bar{x}, \bar{y} tuples of variables of the same sorts such that for any $\mathfrak{M} \models T$ and $\bar{a}, \bar{b}, \bar{c} \in \mathfrak{M}$,*

- $\mathfrak{M} \models E(\bar{a}, \bar{a}),$

- if $\mathfrak{M} \models E(\bar{a}, \bar{b})$, then $\mathfrak{M} \models F(\bar{b}, \bar{a})$, and
- if $\mathfrak{M} \models E(\bar{a}, \bar{b}) \wedge E(\bar{b}, \bar{c})$, then $\mathfrak{M} \models E(\bar{a}, \bar{c})$

(i.e. T entails that $E(\bar{x}, \bar{y})$ is an equivalence relation). Then

- (i) $E(\bar{x}, \bar{y})$ is logically equivalent modulo T to a conjunction of closed formulas $F(\bar{x}, \bar{y})$ which are equivalence relations over T , and
- (ii) if $E(\bar{x}, \bar{y})$ is a closed formula, then there is a real formula $\rho(\bar{x}, \bar{y})$ such that $E(\bar{x}, \bar{y}) \equiv_T (\rho(\bar{x}, \bar{y}) = 0)$ and $\rho(\bar{x}, \bar{y})$ is a pseudo-metric over T .

Proof. (i) Let $U_0(\bar{x}, \bar{y})$ be an open formula such that $\llbracket E(\bar{x}, \bar{y}) \rrbracket_T \subseteq \llbracket U_0(\bar{x}, \bar{y}) \rrbracket$. By compactness, for any i , there exists a closed formula $F_{i+1}(\bar{x}, \bar{y})$ and open formula $U_{i+1}(\bar{x}, \bar{y})$ such that $\llbracket E(\bar{x}, \bar{y}) \rrbracket_T \subseteq \llbracket U_{i+1}(\bar{x}, \bar{y}) \rrbracket \subseteq \llbracket F_{i+1}(\bar{x}, \bar{y}) \rrbracket \subseteq \llbracket U_i(\bar{x}, \bar{y}) \rrbracket$ and such that

$$\begin{aligned} T \models \forall \bar{x} \bar{y} \bar{z} U_{i+1}(\bar{x}, \bar{x}) \\ \wedge (F_{i+1}(\bar{x}, \bar{y}) \rightarrow U_i(\bar{x}, \bar{y})) \\ \wedge (F_{i+1}(\bar{x}, \bar{y}) \wedge F_{i+1}(\bar{y}, \bar{z}) \rightarrow U_i(\bar{x}, \bar{y})). \end{aligned}$$

Therefore the closed formula $E'(\bar{x}, \bar{y}) = \bigwedge_{i < \omega} F_i(\bar{x}, \bar{y})$ is an equivalence relation over T with $T \models \forall \bar{x} \bar{y} E(\bar{x}, \bar{y}) \rightarrow E'(\bar{x}, \bar{y})$. Since we can do this for any open neighborhood of $E(\bar{x}, \bar{y})$, we have that $E(\bar{x}, \bar{y})$ is logically equivalent to the conjunction of such formulas.

(ii) Assume that $E(\bar{x}, \bar{y})$ is a closed formula, and by removing variables that are not actually free in E , assume that \bar{x} and \bar{y} are at most countable tuples of the same sorts. Pass to a countable reduct $\mathcal{L}_0 \subseteq \mathcal{L}$ such that $E(\bar{x}, \bar{y})$ is still an \mathcal{L}_0 -formula. Let T_0 be the set of all \mathcal{L}_0 -sentences entailed by T . Note that $E(\bar{x}, \bar{y})$ is still an equivalence

relation over T_0 . Now consider the type space $S_{\bar{x}\bar{z}}(T_0)$, with \bar{z} a tuple of variables of the same sorts as \bar{x} . Define an equivalence relation \sim on $S_{\bar{x}\bar{z}}(T_0)$ by $p(\bar{x}, \bar{z}) \sim q(\bar{x}, \bar{z})$ if and only if there exists $\mathfrak{M} \models T_0$ with $\bar{a}, \bar{b}, \bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models p(\bar{a}, \bar{c}) \wedge q(\bar{b}, \bar{c}) \wedge E(\bar{a}, \bar{b})$. By considering automorphisms of the monster model of $\text{Th}(\mathfrak{M})$, we have that this is an equivalence relation.

We need to argue that for any open set $U \subseteq S_{\bar{x}\bar{z}}(T_0)$, the set

$$U^\sim = \{p \in S_{\bar{x}\bar{z}}(T_0) : (\exists q \in U)p \sim q\}$$

is also an open set. (We will use this notation for arbitrary sets of types.) Fix $U(\bar{x}, \bar{z}) = (\varphi(\bar{x}, \bar{z}) > 0)$. Since \mathcal{L}_0 and $\bar{x}\bar{z}$ are countable, every open subset of $S_{\bar{x}\bar{z}}(T_0)$ is of this form (Facts A.1.3 and A.2.12). Let $\psi(\bar{x}, \bar{y})$ be a $[0, 1]$ -valued formula such that $E(\bar{x}, \bar{y}) \equiv (\psi(\bar{x}, \bar{y}) = 0)$. By compactness, for each $k < \omega$, there exists $\varepsilon_k > 0$ such that if $p \notin \llbracket \varphi \geq 2^{-k} \rrbracket_{T_0}^\sim$ (which is closed, by compactness), then for every $\bar{a}\bar{b}\bar{c} \in \mathfrak{M} \models T_0$, if $\mathfrak{M} \models p(\bar{a}, \bar{c}) \wedge \varphi(\bar{b}, \bar{c}) \geq 2^{-k}$, then $\mathfrak{M} \models \psi(\bar{a}, \bar{b}) \geq \varepsilon_k$. This implies that $\llbracket \exists \bar{y} \varphi(\bar{y}, \bar{z}) > 2^{-k} \wedge \psi(\bar{x}, \bar{y}) < \varepsilon_k \rrbracket_{T_0} \subseteq \llbracket U(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim$. Therefore in particular we have that

$$\left[\bigvee_{k < \omega} \exists \bar{y} \varphi(\bar{y}, \bar{z}) > 2^{-k} \wedge \psi(\bar{x}, \bar{y}) < \varepsilon_k \right]_{T_0} \subseteq \llbracket U(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim.$$

For the reverse inclusion, find $q \in \llbracket U(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim$ and $p \in \llbracket U(\bar{x}, \bar{z}) \rrbracket_{T_0}$ with $p \sim q$. For any $\mathfrak{M} \models T_0$ and $\bar{b}\bar{c} \in \mathfrak{M}$ with $\mathfrak{M} \models q(\bar{b}, \bar{c})$, we can find an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ such that for some $\bar{a} \in \mathfrak{N}$, $\mathfrak{N} \models p(\bar{a}\bar{b}) \wedge E(\bar{a}, \bar{b})$. Therefore

$$\mathfrak{N} \models \bigvee_{k < \omega} \exists \bar{y} \varphi(\bar{y}, \bar{c}) > 2^{-k} \wedge \psi(\bar{b}, \bar{y}) < \varepsilon_k,$$

so by elementarity \mathfrak{M} does as well. Since we can do this for any $q \in \llbracket U(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim$, we have the reverse inclusion and the two sets are the same.

Let X be $S_{\bar{x}\bar{z}}(T_0)/\sim$, the quotient space, with quotient map $f : S_{\bar{x}\bar{z}}(T_0) \rightarrow X$. X is automatically compact. We need to argue that X is Hausdorff. For any pair of distinct $a, b \in X$, we have that $f^{-1}(a)$ and $f^{-1}(b)$ are disjoint. Furthermore, they are each the equivalence class of a single type, so they are in fact closed, by compactness. By compactness we have that $f^{-1}(a) = \bigcap \{(\text{cl } U)^\sim : U \supseteq f^{-1}(a), U \text{ open}\}$, and likewise for b . Therefore, since $f^{-1}(a)$ and $f^{-1}(b)$ are disjoint, there must be $U \supseteq f^{-1}(a)$ and $V \supseteq f^{-1}(b)$ such that $(\text{cl } U)^\sim$ and $(\text{cl } V)^\sim$ are disjoint. Therefore U^\sim and V^\sim are disjoint. Since these are open we have that $f(U^\sim)$ and $f(V^\sim)$ are disjoint open neighborhoods of p and q and X is Hausdorff.

Since \mathcal{L} and $\bar{x}\bar{z}$ are countable, X is metrizable as well. Therefore the collection of continuous functions $X \rightarrow [0, 1]$ is separable under the uniform norm. Let $\{g_i\}_{i < \omega}$ be a countable sequence of continuous functions from X to $[0, 1]$ dense in the uniform norm. For each $i < \omega$, the function $g_i \circ f$ is continuous on $S_{\bar{x}\bar{y}}(T_0)$, so we can find an $\mathcal{L}(\bar{x}\bar{z})$ -formula $\varphi_i(\bar{x}, \bar{z})$ with $I(\varphi_i) \subseteq [0, 1]$ such that for any $p \in S_{\bar{x}\bar{y}}(T_0)$, $\varphi_i(p) = g_i \circ f(p)$. Now let

$$\rho(\bar{x}, \bar{y}) = \sup_{i < \omega} \sup_{\bar{z}} 2^{-i} |\varphi_i(\bar{x}, \bar{z}) - \varphi_i(\bar{y}, \bar{z})|.$$

Note that since each φ_i has $I(\varphi_i) \subseteq [0, 1]$, this is a real formula, and it is a pseudo-metric by construction. We want to argue that $E(\bar{x}, \bar{y})$ and $(\rho(\bar{x}, \bar{y}) = 0)$ are logically equivalent over T_0 . For any $\mathfrak{M} \models T_0$ and $\bar{a}\bar{b} \in \mathfrak{M}$, if $\mathfrak{M} \models E(\bar{a}, \bar{b})$, then for any $\bar{c} \in \mathfrak{M}$ of the same sorts as \bar{a} , we have that $f(\text{tp}(\bar{a}\bar{c})) = f(\text{tp}(\bar{b}\bar{c}))$ by construction, so $\mathfrak{M} \models \varphi_i(\bar{a}, \bar{c}) = \varphi_i(\bar{b}, \bar{c})$ for each $i < \omega$ and we have that $\mathfrak{M} \models \rho(\bar{a}, \bar{b}) = 0$. Conversely, assume that $\mathfrak{M} \models$

$\rho(\bar{a}, \bar{b}) = 0$. This implies that for every $i < \omega$, $\mathfrak{M} \models \varphi_i(\bar{a}, \bar{b}) = \varphi_i(\bar{b}, \bar{b})$. By assumption, the closed formula $E(\bar{x}, \bar{z})$ satisfies $\llbracket E(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim = \llbracket E(\bar{x}, \bar{z}) \rrbracket_{T_0}$, so let $F = f(\llbracket E(\bar{x}, \bar{z}) \rrbracket_{T_0})$. This is a closed subset of X . Since X is a compact metrizable space, we have that there is a continuous function $h : X \rightarrow [0, 1]$ whose zeroset is precisely F . Since the g_i 's are dense in the uniform norm, there is a sequence $\{g_{i(n)}\}_{n < \omega}$ that limits to h in the uniform norm. This implies that $h(f(\text{tp}(\bar{a}, \bar{b}))) = h(f(\text{tp}(\bar{b}, \bar{b}))) = 0$. Therefore $f(\text{tp}(\bar{a}, \bar{b})) \in F$ and we have that $\text{tp}(\bar{a}, \bar{b}) \in \llbracket E(\bar{x}, \bar{z}) \rrbracket_{T_0}^\sim = \llbracket E(\bar{x}, \bar{z}) \rrbracket_{T_0}$, therefore $\mathfrak{M} \models E(\bar{a}, \bar{b})$, and we have the required logical equivalence. \square

The following lemma was originally shown in [BY10a].

Lemma 3.4.5. *Let $F(\bar{x})$ be a closed type-set formula, and let $\rho(\bar{x}, \bar{y})$ be a real formula with \bar{x} and \bar{y} at most countable tuples of variables of the same sorts. Suppose that for any \mathfrak{M} and $\bar{a}, \bar{b}, \bar{c}$ such that $\mathfrak{M} \models F(\bar{a}) \wedge F(\bar{b}) \wedge F(\bar{c})$, $\mathfrak{M} \models \rho(\bar{a}, \bar{a}) = 0 \wedge \rho(\bar{a}, \bar{b}) = \rho(\bar{b}, \bar{a}) \geq 0 \wedge \rho(\bar{a}, \bar{c}) \leq \rho(\bar{a}, \bar{b}) + \rho(\bar{b}, \bar{c})$ (i.e. ρ is a pseudo-metric on realizations of F), then there is a real formula $\rho'(\bar{x}, \bar{y})$ such that $\rho'(\bar{x}, \bar{y})$ is a pseudo-metric and $\rho(\bar{x}, \bar{y})$ and $\rho'(\bar{x}, \bar{y})$ are uniformly equivalent on realizations of $F(\bar{x})$, i.e. for any $\varepsilon > 0$ there exists a $\delta > 0$ such that if \bar{a}, \bar{b} are realizations of $F(\bar{x})$ and $\rho(\bar{a}, \bar{b}) < \delta$, then $\rho'(\bar{a}, \bar{b}) < \varepsilon$, and vice versa.*

Proof. Let $U_0(\bar{x}) = \top$. Note that $\llbracket F(\bar{x}) \rrbracket \subseteq \llbracket U_0(\bar{x}) \rrbracket$. For each $i < \omega$, by compactness, there is a closed formula $G_i(\bar{x})$ and an open formula $U_{i+1}(\bar{x})$ such that $\llbracket U_i(\bar{x}) \rrbracket \supseteq \llbracket G_i(\bar{x}) \rrbracket \supseteq \llbracket U_{i+1}(\bar{x}) \rrbracket \supseteq \llbracket F(\bar{x}) \rrbracket$ and such that if $\bar{a}, \bar{b}, \bar{c}$ satisfy $U_{i+1}(\bar{x})$, then $\rho(\bar{a}, \bar{a}) < 2^{-i}$, $\rho(\bar{a}, \bar{b}) > -2^{-i}$, $|\rho(\bar{a}, \bar{b}) - \rho(\bar{b}, \bar{a})| < 2^{-i}$, and $\rho(\bar{a}, \bar{c}) - \rho(\bar{a}, \bar{b}) - \rho(\bar{b}, \bar{c}) < 2^{-i}$. Therefore $G(\bar{x}) = \bigwedge_{i < \omega} G_i(\bar{x})$ is the required closed formula.

We now have that

$$E(\bar{x}, \bar{y}) = (\bar{x} = \bar{y} \vee (G(\bar{x}) \wedge G(\bar{y}) \wedge \rho(\bar{x}, \bar{y}) = 0))$$

is a closed formula that is an equivalence relation. Therefore, by Lemma 3.4.4, we can find a real formula $\rho'(\bar{x}, \bar{y})$ such that $\rho'(\bar{x}, \bar{y})$ is a pseudo-metric and $E(\bar{x}, \bar{y}) \equiv (\rho'(\bar{x}, \bar{y}) = 0)$. By compactness, $\rho(\bar{x}, \bar{y})$ and $\rho'(\bar{x}, \bar{y})$ must be uniformly equivalent for realizations of $G(\bar{x})$, and therefore for realizations of $F(\bar{x})$ as well. \square

Remark 3.4.6. In [BY10a], Ben Yaacov mentions that the problem of whether or not Lemma 3.4.5 holds with $\rho(\bar{x}, \bar{y}) = \rho'(\bar{x}, \bar{y})$ for realizations of $F(\bar{x})$ is open. As far as I know it is still.

Proposition 3.4.7. *For any theory T , any sort O , any quotient $R = O/\rho$ with quotient map $q : O \rightarrow R$, any tuple of variables \bar{x} of sort O , and tuple \bar{y} of the same length of tuples of sort R , the natural induced map $S_{\bar{x}}(T) \rightarrow S_{\bar{y}}(T)$ induced by $\text{tp}(\bar{a}) \mapsto \text{tp}(q(\bar{a}))$ is open and continuous.*

Proof. It's enough to show this for open formulas. Let $U(\bar{x})$ be an open formula logically equivalent to $(\varphi(\bar{x}) > 0)$. Just as in the proof of Lemma 3.4.4, for any $k < \omega$ there is an $\varepsilon_k > 0$ such that $\llbracket \exists \bar{x}' \varphi(\bar{x}') > 2^{-k} \wedge \rho(\bar{x}, \bar{x}') < \varepsilon_k \rrbracket_T \subseteq q^{-1}(q(U))$. So again by the same reasoning we have

$$\llbracket \bigvee_{k < \omega} \exists \bar{x}' \varphi(\bar{x}') > 2^{-k} \wedge \rho(\bar{x}, \bar{x}') < \varepsilon_k \rrbracket_T = q^{-1}(q(U)).$$

Therefore, $\llbracket \bigvee_{k < \omega} \exists \bar{x} \varphi(\bar{x}) > 2^{-k} \wedge \rho(q(\bar{x}), \bar{y}) < \varepsilon_K \rrbracket_{T, \bar{y}} = q(U)$, and we have that q is an open map.

To see that these maps are continuous, note that for any closed formula $F(\bar{y})$, $\llbracket \exists \bar{y} F(\bar{y}) \wedge q(\bar{x}) = \bar{y} \rrbracket$ is the preimage of $\llbracket F(\bar{y}) \rrbracket$ under the induced map, so there is a basis of closed sets in which the preimage of every closed set is closed, and it follows that the induced function is continuous. \square

Proposition 3.4.8 (Limited Unions of Chains of Sorts). *For any theory T , if $\{O_i\}_{i < \omega}$ is a sequence of imaginary sorts over T , $\{f_i\}_{i < \omega}$ is a sequence of definable functions over T with $f_i : O_i \rightarrow O_{i+1}$, and $\{r_i\}_{i < \omega}$ is a sequence of positive real numbers with $\sum_{i < \omega} r_i < \infty$, such that for each $i < \omega$, $T \models (\forall xy:O_i) dxy = d(f_i(x), f_i(y))$ and $T \models (\forall x:O_{i+1})(\exists y:O_i) d(x, f(y)) \leq r_i$, then there exists an imaginary sort S with definable functions $g_i : O_i \rightarrow S$ such that for each $i < \omega$,*

$$\begin{aligned} T &\models (\forall x:O_i) g_i(x) = g_{i+1}(f_i(x)), \\ T &\models (\forall xy:O_i) d(x, y) = d(g_i(x), g_i(y)), \text{ and} \\ T &\models (\forall x:S)(\exists y:O_i) d(x, g_i(y)) \leq \sum_{j \geq i} r_j. \end{aligned}$$

So in particular, in every model $\mathfrak{M} \models T$, $S(\mathfrak{M})$ is the metric closure of the union of the isometric images of $O_i(\mathfrak{M})$ under the maps g_i , and $R(\mathfrak{M})$ is isometric to the metric closure of the union of $\{O_i\}_{i < \omega}$ interpreted as a chain with the maps f_i .

Proof. Consider the imaginary sort $\prod_{i < \omega} O_i$. Let

$$C(\bar{x}) = \bigwedge_{i < \omega} d(f_i(x_i), x_{i+1}) \leq r_i \text{ and}$$

$$\psi(\bar{x}, \bar{y}) = d(x_0, y_0) + \sum_{i < \omega} [d(x_{i+1}, y_{i+1}) - d(x_i, y_i)]_{-r_i}^{r_i}.$$

To unpack, $C(\bar{x})$ says that \bar{x} is (morally speaking) a Cauchy sequence converging at a certain prescribed rate. By construction, we have that $\psi(\bar{x}, \bar{y})$ is a pseudo-metric over T for realizations of $C(\bar{x})$. By Lemma 3.4.5 there is a pseudo-metric $\rho(\bar{x}, \bar{y})$ on $\prod_{i < \omega} O_i$ that is uniformly equivalent to $\psi(\bar{x}, \bar{y})$ on realizations of $C(\bar{x})$. Let $R = \prod_{i < \omega} O_i / \rho$ with quotient map q . Let $q_1 : S_{\bar{x}}(T) \rightarrow S_z(T)$ and $q_2 : S_{\bar{x}\bar{y}}(T) \rightarrow S_{zw}(T)$ be the induced map on type spaces, where z and w are variables of sort R .

For each $i < \omega$, consider the closed formula

$$\begin{aligned} F_i(x_i, z) &= (\exists x_0 \dots x_{i-1}) C(x_0, \dots, x_{i-1}, x_i, f_i(x_i), f_{i+1}(f_i(x_i)), \dots) \\ &\quad \wedge q(x_0, \dots, x_{i-1}, x_i, f_i(x_i), f_{i+1}(f_i(x_i)), \dots) = z. \end{aligned}$$

By assumption, for any $\mathfrak{M} \models T$ and $a_i \in O_i(\mathfrak{M})$, there is an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ in which there are a_0, \dots, a_{i-1} with $a_j \in O_j(\mathfrak{M})$ and $d(f_j(a_j), a_{j+1}) = r_j$ for each $j < i$. Therefore $b = q(a_0, \dots, a_i, f_i(a_i), f_{i+1}(f_i(a_i)), \dots)$ is a realization of $C(\bar{x})$. By construction, b is unique among the elements of $R(\mathfrak{M})$. Therefore, by Proposition 2.3.79, $F_i(x_i, y)$ is the graph of a definable function, $h_i : O_i \rightarrow R$, over T . By Proposition 2.3.78, for each $i < \omega$, the metric closure of the image of O_i in R under h_i is definable. Let $D_i(z)$ be that image. Note that by construction, $D_i(\mathfrak{M}) \subseteq D_{i+1}(\mathfrak{M})$ for any $\mathfrak{M} \models T$. Also, since $\rho(\bar{x}, \bar{y})$ is uniformly equivalent to $\psi(\bar{x}, \bar{y})$, we have that in each D_i , $\rho(\bar{x}, \bar{y})$ is uniformly equivalent to d_{O_i} , in the sense that for every $\varepsilon > 0$ there is a $\delta > 0$ such that $T \models \forall x_i y_i d(x_i, y_i) < \delta \rightarrow d(h_i(x_i), h_i(y_i)) \leq \varepsilon$ and $T \models \forall x_i y_i d(h_i(x_i), h_i(y_i)) < \delta \rightarrow d(x_i, y_i) \leq \varepsilon$. Therefore in particular, there is some sequence of real numbers s_i

with $\lim_{i < \omega} s_i = 0$ such that for each i , $T \models d_H(D_i, D_{i+1}) \leq s_i$. By Corollary 2.3.30, the metric closure of the union of the chain $\{D_i(z)\}_{i < \omega}$ is definable. Let this set be $D(z)$. Note that it must be the case that $T, D(q(\bar{x})) \models C(\bar{x})$. Take the imaginary sort $D(R)$, and let $e : D(R) \rightarrow R$ be the inclusion map. For each $i < \omega$, we have that for any $a_i \in O_i(\mathfrak{M})$ with $\mathfrak{M} \models T$, there is a unique $c_i \in D(R)$ such that $h_i(a_i) = e(c_i)$. Therefore, by Proposition 2.3.79, there is a definable function, $h'_i : O_i \rightarrow D(R)$, over T , such that for any $a_i \in O_i(\mathfrak{M})$, $h_i(a_i) = e(h'_i(a_i))$.

By construction, for any $p \in S_{zw}(T)$, there is a unique $r \in \mathbb{R}$ such that if $\text{tp}(q(\bar{a}), q(\bar{b})) = p$ (for $\bar{a}, \bar{b} \in \mathfrak{M} \models T \wedge C(\bar{a}) \wedge C(\bar{b})$), then $\psi(\bar{a}, \bar{b}) = r$. Therefore we can define a function, ψ^* on $q_2(\llbracket C(\bar{x}) \wedge C(\bar{y}) \rrbracket)$ by this. By Proposition 3.4.7, for any r , the sets $q_2(\llbracket \psi(\bar{x}, \bar{y}) < r \rrbracket_T)$ and $q_2(\llbracket \psi(\bar{x}, \bar{y}) > r \rrbracket_T)$ are open, therefore the function ψ^* is continuous on ψ^* on $q_2(\llbracket C(\bar{x}) \wedge C(\bar{y}) \rrbracket)$, and by the Tietze extension theorem and Proposition 1.5.9 we can find a formula $\psi^{**}(z, w)$ extending ψ^* to a continuous function on all of $S_{zw}(T)$. Now we have that $\rho'(u, v) = \psi^{**}(e(u), e(v))$ is a pseudo-metric on $D(R)$. Moreover, it is uniformly equivalent to the metric on $D(R)$. So now consider the imaginary sort $S = D(R)/\rho'$, and let $q' : D(R) \rightarrow S$ be the quotient map (although this isn't so much a quotient as just an adjustment of the metric). Now finally $g_i = q' \circ h'_i : O_i \rightarrow S$ are the required maps. \square

Corollary 3.4.9 (Countable Disjoint Union of Sorts). *For any countable sequence of sorts $\{O_i\}_{i < \omega}$ in a theory T (such that for each $i < \omega$, $T \models (\exists x : O_i)x = x$) there is an imaginary sort R , a formula $\varphi(x)$ with x a variable of sort R , definable maps $f_i : O_i \rightarrow R$,*

and a definable zero-ary function ∞ of sort R , such that for each $i < j < \omega$,

$$T \models (\forall xy:O_i)d(f_i(x), f_i(y)) = \frac{d(x, y)}{2^i \text{db}(O_i)},$$

$$T \models (\forall x:O_i)(\forall y:O_j)d(f_i(x), f_j(y)) = |2^{-i} - 2^{-j}|,$$

$$T \models (\forall x:O_i)d(f_i(x), \infty) = 2^{-i},$$

$$T \models (\forall x:O_i)\varphi(f_i(x)) = 2^{-i},$$

$$T \models \varphi(\infty) = 0,$$

and such that in any $\mathfrak{M} \models T$, $R(\mathfrak{M}) = \{\infty^{\mathfrak{M}}\} \cup \bigcup_{i < \omega} f_i(O_i(\mathfrak{M}))$.

Proof. By passing to quotients by scaled metrics, we may assume that $\text{db}(O_i) = 2^{-i}$ for each $i < \omega$. By using Proposition 3.4.3, we can construct finite disjoint unions of $\{O_i\}_{i < n}$ for each $n < \omega$ with metrics that satisfy the statements in the proposition for small enough i . (It is not hard to check that the conditions in the proposition are consistent with the axioms of a metric.) By Proposition 2.3.79 we get inclusion maps $f_i : \bigsqcup_{i < n} O_i \rightarrow \bigsqcup_{i \leq n} O_i$ and now we can apply Proposition 3.4.8 to get the required sort. The density of the images of the sorts O_i in R , together with the specification of the sort and the fact that the sorts O_i are non-empty, implies that there is always a unique point not in the image of any of the inclusion maps. This element is the unique element that satisfies $\bigwedge_{i < \omega} (\exists y:O_i)d(x, f_i(y)) = 2^{-i}$. Therefore by Proposition 2.3.79 there is a zero-ary definable function ∞ which is equal to this point in every model of T . It is not hard to check that the required properties of $\varphi(x)$ specify a unique continuous function on $S_{x:R}(T)$. Therefore by Proposition 1.5.9 the required formula exists. \square

Definition 3.4.10 (Gromov-Hausdorff Metric on Imaginaries). For any theory T and

any two (imaginary) sorts O_0 and O_1 over T , a *metric correspondence between O_0 and O_1 over T* , written $\varphi : O_0 \approx_T O_1$ is a formula $\varphi(x:O_0, y:O_1)$ such that

$$\delta_\varphi(x, y) = (1 - |\psi(x) - \psi(y)|)d(x, y) + (\psi(x) + \psi(y) - \psi(x)\psi(y))\varphi'(x, y)$$

is a pseudo-metric on $O_0 \sqcup_r O_1$ (the specific r is unimportant), where x, y are variables of sort $O_0 \sqcup_r O_1$, $\psi(x)$ is a formula that is 0 on the image of O_0 and 1 on the image of O_1 , and $\varphi'(x, y)$ is some extension of $\varphi(x, y)$ to a formula on $O_0 \sqcup_r O_1$ (the specific extension is unimportant, but note that such an extension always exists).

For any two (imaginary) sorts O_0, O_1 over T , the *definable Gromov-Hausdorff distance between O_0 and O_1 over T* , written $d_{dGH}^T(O_0, O_1)$, is

$$\inf\{r \geq 0 : (\exists \varphi : O_0 \approx_T O_1)(\forall i < 2)T \models (\forall x_i:O_i)(\exists x_{1-i}:O_{1-i})\varphi(x_0, x_1) \leq r\}. \quad \triangleleft$$

Such a function for metric spaces in general is also called a bi-Katětov function. Note that it is possible for two sorts to have $d_{dGH}^T(O_0, O_1) = 0$ without there being a definable isometry between O_0 and O_1 over T (see Counterexample C.1.4).

Corollary 3.4.11 (Uniform Limits of Sorts). *For any theory T , the collection of (imaginary) sorts over T is a complete pseudo-metric space under d_{dGH}^T .*

Proof. $d_{dGH}^T(O_0, O_0) = 0$ and $d_{dGH}^T(O_0, O_1) = d_{dGH}^T(O_1, O_0) \geq 0$ are immediate from the definition. Assume that $d_{dGH}^T(O_0, O_1) < r$ and $d_{dGH}^T(O_1, O_2) < s$. Let $\varphi : O_0 \approx_T O_1$ and $\psi : O_1 \approx_T O_2$ be witnesses of these bounds. Then $\chi(x, y) = \inf_{z:O_1} \varphi(x, z) + \psi(z, y)$ is a witness that $d_{dGH}^T(O_0, O_2) \leq r + s$. Since we can do this for any upper bounds of $d_{dGH}^T(O_0, O_1)$ and $d_{dGH}^T(O_1, O_2)$, we have that $d_{dGH}^T(O_0, O_2) \leq d_{dGH}^T(O_0, O_1) +$

$d_{dGH}^T(O_1, O_2)$, as required. Therefore d_{dGH}^T is a pseudo-metric.

For completeness, assume that $\{O_i\}_{i < \omega}$ is a Cauchy sequence in d_{dGH}^T . By thinning the sequence, we may assume that $d_{dGH}^T(O_i, O_{i+1}) < 2^{-i}$, for each $i < \omega$. For each $i < \omega$, let $\varphi_i(x_i:O_i, x_{i+1}:O_{i+1})$ be a witness that $d_{dGH}^T(O_i, O_{i+1}) < 2^{-i}$. For each $n < \omega$, we can construct a sort that is a quotient of $\bigsqcup_{i < n} O_i$ by a metric δ equal to the normal metric on O_i for each $i < n$ and such that for $i < j < n$, $x_i \in O_i$, and $x_j \in O_j$, $\delta(x_i, x_j) = \inf_{x_{i+1}} \dots \inf_{x_{j-1}} \varphi_i(x_i, x_{i+1}) + \dots + \varphi_{j-1}(x_{j-1}, x_j)$. By Proposition 2.3.79 we get isometric inclusion maps $f_i : \bigsqcup_{i < n} O_i \rightarrow \bigsqcup_{i \leq n} O_i$, so we can apply Proposition 3.4.8 to get a sort R that is $\bigsqcup_{i < \omega} O_i$ modulo the appropriate metric. For each $i < \omega$, we have an isometric inclusion $g_i : O_i \rightarrow R$. By construction, the sequence of definable sets $\{g_i(O_i)\}_{i < \omega}$ are a Hausdorff-Cauchy sequence in R . Therefore by Proposition 2.3.29, the limit is a definable set $D(x)$ as well. Taking the definable set sort $D(R)$, we have that for each $i < \omega$, $d_R(g_i(x), h(y))$, where $h : D(R) \rightarrow R$ is the inclusion map, is a metric correspondence between O_i and $D(R)$ witnessing that $d_{dGH}^T(O_i, D(R)) \leq \sum_{i \leq j < \omega} 2^{-j} = 2^{-i+1}$. Therefore $D(R)$ is a limit of the sequence $\{O_i\}_{i < \omega}$ under d_{dGH}^T . \square

3.5 Compact Parameter Sorts and Almost Uniformity

Definition 3.5.1. For any theory T , a *compact sort over T* is a (possibly imaginary) sort K such that for any $\mathfrak{M} \models T$, $K(\mathfrak{M}) \subseteq \text{acl}(\emptyset)$. A *compact parameter sort over T* is a (possibly imaginary) sort K such that for any $\mathfrak{M} \models T$, $K(\mathfrak{M}) \subseteq \text{dcl}(\emptyset)$. \triangleleft

By Proposition 2.3.63, in both cases this implies that when \mathfrak{M} is a model of T , $O(\mathfrak{M})$ is metrically compact. Just as how finite imaginary sorts are important in discrete stability theory (e.g. for the notion of strong types), compact imaginary sorts are important

in continuous stability theory, but here we are more concerned with compact parameter sorts.

Commonly in discrete logic, a formula or definable function is constructed by patching together a finite collection of definable formulas, where the collection is guaranteed to be finite by a compactness argument. A good example of this is in the definable closure. In discrete logic, if we have a formula $\varphi(\bar{x}, y)$ and we know that for any tuple \bar{a} , if $\exists y\varphi(\bar{a}, y)$, then there is $b \in \text{dcl}(\bar{a})$ such that $\varphi(\bar{a}, b)$, then we can show that this is witnessed by a definable Skolem function. Compactness gives us that there's a finite list of definable partial functions f_0, \dots, f_{n-1} such that if $\exists y\varphi(\bar{a}, y)$, then for some $i < n$, $\varphi(\bar{a}, f_i(\bar{a}))$ holds. Then we can patch these together with

$$F(\bar{x}) = \begin{cases} f_i(\bar{x}) & i \text{ is the smallest such that } \varphi(\bar{x}, f_i(\bar{x})) \text{ holds} \\ x_0 & \text{otherwise} \end{cases}. \quad (*)$$

With this, $F(\bar{x})$ is a definable total function and we have that $\forall \bar{x}(\exists y\varphi(\bar{x}, y)) \rightarrow \varphi(\bar{x}, F(\bar{x}))$ holds.

These sorts of discrete shenanigans simply do not fly in continuous logic. Firstly, we are not guaranteed a clean clopen separation between the case in which something we want is happening and the case in which it is not, even if it is first-order. Secondly, the formula corresponding to our desired situation (analogous to $\exists y\varphi(\bar{x}, y)$) will typically be a closed formula, so compactness does not apply. We do have an extremely weak set theoretic form of compactness, specifically $\llbracket F \rrbracket$ for a closed formula F is a closed G_δ set. By results of [Usu19], any cover of a compact Hausdorff space by G_δ sets has a subcover of size at most the smallest ω_1 -strongly compact cardinal (and furthermore this property

characterizes the first ω_1 -strongly compact cardinal which consistently does not exist). Needless to say, this bound is slightly too large for comfort.

An abstract perspective on the example (*) is to think of $F(\bar{x})$ as being given by a formula $\psi(\bar{x}, i, y)$ where i is a variable in an imaginary sort with precisely $n+1$ elements, which are each distinguishable. If we take each partial function f_i to be given by the formula $\chi_i(\bar{x}, y)$ then we can construct a formula

$$\psi(\bar{x}, i, y) = \bigvee_{j < n} i = j \wedge \chi_j(\bar{x}, y) \vee (i = n + 1 \wedge x_0 = y), \quad (\heartsuit)$$

and now we have that

$$\forall \bar{x} \exists i \exists! y (\exists z \varphi(\bar{x}, z)) \rightarrow \psi(\bar{x}, i, y) \wedge \varphi(\bar{x}, y).$$

The point of this odd way of looking at (*) is that (\heartsuit) does generalize to continuous logic, with the familiar stipulation that the imaginary sort of the parameter i is in general metrically compact, rather than actually finite. The only barrier to stitching together a function like F is that while there is a way of picking a canonical witness, there is no continuous way.

A simple, but important, observation is that in compact sorts, weak existence implies existence.

Proposition 3.5.2. *If y is a variable in a compact sort and a structure \mathfrak{M} has $\mathfrak{M} \models \exists y F(\bar{a}, y)$ for some closed formula F and some tuple of parameters $\bar{a} \in \mathfrak{M}$, then there exists $b \in \mathfrak{M}$ such that $\mathfrak{M} \models F(\bar{a}, b)$.*

Proof. $\exists y F(\bar{x}, y) \equiv (\sup_y \varphi(\bar{x}, y) \geq 0)$ for some real formula φ . Continuous function

attain their suprema on compact spaces, so we have that the required witness exists. \square

Notation 3.5.3. If we are working in the context of a theory T over which some sort K is compact, we may write $(\exists x:K)$ or just $\exists x$ in closed formulas for variables from the sort K , instead of $(\exists x:K)$ or $\exists x$. \triangleleft

The two kinds of ‘almost uniformity’ that will occur in this thesis are almost uniformly definable sets and almost functions. Really we only care about this definition for compact parameter sorts, but only one of the results in this subsection actually depends on this extra assumption.

Definition 3.5.4. Fix a theory T , partial type $\Sigma(\bar{x}) \supseteq T$, compact sort K over T , and variable r in sort K .

- A real formula $\varphi(\bar{v}, \bar{x}, r)$ gives an *almost \bar{x} -uniformly definable set over $\Sigma(\bar{x})$ with parameter r* if $\Sigma(\bar{x}) \models \exists r \text{DEF} \bar{v} \varphi(\bar{v}, \bar{x}, r)$. We say φ gives an *almost \bar{x} -uniformly definable set over T with parameter r* if $T \models \forall \bar{x} \exists r \text{DEF} \bar{v} \varphi(\bar{v}, \bar{x}, r)$.
- A real formula $\varphi(\bar{x}, y, r)$ gives a *definable partial almost function over $\Sigma(\bar{x})$ with parameter r* if $\Sigma(\bar{x}) \models \exists r \text{FUN} y \varphi(\bar{x}, y, r)$. We say φ gives a *definable almost function over T with parameter r* if $T \models \forall \bar{x} \exists r \text{FUN} y \varphi(\bar{x}, y, r)$. \triangleleft

Note that a definable partial almost function on \bar{x} is actually just a special kind of definable partial function on $\bar{x}r$.

These concepts are only non-trivial if there are parameters \bar{a} for which multiple parameters from the compact parameter sort work.

Lemma 3.5.5. *For any partial type $\Sigma(\bar{x})$, any formula $\varphi(\bar{v}, \bar{x}, r)$ that gives an almost \bar{x} -uniformly definable set over $\Sigma(\bar{x})$, and any parameter $\bar{a} \in \mathfrak{M} \models \Sigma(\bar{a})$, the formula $F(\bar{v}, \bar{a}) = \exists r(\varphi(\bar{v}, \bar{a}, r) = 0) \wedge \text{DEF}_{\bar{u}}\varphi(\bar{u}, \bar{a}, r)$ gives a definable set.*

Proof. By converting parameters to constants, it is sufficient to show that for any complete theory T with compact sort K and variable r of sort K and any formula $\varphi(\bar{v}, r)$, the closed formula $F(\bar{v}) = \exists r(\varphi(\bar{v}, r) = 0) \wedge \text{DEF}_{\bar{u}}\varphi(\bar{u}, r)$ gives a definable set.

The closed formula $G(r) = \text{DEF}_{\bar{u}}\varphi(\bar{u}, r)$ gives a definable set by Proposition 2.3.65, so we have that $\psi(\bar{v}) = \inf_{r \in G} \varphi(\bar{v}, r)$ is a distance predicate such that $F(\bar{v}) \equiv (\psi(\bar{v}) = 0)$, hence $F(\bar{v})$ gives a definable set. \square

Specifically, in the previous proof the definable set F has $F(\mathfrak{M}) = \bigcup\{(\varphi(\mathfrak{M}, r) = 0) : r \in K(\mathfrak{M}), \mathfrak{M} \models \text{DEF}_{\bar{v}}\varphi(\bar{v}, r)\}$. Note that it actually is the union, not the metric closure of the union, as the union of ‘compactly many’ closed sets is closed (Fact A.2.15).

Proposition 3.5.6. *Fix a theory T and a partial type $\Sigma(\bar{x}) \supseteq T$.*

(i) *Suppose that $\varphi(\bar{v}, \bar{x}, r)$ gives an almost \bar{x} -uniformly definable set over $\Sigma(\bar{x})$ with parameter r of sort K and furthermore suppose that for any $\bar{a} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, there is a unique $b \in K(\mathfrak{M})$ such that $\mathfrak{M} \models \text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{a}, b)$, then there is a \bar{x} -uniformly definable set over $\Sigma(\bar{x})$, $D(\bar{v}, \bar{x})$ such that for any $\bar{a}b \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$ and $\mathfrak{M} \models \text{DEF}_{\bar{v}}\varphi(\bar{v}, \bar{a}, b)$, $(\varphi(\mathfrak{M}, \bar{a}, b) = 0)$ and $D(\mathfrak{M}, \bar{a})$ are the same set.*

(ii) *Suppose that $\varphi(\bar{x}, y, r)$ gives a definable partial almost function over $\Sigma(\bar{x})$ and furthermore suppose that for any $\bar{a} \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$, there is a unique $b \in K(\mathfrak{M})$ such that $\mathfrak{M} \models \text{FUN}_y\varphi(\bar{a}, y, b)$, then there is a definable partial function*

f such that $\Sigma(\bar{x}) \models f(\bar{x})\downarrow$ and for any $\bar{a}b \in \mathfrak{M}$ such that $\mathfrak{M} \models \Sigma(\bar{a})$ and $\mathfrak{M} \models \text{FUN}y\varphi(\bar{a}, y, b)$, we have $\mathfrak{M} \models \varphi(\bar{a}, f(\bar{a})\downarrow, b) = 0$.

Proof. Part (ii) follows easily from part (i), so we just need to show part (i).

Consider $S_{\bar{v}\bar{x}r}(T)$. Let $f : S_{\bar{v}\bar{x}r}(T) \rightarrow S_{\bar{v}\bar{x}}(T)$ be the natural projection map. Consider $F(\bar{v}, \bar{x}, r) = \text{DEF}\bar{u}\varphi(\bar{u}, \bar{x}, r) \wedge (\varphi(\bar{v}, \bar{x}, r) = 0)$. By assumption we have that $f \upharpoonright \llbracket F \rrbracket$ is a continuous bijection between compact Hausdorff spaces. Therefore by Fact A.2.11 $f \upharpoonright \llbracket F \rrbracket$ is a homeomorphism. Let g be its inverse, and let $\psi(\bar{v}, \bar{x})$ be a formula corresponding to the continuous function $\varphi(g(\text{tp}(\bar{v}\bar{x})))$ on $S_{\bar{v}\bar{x}}(T)$. By construction, $\psi(\bar{v}, \bar{x})$ is the required distance predicate of the required \bar{x} -uniformly definable set. \square

Remark 3.5.7. Another application of almost uniformity, which we did not have the chance to develop in this thesis, is to the issue of Skolemization in continuous logic. There are several different ways of generalizing the notion of Skolemization to continuous logic, but a seemingly natural one is this: A theory T is *weakly Skolemized* if for any set of parameters A , $\text{dcl}A$ is a model of T . It is not hard to show in discrete logic that this is equivalent to ordinary Skolemization.

It is possible to prove, using an argument similar to the one used in the proof of Theorem 2.4.14, that a continuous theory is weakly Skolemized if and only if for each open formula $U(\bar{x}, y)$ such that $T \models \forall \bar{x} \exists y U(\bar{x}, y)$, it has definable almost Skolem function f , i.e. a definable almost function $f_z(\bar{x})$, with z a variable in a compact parameter sort \mathbf{K} , such that for every \bar{a} , there exists a $b \in \mathbf{K}$ such that $U(\bar{a}, f_b(\bar{a})\downarrow)$ holds.

3.6 Interpretations

Definition 3.6.1. Given an \mathcal{L}_0 -theory T_0 and an \mathcal{L}_1 -theory T_1 , an *interpretation of T_1 in T_0* , \mathcal{I} , consists of

- a sort $\mathcal{I}(O) \in \mathcal{S}(\mathcal{L}^{eq}(T_0))$ for each sort $O \in \mathcal{S}(\mathcal{L}_1)$,
- an $\mathcal{L}^{eq}(T_0)$ -formula $\mathcal{I}(P)(\bar{x})$, where $S(x_i) = \mathcal{I}(a(P)(i))$ and $I(\mathcal{I}(P)) \subseteq I(P)$, for each predicate symbol $P \in \mathcal{P}(\mathcal{L}_1)$, with $\mathcal{I}(d_O) = d_{\mathcal{I}(O)}(x, y)$, and
- an $\mathcal{L}^{eq}(T_0)$ -formula $\mathcal{I}(f)(y, \bar{x})$, where $S(x_i) = \mathcal{I}(a(P)(i))$ and $S(y) = \mathcal{I}(S(f))$ and $I(\mathcal{I}(P)) \subseteq I(d_{S(y)})$, for each function symbol $f \in \mathcal{F}(\mathcal{L}_1)$,

such that if we extend the map to arbitrary unnested \mathcal{L}_1 -formulas by

- $\mathcal{I}(P\bar{x}) = \mathcal{I}(P)(\mathcal{I}(x_0), \mathcal{I}(x_1), \dots)$ (note that P may be d_O), where for each variable symbol x of sort $O \in \mathcal{S}(\mathcal{L}_1)$, $\mathcal{I}(x)$ is a corresponding variable symbol of sort $\mathcal{I}(O)$,
- $\mathcal{I}(dyf\bar{x}) = \mathcal{I}(f)(\mathcal{I}(y), \mathcal{I}(x_0), \mathcal{I}(x_1), \dots)$,
- $\mathcal{I}(F\bar{\varphi}) = F(\mathcal{I}(\varphi_0), \mathcal{I}(\varphi_1), \dots)$, for any connective F ,
- $\mathcal{I}(Qv\varphi) = Q.\mathcal{I}(v).\mathcal{I}(\varphi)$, for $Q \in \{\text{sup}, \text{inf}\}$, and
- $\mathcal{I}(\varphi \square \psi) = \mathcal{I}(\varphi) \square \mathcal{I}(\psi)$, for any $\square \in \{<, >, \leq, \geq, =, \neq\}$,

then we have that for any closed \mathcal{L}_1 -sentence G , if $T_0 \models G$, then $T_1^{eq} \models \mathcal{I}(G)$.

We say that T_1 is *interpretable in T_0* if there exists an interpretation of T_1 in T_0 .

We say that T_0 and T_1 are *mutually interpretable* if T_0 is interpretable in T_1 and T_1 is interpretable in T_0 . ◁

Proposition 3.6.2. *For any \mathcal{L}_0 -theory T_0 and \mathcal{L}_1 -theory T_1 , any interpretation \mathcal{I} of T_1 in T_0 extends to an interpretation \mathcal{I}^{eq} of T_1^{eq} in T_0 .*

Proof. Define $\mathcal{I}^{eq}(O)$ iteratively in the obvious way (note that if D is a definable set over T_1^{eq} , then by Lemma B.1.1 $\mathcal{I}^{eq}(D)$ will be a definable set over T_0^{eq} , so the obvious definition is well defined). A straightforward but tedious induction shows that \mathcal{I}^{eq} is an interpretation of T_1^{eq} in T_0 . \square

Definition 3.6.3. Given an \mathcal{L}_0 -theory T_0 and an \mathcal{L}_1 -theory T_1 , a *bi-interpretation* between T_0 and T_1 is a pair $(\mathcal{I}_0, \mathcal{I}_1)$ such that for both $i < 2$, \mathcal{I}_i is an interpretation of T_{1-i} in T_i and for each sort $O \in \mathcal{S}(\mathcal{L}_i)$, there is a definable isometry $f_O : O \rightarrow \mathcal{I}_{1-i}(\mathcal{I}_i(O))$ with definable inverse over T_i^{eq} such that for any definable function $g : O \rightarrow O'$, $T_i^{eq} \models (\forall x:O) f_{O'}(g(x)) = \mathcal{I}_{1-i}(\mathcal{I}_i(g))(f_O(x))$.

We say that T_0 and T_1 are *bi-interpretable* if there is a bi-interpretation between them. \triangleleft

There are two special kinds of bi-interpretation that we will commonly use.

Definition 3.6.4. An \mathcal{L}_0 -theory T_0 and an \mathcal{L}_1 -theory T_1 are *inter-definable* if there is a bi-interpretation $(\mathcal{I}_0, \mathcal{I}_1)$ between them such that as \mathcal{I}_0 maps the sorts of \mathcal{L}_0 bijectively onto the sorts of \mathcal{L}_1 , \mathcal{I}_1 is its inverse, and for both $i < 2$, for every \mathcal{L}_i -formula φ , $\mathcal{I}_{1-i}(\mathcal{I}_i(\varphi))$ is logically equivalent to φ . \triangleleft

Usually we will consider pairs of inter-definable theories where the two languages have the same set of sorts (such as a theory and a reduct of that theory with the same sorts).

The second special kind of bi-interpretation is the one guaranteed by the following proposition.

Proposition 3.6.5. *Fix a signature \mathcal{L} and an \mathcal{L} -theory T . If $\{H_i\}_{i \in I}$ and $\{O_j\}_{j \in J}$ are two families of sorts in \mathcal{L} such that for any $\mathfrak{M} \models T$, $\bigcup_{i \in I} H_i(\mathfrak{M}) \subseteq \text{dcl}\left(\bigcup_{j \in J} O_j(\mathfrak{M})\right)$ and $\bigcup_{i \in I} O_i(\mathfrak{M}) \subseteq \text{dcl}\left(\bigcup_{j \in J} H_j(\mathfrak{M})\right)$, then $T_H = T^{\text{Mor}} \upharpoonright \{H_i\}_{i \in I}$ and $T_O = T^{\text{Mor}} \upharpoonright \{O_j\}_{j \in J}$ are bi-interpretable.*

Proof. After passing to T^{Mor} , we may assume that $\{H_i\}_{i \in I}$ and $\{O_j\}_{j \in J}$ exhaust the sorts of \mathcal{L} by passing to the reduct containing only those sorts.

By Theorem 3.2.18, we have that every H_i is equivalent to an imaginary of T_O and that every O_j is equivalent to an imaginary of T_H (since the theory is Morleyized). This gives interpretations $\mathcal{I}_0 : T_H \rightarrow T_O$ and $\mathcal{I}_1 : T_O \rightarrow T_H$. The required maps witnessing that $(\mathcal{I}_0, \mathcal{I}_1)$ is a bi-interpretation come from the fact that T_H and T_O are themselves Morleyized. □

Chapter 4

Minimality and Categoricity

The classical Baldwin-Lachlan characterization of uncountably categorical theories gives detailed structural information about such theories. In particular each such theory has a strongly minimal set definable over its prime model and every model of the theory is ‘controlled’ by any such strongly minimal set. This structural picture fails in continuous logic in the context of inseparably categorical theories. In particular the theory of infinite dimensional Hilbert spaces (IHS) does not contain anything resembling a strongly minimal set.

Nevertheless, there is a meaningful notion of strongly minimal sets in continuous logic which is a non-trivial generalization of the notion in discrete logic. In this chapter we will examine this notion and related machinery.

Section 4.3 deals with a few mild variations on the notion of Vaughtian pairs in continuous logic. It is shown that none of them can occur in an inseparably categorical theory.

Section 4.1 deals with (strongly) minimal sets in continuous logic. The theory of $(\mathbb{R}, +)$ is shown to be strongly minimal and to not interpret any infinite discrete structures. The general development of strongly minimal sets in continuous logic resembles their development in discrete logic, so much so that some of the proofs here follow Tent and Ziegler [TZ12] closely, but there are a few important differences:

- The definition needs to be slightly different from the most obvious direct translation—“Every definable subset is either compact or co-pre-compact.”—which does not work. There are two problems with this definition:
 - It might be that the set in question has many distinct non-algebraic types but not enough definable subsets to distinguish between them.
 - Under the correct definition, a strongly minimal set can have a definable subset that is neither compact nor co-pre-compact.
- It’s not clear whether or not the property of having no Vaughtian pairs is sufficient to ensure that any minimal set is strongly minimal in an arbitrary theory. There are two stronger hypotheses, both satisfied by \aleph_1 -categorical theories, both of which are strong enough to ensure that any minimal set is strongly minimal. In particular if T is dictionary and has no Vaughtian pairs, or if T has no open-in-definable Vaughtian pairs, then any minimal set is strongly minimal.
- The proof that minimal sets over \aleph_0 -saturated structures are strongly minimal still works, but in continuous logic sometimes we need to work with approximately \aleph_0 -saturated structures. Fortunately minimal sets over approximately \aleph_0 -saturated structures are strongly minimal.
- In general, algebraic closure in continuous logic does not have finite character, only countable character and approximate finite character, but in strongly minimal sets algebraic closure does have finite character. This is notable because there are \aleph_1 -categorical theories in which acl is a pregeometry that does not have finite character, namely IHS.

- In discrete logic if p is a strongly minimal type based on some set A , then a corresponding strongly minimal set is definable over A . In continuous logic if A is not a model we need to invoke an extra assumption and work with a weaker notion than strongly minimal set.
 - If T is dictionaric then for any strongly minimal p based on some set A , there is an A -definable ‘approximately strongly minimal set’ corresponding to p .
 - Conversely (with no assumptions about T) if D is an approximately strongly minimal set definable over A , then it corresponds to a strongly minimal type based on A .
 - If D is an approximately strongly minimal set definable over A and $\mathfrak{M} \supseteq A$ is any model, then there is a $(D(\mathfrak{M}) \cup A)$ -definable strongly minimal set $E \subseteq D$ corresponding to the same strongly minimal type (again with no assumptions about T).
 - Every known example of an A -definable approximately (strongly) minimal set has an A -definable (strongly) minimal imaginary quotient, but it is not clear that this is always true.
- In discrete logic if $S_n(\mathfrak{M})$ is topologically scattered (if T is totally transcendental, for instance), then every non-algebraic open set contains a type that is minimal over \mathfrak{M} . In continuous logic the correct topometric analog of topologically scattered (i.e. CB-analyzable [BY08c]) is not strong enough to guarantee the existence of any minimal or strongly minimal types. In particular there is an \aleph_1 -categorical theory that does not even interpret a strongly minimal set, namely IHS.

Each of these differences with the exception of the last one are either mild technical issues or the fortuitous lack of a potential mild technical issue. The last difference is the most important one in that not every inseparably categorical theory is in some way ‘based on’ a strongly minimal set.

Section 4.4 contains a partial generalization of the Baldwin-Lachlan characterization to the context of continuous logic. A counterexample to the most direct translation is presented. Examples of inseparably categorical theories that do have strongly minimal sets in their home sort, but only over models of sufficiently high dimension, are also presented. Finally the issue of the number of separable models of an inseparably categorical theory is discussed, with some mild progress.

In another chapter (Chapter 5), we will present a couple of common conditions that ensure the presence of minimal sets in an ω -stable theory. Namely, theories with a locally compact model have minimal sets and ultrametric theories have minimal imaginaries. The relationship between ultrametric theories and theories with totally disconnected type spaces is characterized in this section. This chapter also contains a novel proof that infinite dimensional Hilbert spaces do not interpret any non-trivial locally compact theories and in particular any strongly minimal theories, a fact which is implicit in [BYB04].

4.1 (Strongly) Minimal Sets in Continuous Logic

There are existing definitions of minimal set and strongly minimal set in the literature given in [Noq17], but both of these definitions are too weak. The definition of strongly minimal set given there fails to generalize the notion of strongly minimal in discrete

logic and instead corresponds to a countably type-definable set containing a unique non-algebraic type which furthermore is minimal (i.e. a type for which every forking extension is algebraic). The definition of minimal set given there is trivial in the sense that under it every countable theory with non-compact models has a minimal zeroset over any given separable model. The definitions given here are equivalent to the definitions given in [Noq17] with the extra stipulation that the sets are definable rather than just zerosets. In particular this means that the notion of a strongly minimal *theory* given in [Noq17] is equivalent to the one given here.

Definition 4.1.1.

- (i) A non-algebraic definable set D is *minimal* (over the set A) if for each pair $F, G \subseteq D$ of disjoint A -zerosets, at most one of F or G is non-algebraic.
- (ii) A definable set is *strongly minimal* if it is minimal over every set of parameters. \triangleleft

The naïve translation of the definition of minimal set—every set $\llbracket P(\mathfrak{M}) = 0 \rrbracket \cap D(\mathfrak{M})$ is either compact or co-pre-compact (i.e. has a complement with a compact closure)—does not work:

Example 4.1.2. *A strongly minimal set D with a definable set $E \subseteq D$ that is neither compact nor co-pre-compact.*

Verification. Let \mathfrak{M} be a structure whose universe is $\omega \times S^1$, where $S^1 \subseteq \mathbb{C}$ is the unit circle with the standard Euclidean metric. Let the distance between any points in distinct circles be 1. Let D be the entire structure, and let E be the subset of $S_1(\mathfrak{M})$ given by $\{(n, e^{2\pi ki/(n+1)}) : k \leq n\}$ (where we are identifying elements of \mathfrak{M} with their types in $S_1(\mathfrak{M})$) together with the unique non-algebraic type. This set is clearly closed.

To see that it is a definable set, pick $\varepsilon > 0$, and consider $E^{<\varepsilon}$. For any $n > \frac{4\pi}{\varepsilon}$, $E^{<\varepsilon}$ contains all of the circle $\{n\} \times S^1$. There are only finitely many $n \leq \frac{4\pi}{\varepsilon}$, and on each of these $E^{<\varepsilon}$ is an open set since the logic topology agrees with the metric topology on each individual circle in \mathfrak{M} . Therefore $E^{<\varepsilon}$ is an open set, and so E is definable. \square

Another example is $(-\infty, 0] \cup \{\ln(1+n) : n < \omega\}$, which is an \mathbb{R} -definable subset of \mathbb{R} (which will be shown to be strongly minimal as a metric space with the appropriate metric in Theorem 5.3.3).

4.1.1 Some Characterizations

Here we present some more traditional characterizations of minimal sets.

Proposition 4.1.3. *For a definable set D over the structure \mathfrak{M} the following are equivalent:*

- (i) D is minimal.
- (ii) For each restricted M -formula φ , at most one of $\llbracket \varphi(\mathfrak{M}) \leq \frac{1}{3} \rrbracket$ or $\llbracket \varphi(\mathfrak{M}) \geq \frac{2}{3} \rrbracket$ is non-compact.¹ (In particular we only need to check compactness in \mathfrak{M} , not in arbitrary elementary extensions of \mathfrak{M} .)
- (iii) D is dictionaric (as a subset of $S_n(\mathfrak{M})$) and for every M -definable subset E of D , either $E(\mathfrak{M})$ is compact or $D(\mathfrak{M}) \cap \llbracket d(\mathfrak{M}, E) \geq \varepsilon \rrbracket$ is compact for every $\varepsilon > 0$.

Proof. In this proof we will use the notation $A^{\geq\varepsilon}$ for the set types with distance from A greater than or equal to ε .

¹Note that there is nothing special about $\frac{1}{3}$ and $\frac{2}{3}$ (beyond the fact that they are distinct) or the fact that these are inequalities rather than equalities, so we get the same statement with $\llbracket \varphi = 0 \rrbracket$ and $\llbracket \varphi = 1 \rrbracket$. The proof of the statement in the proposition is conceptually clearer, however.

(i) \Rightarrow (ii). This is obvious.

(ii) \Rightarrow (iii). First we will show that D has a network of definable sets, which is sufficient by Proposition 2.4.2 part (v). Let $p \subseteq D$ be a type, and let U be an open-in- D neighborhood of p . If p is algebraic, then we are done, as $\{p\}$ is a definable set. If p is non-algebraic, then find V such that $p \in V \subseteq \bar{V} \subseteq U$.

We want to argue that every $q \in \partial V$ is an algebraic type. For each $q \in \partial V$ find restricted formula ψ such that $p \in \llbracket \psi < \frac{1}{3} \rrbracket$ and $q \in \llbracket \psi > \frac{2}{3} \rrbracket$. Since p is non-algebraic and V is open, it must be the case that $\llbracket \psi(\mathfrak{M}) \leq \frac{1}{3} \rrbracket$ is non-compact. Therefore $\llbracket \psi(\mathfrak{M}) \geq \frac{2}{3} \rrbracket$ is compact. Since q is contained in the interior of a zerset that is compact in some model, it is algebraic by Lemma 2.3.62. Since every $q \in \partial V$ is algebraic, and therefore definable, \bar{V} is a union of open and definable sets that is topologically closed, and therefore definable. Hence D has a network of definable sets and is dictionaric.

If $E \subseteq D$ is an M -definable set, then for any $\varepsilon > 0$, we can find a restricted formula φ such that $E \subseteq \llbracket \varphi < \frac{1}{3} \rrbracket$ and $D \cap E^{\geq \varepsilon} \subseteq \llbracket \varphi > \frac{2}{3} \rrbracket$. Either $\llbracket \varphi(\mathfrak{M}) \leq \frac{1}{3} \rrbracket$ and therefore $E(\mathfrak{M})$ is compact, or $\llbracket \varphi(\mathfrak{M}) \geq \frac{2}{3} \rrbracket$ and therefore $D(\mathfrak{M}) \cap \llbracket d(\mathfrak{M}, E) \geq \varepsilon \rrbracket$ is compact. This implies that either $E(\mathfrak{M})$ is compact or $D(\mathfrak{M}) \cap \llbracket d(\mathfrak{M}, E) \geq \varepsilon \rrbracket$ is compact for every $\varepsilon > 0$, as required.

(iii) \Rightarrow (i). If $F, G \subseteq D$ are disjoint, M -zerset then we can find $E \supseteq F$ such that $E \subseteq D$ is a definable set disjoint from G . $E(\mathfrak{M})$ is either compact or $D(\mathfrak{M}) \cap \llbracket d(\mathfrak{M}, E) \geq \varepsilon \rrbracket$ is compact for every $\varepsilon > 0$. If $E(\mathfrak{M})$ is compact then $E(\mathfrak{N})$ and therefore $F(\mathfrak{N})$ is compact in every model $\mathfrak{N} \succeq \mathfrak{M}$, since E is definable. Otherwise there is some $\varepsilon > 0$ small enough that $G \cap E^{< \varepsilon} = \emptyset$. In that case find $H \supseteq E^{\geq \varepsilon}$ such that $H \subseteq D$ is a definable set disjoint from E . Since H is disjoint from E there is some $\delta > 0$ small enough that $H \cap E^{< \delta} = \emptyset$, so $H(\mathfrak{M}) \subseteq D(\mathfrak{M}) \cap \llbracket d(\mathfrak{M}, E) \geq \delta \rrbracket$ is a compact set. Since

H is definable we have that $H(\mathfrak{N})$ is compact in every model $\mathfrak{N} \succeq \mathfrak{M}$, therefore $G(\mathfrak{N})$ is as well. \square

We will ultimately show that any minimal set contains a unique non-algebraic type, but this will be a corollary of something slightly more technical so we will defer this to later (Proposition 4.1.12 and Corollary 4.1.13). The condition that D be dictionary when restricting attention to definable sets is necessary in light of Example 2.3.32, since the theory there has more than one non-algebraic \emptyset -type, but appears ‘strongly minimal’ with regards to definable sets in that every definable set is either finite or co-finite.

Compare the following Proposition 4.1.4 to the classical facts that minimal sets in theories with no Vaughtian pairs or over ω -saturated models are strongly minimal.

Proposition 4.1.4.

(i) *If D is a minimal set over a model \mathfrak{M} and either*

- *T is dictionary and has no Vaughtian pairs, or*
- *T has no open-in-definable Vaughtian pairs,*

then D is strongly minimal.

(ii) *If D is a minimal set over \mathfrak{M} , an approximately \aleph_0 -saturated model, then D is strongly minimal.*

Proof. (i) Assume without loss of generality that T is countable (if we prove this for each countable reduct of T over which D is definable, then it will be true for the whole theory). Assume that D is a minimal set containing the non-algebraic type p and that p is not a strongly minimal type. Let q_0, q_1 be distinct non-algebraic extensions of p

to some parameter set $A \supseteq \mathfrak{M}$. Let $\varphi(\bar{x}; \bar{a})$ be a restricted formula with $\bar{a} \in A$ such that $\varphi(q_i; \bar{a}) = i$ for both $i < 2$. Let $\{\bar{b}_i\}_{i < \omega}$ be a sequence of tuples from \mathfrak{M} such that $\text{tp}(\bar{b}_i/\mathfrak{M}) \rightarrow \text{tp}(\bar{a}/\mathfrak{M})$. It must be the case that either infinitely many $i < \omega$ have $\varphi(p; \bar{b}_i) < \frac{2}{3}$ or infinitely many $i < \omega$ have $\varphi(p; \bar{b}_i) > \frac{1}{3}$. Without loss of generality assume that infinitely many $i < \omega$ have $\varphi(p; \bar{b}_i) > \frac{1}{3}$ and restrict to a sub-sequence on which this is always true.

Now we have that for each $i < \omega$, $D \cap \llbracket \varphi(\cdot; \bar{b}_i) \leq \frac{1}{3} \rrbracket$ is algebraic. Let $\mathfrak{N} \succ \mathfrak{M}$ be a proper elementary extension. Note that since $D \cap \llbracket \varphi(\cdot; \bar{b}_i) \leq \frac{1}{3} \rrbracket$ is algebraic we have that $D(\mathfrak{M}) \cap \llbracket \varphi(\mathfrak{M}; \bar{b}_i) \leq \frac{1}{3} \rrbracket = D(\mathfrak{N}) \cap \llbracket \varphi(\mathfrak{N}; \bar{b}_i) \leq \frac{1}{3} \rrbracket$ for each $i < \omega$.

Let $\mathfrak{A}_i = (\mathfrak{N}, \mathfrak{M}, \bar{b}_i)$, and take a non-principal ultraproduct $\mathfrak{A}_\omega = \prod_{i < \omega} \mathfrak{A}_i/\mathcal{U} = (\mathfrak{N}_\omega, \mathfrak{M}_\omega, \bar{b}_\omega)$. Note that $\mathfrak{N}_\omega \succ \mathfrak{M}_\omega$,² that $D(\mathfrak{M}_\omega) \cap \llbracket \varphi(\mathfrak{M}_\omega; \bar{b}_\omega) \leq \varepsilon \rrbracket = D(\mathfrak{N}_\omega) \cap \llbracket \varphi(\mathfrak{N}_\omega; \bar{b}_\omega) \leq \varepsilon \rrbracket$ for any $0 < \varepsilon < \frac{1}{3}$ and that $\text{tp}(\bar{b}_\omega/\mathfrak{M}) = \text{tp}(\bar{a}/\mathfrak{M})$. Therefore p has two non-algebraic extensions r_0, r_1 to $S_n(\mathfrak{M}_\omega)$ such that $\varphi(r_i; \bar{b}_\omega) = i$ for both $i < 2$. This implies that $D(\mathfrak{M}_\omega) \cap \llbracket \varphi(\mathfrak{M}_\omega; \bar{b}_\omega) \leq \varepsilon \rrbracket$ is not metrically compact for any $\varepsilon > 0$.

If T is dictionaric we can find a definable set E such that $D \cap \llbracket \varphi(\cdot; \bar{b}_\omega) \leq \frac{1}{6} \rrbracket \subseteq E \subseteq D \cap \llbracket \varphi(\cdot; \bar{b}_\omega) < \frac{1}{3} \rrbracket$, then we will have that $E(\mathfrak{M}_\omega) = E(\mathfrak{N}_\omega)$ is a non-compact definable set, witnessing that $(\mathfrak{N}_\omega, \mathfrak{M}_\omega)$ is a Vaughtian pair.

Otherwise the set $D \cap \llbracket \varphi(\cdot; \bar{b}_\omega) < \frac{1}{3} \rrbracket$ is open-in-definable and unbounded (it is unbounded because $D \cap \llbracket \varphi(\cdot; \bar{b}_\omega) \leq \frac{1}{6} \rrbracket$ is closed and non-algebraic), witnessing that $(\mathfrak{N}_\omega, \mathfrak{M}_\omega)$ is an open-in-definable Vaughtian pair.

(ii) Let $D(x; \bar{a})$ with $\bar{a} \in \mathfrak{M}$, our approximately \aleph_0 -saturated model, be a definable set such that over some set of parameters B there are distinct non-algebraic types $p_0, p_1 \in D$.

²This is not true for arbitrary ultraproducts of proper elementary pairs; you need a uniform ‘width’ of the extensions, as in you need a fixed $\varepsilon > 0$ such that most \mathfrak{N}_i contain an a with $d_{\text{inf}}(a, \mathfrak{M}_i) > \varepsilon$.

Let $\varphi(x; \bar{b})$ be a formula such that $\varphi(p_i; \bar{b}) = i$ for both $i < 2$. Since p_0 and p_1 are non-algebraic we have that $D \cap \llbracket \varphi(\cdot; \bar{b}) \leq \frac{1}{6} \rrbracket$ and $D \cap \llbracket \varphi(\cdot; \bar{b}) \geq \frac{5}{6} \rrbracket$ are both non-algebraic. Let $\varepsilon > 0$ be small enough that $\#_{>\varepsilon}^{\text{ent}} D(\mathfrak{N}) \cap \llbracket \varphi(\mathfrak{N}; \bar{b}) \leq \frac{1}{3} \rrbracket$ and $\#_{>\varepsilon}^{\text{ent}} D(\mathfrak{N}) \cap \llbracket \varphi(\mathfrak{N}; \bar{b}) \geq \frac{2}{3} \rrbracket$ are both infinite in any model \mathfrak{N} (this is always possible because $D \cap \llbracket \varphi(\cdot; \bar{b}) < \frac{1}{3} \rrbracket$ and $D \cap \llbracket \varphi(\cdot; \bar{b}) > \frac{2}{3} \rrbracket$ are relatively open in D).

Find $\delta > 0$ small enough that $\delta < \frac{1}{3}\varepsilon$ and that if $d(x, y) < \delta$, then $|\varphi(x; \bar{b}) - \varphi(y; \bar{b})| < \frac{1}{9}$. Find $\gamma > 0$ small enough that if $d(\bar{z}, \bar{w}) < \gamma$, then $\sup_x |D(x; \bar{z}) - D(x; \bar{w})| < \delta$. By approximate \aleph_0 -saturation we can find $\bar{c}\bar{e} \equiv \bar{a}\bar{b}$ such that $d(\bar{a}, \bar{c}) < \gamma$ and $\bar{c}\bar{e} \in \mathfrak{M}$. Let $\{u_i\}_{i < \omega}$ be an infinite ($>\varepsilon$)-separated set in $D(\mathfrak{M}; \bar{c}) \cap \llbracket \varphi(\mathfrak{M}; \bar{e}) \leq \frac{1}{3} \rrbracket$, and let $\{v_i\}_{i < \omega}$ be an infinite ($>\varepsilon$)-separated set in $D(\mathfrak{M}; \bar{c}) \cap \llbracket \varphi(\mathfrak{M}; \bar{e}) \geq \frac{2}{3} \rrbracket$. By construction we have that $d_H(D(\mathfrak{M}; \bar{a}), D(\mathfrak{M}; \bar{c})) < \delta$ (where d_H is the Hausdorff metric), so we can find $\{w_i\}_{i < \omega} \subseteq D(\mathfrak{M}; \bar{a})$ and $\{t_i\}_{i < \omega} \subseteq D(\mathfrak{M}; \bar{a})$ such that $d(v_i, w_i) < \delta$ and $d(u_i, t_i) < \delta$ for each $i < \omega$. By construction this implies that $\varphi(w_i; \bar{e}) < \frac{4}{9}$ and $\varphi(t_i; \bar{e}) > \frac{5}{9}$ for every $i < \omega$, and hence we have that $D(\mathfrak{M}; \bar{a}) \cap \llbracket \varphi(\mathfrak{M}; \bar{e}) \leq \frac{4}{9} \rrbracket$ and $D(\mathfrak{M}; \bar{a}) \cap \llbracket \varphi(\mathfrak{M}; \bar{e}) \geq \frac{5}{9} \rrbracket$ are both not metrically compact and therefore not algebraic. Therefore $D(x; \bar{a})$ is not minimal over \mathfrak{M} .

So by the converse we have that if D is a minimal set over an approximately \aleph_0 -saturated model, then it is strongly minimal. \square

It is unclear whether or not Proposition 4.1.4 part (i) can be improved to theories with no Vaughtian pairs in general, in that while a minimal set over A is dictionaric over A , it may not be dictionaric over some $B \supseteq A$.

4.1.2 The Pregeometry of a Strongly Minimal Set

Proposition 4.1.5. *If D is a strongly minimal set definable over the set A , then $X \mapsto \text{acl}(XA)$ restricted to D is the closure operator of a pregeometry with finite character.*

Proof. acl automatically obeys reflexivity and transitivity, so all we need to verify is finite character and exchange.

For finite character, suppose that $a \in D(\mathfrak{C})$ and $a \in \text{acl}(B)$ for some set $B \supseteq A$ (not necessarily in D). Let $\chi(\bar{x}; \bar{b})$ be a distance predicate for an algebraic subset of D containing a , with $\bar{b} \in B$. Let p be the unique non-algebraic global type in D . Find a restricted formula φ such that $\sup_{\bar{x}, \bar{y}} |\chi(\bar{x}; \bar{y}) - \varphi(\bar{x}; \bar{y})| < \frac{1}{4}d(p, \mathfrak{C})$. $\varphi(\bar{x}; \bar{b})$ only depends on finitely many parameters in the tuple \bar{b} . Note the following: $\llbracket \chi(\cdot; \bar{b}) = 0 \rrbracket \subseteq D \cap \llbracket \varphi(\cdot; \bar{b}) \leq \frac{1}{4}d(p, \mathfrak{C}) \rrbracket \subseteq \llbracket \chi(\cdot; \bar{b}) \leq \frac{1}{2}d(p, \mathfrak{C}) \rrbracket \not\equiv p$. Therefore $D \cap \llbracket \varphi(\cdot; \bar{b}) \leq \frac{1}{4}d(p, \mathfrak{C}) \rrbracket$ contains a and is algebraic, so $a \in \text{acl}(B_0)$ for some finite $B_0 \subseteq B$.

For exchange, let $b \in D(\mathfrak{C}) \setminus \text{acl}(A)$, and let $c \in D(\mathfrak{C}) \setminus \text{acl}(Ab)$. We want to show that $b \notin \text{acl}(Ac)$. Since D has a unique non-algebraic type over any parameter set, any such pair like bc has the same type. Let B' be a set of realizations of $\text{tp}(b/A)$ of cardinality $(|\mathcal{L}| + |A| + 2^{\aleph_0})^+$, and let c' be a realization of the unique non-algebraic type in D over the set AB' . For any $b'_0, b'_1 \in B'$ we have that $b'_0 c' \equiv_A b'_1 c'$, but there are too many realizations of $\text{tp}(b'_0/Ac')$ for it to be algebraic, so we must have $b' \notin \text{acl}(Ac')$ for any $b' \in B'$. Since $b' c' \equiv bc$, the same must be true for b and c , so $b \notin \text{acl}(Ac)$, as required. \square

Note that in a strongly minimal set, acl has finite character relative to arbitrary parameters, not just those in the strongly minimal set. In fact it has finite character in a particularly strong sense in that if $c \in \text{acl}(AB)$, where c is in a strongly minimal

set over A , then there is a finite tuple $\bar{c} \in C$ and an $\varepsilon > 0$ such that for any \bar{e} with $d(\bar{c}, \bar{e}) < \varepsilon$, $b \in \text{acl}(A\bar{e})$. The following corollary summarizes this more precisely.

It follows from this proposition that all of the machinery of pregeometries, such as bases and invariant dimension number, works in continuous logic just as it does in discrete logic. The following corollary summarizes this more precisely.

Corollary 4.1.6. *If D is a strongly minimal set definable over the set A (which we may assume without loss of generality is countable), then for any model $\mathfrak{M} \supseteq A$, we have that $D(\mathfrak{M}) \subseteq \text{acl}(AB)$ for any $B \subseteq D(\mathfrak{M})$ which is a basis with regards to the pregeometry induced by the closure operator $X \mapsto \text{acl}(X \cup A)$.*

Any two bases of $D(\mathfrak{M})$ have the same cardinality, any two bases of the same cardinality are elementarily equivalent, and any basis is an A -indiscernible set. So in particular if $B_i \subseteq D(\mathfrak{M}_i)$ for both $i < 2$ are bases of the same cardinality, then any bijection $f : B_0 \rightarrow B_1$ extends to an elementary map $f' : D(\mathfrak{M}_0) \equiv D(\mathfrak{M}_1)$.

Finally, for any $X \subseteq D(\mathfrak{C})$, if $\#^{\text{dc}}X > \#^{\text{dc}}\text{acl}(A) + |\mathcal{L}|$ then $\dim(X) = \#^{\text{dc}}X$. So in particular, if T and A are countable then for any $X \subseteq D(\mathfrak{C})$ with $\#^{\text{dc}}X$ uncountable, $\dim(X) = \#^{\text{dc}}X$.

The following lemma is useful for understanding the metric properties of strongly minimal sets. Perhaps unsurprisingly, the metric in a strongly minimal set always behaves somewhat like a locally finite edge relation in a discrete strongly minimal set. We are including it in this subsection because its proof relies on the pregeometric structure of strongly minimal sets.

Lemma 4.1.7 (Uniform Local Compactness). *Let D be a strongly minimal set definable over the set A . There exists an $\varepsilon > 0$ such that for every $\delta > 0$, there exists an $n < \omega$*

such that for any model $\mathfrak{M} \supseteq A$, every closed ε -ball B in $D(\mathfrak{M})$ can be covered by at most n open δ balls.

So in particular for any $a, b \in D(\mathfrak{M})$ with $d(a, b) < \varepsilon$, we have that $a \in \text{acl}(b)$.

Proof. For the purposes of this proof we will use the notations $B_{\leq \alpha}(x)$ and $B_{< \alpha}(x)$ to represent closed and open balls in D (rather than in the ambient structure).

Let p be the generic type in D over A . Let b be a realization of p . Let q be the generic type in D over Ab . We have that $d(b, q) > 0$. Find $\varepsilon > 0$ satisfying $\varepsilon < \frac{d(b, q)}{2}$. The set $\llbracket d(x, b) \leq \varepsilon \rrbracket$ does not contain q , so it must be algebraic. This implies that for every $\delta > 0$ there is an $m_\delta < \omega$ such that $B_{\leq \varepsilon}(b)$ can be covered by m_δ open δ -balls. Consider the formula

$$\chi_\delta(x) := \sup_{y_0, \dots, y_{m_\delta-1}} \inf_z \max\{d(x, z) \div \varepsilon, \max_{i < m_\delta} \delta \div d(y_i, z)\}.$$

Note that $\chi_\delta(x)$ cannot take on negative values. Let c_0, \dots, c_{m_δ} be chosen so that the open δ -balls centered at these points cover $B_{\leq \varepsilon}(b)$. Since $B_{\leq \gamma}$ is compact for some $\gamma > \varepsilon$, by compactness there is a $\gamma > \varepsilon$ such that $B_{\leq \gamma}(b)$ is also covered by $\bigcup_{i < m_\delta} B_{< \delta}(c_i)$. Then by compactness again we get that there is some $\theta < \delta$ such that $\bigcup_{i < m_\delta} B_{< \theta}(c_i)$ covers $B_{\leq \gamma}(b)$. By plugging in these c_i 's for the y_i 's in $\gamma_\delta(b)$, we get a witness that $\chi_\delta(a) > \min\{\gamma - \varepsilon, \delta - \theta\} > 0$. Conversely, we have that for any $b \models D$, if $\chi_\delta(b) > 0$, then $B_{\leq \delta}(b)$ can be covered by m_δ open δ -balls. The set $\llbracket \chi_\delta(x) > 0 \rrbracket$ is an open neighborhood of p , so we have that $D \cap \llbracket \chi_\delta(x) = 0 \rrbracket$ is algebraic. For any $e \in D$ with $\chi_\delta(e) = 0$, we must have that $d(e, p) > \varepsilon$. To see this, let f be a realization of q . We have that $d(b, f) \geq d(b, q) > 2\varepsilon$, by the choice of ε . Since e is in $\text{acl}(\emptyset)$, b and f have the same type over Ae , implying that $d(b, e) = d(f, e)$. By the triangle inequality, these must both be

greater than ε . This implies that the set $F = (\llbracket \chi_\delta(x) = 0 \rrbracket \cap D)^{\leq \varepsilon} \cap D$ does not contain p and is therefore also algebraic. This implies that there is some $k_\delta < \omega$ such that $F(\mathfrak{C})$ is covered by k_δ open δ -balls, which implies that for any $e \in D$ with $\chi_\delta(e) = 0$, $B_{\leq \varepsilon}(e)$ is also covered by k_δ open δ -balls (since it is a subset of $F(\mathfrak{C})$). Therefore we have that for any $g \in D$, $B_{\leq \varepsilon}(g)$ is covered by no more than $n_\delta = \max\{m_\delta, k_\delta\}$ open δ -balls.

Since we can do this for any $\delta > 0$, and since this property is clearly preserved under passing to substructures (possibly increasing n as a function of δ), we have the required result. \square

4.1.3 Strongly Minimal Types

Now we would like to give definitions of strongly minimal types directly and relate them to the notion of strongly minimal sets. As part of this we will need a notion of ‘approximate algebraicity.’

Definition 4.1.8.

- (i) A type $p \in S_n(A)$ is said to be $(< \varepsilon)$ -algebraic if for any $\mathfrak{M} \supseteq A$ and any $q \in S_n(\mathfrak{M})$ such that $q \upharpoonright A = p$, $d(q, \mathfrak{M}) < \varepsilon$ (thinking of \mathfrak{M} as the set of types it realizes in $S_n(\mathfrak{M})$).
- (ii) A set of types is said to be $(< \varepsilon)$ -algebraic if every type in it is.
- (iii) A type $p \in S_n(A)$ is *pre-minimal (over A)* if for all sufficiently small $\varepsilon > 0$, p is d -atomic in the set of non- $(< \varepsilon)$ -algebraic types in $S_n(A)$.³

³Pre-minimal types have the same relationship with minimal sets that strongly minimal types have with strongly minimal sets. There is a terminological issue we are inheriting from discrete logic here. In discrete logic, strongly minimal sets have strongly minimal types, and weakly minimal sets have minimal types, but minimal sets do not get a name for their special types (which are admittedly not

(iv) A type $p \in S_n(A)$ is *strongly minimal* if it has a unique pre-minimal global extension. ◁

Note that if a zeroset F is $(< \varepsilon)$ -algebraic, then in any model \mathfrak{M} , by compactness $F(\mathfrak{M})$ is covered by finitely many $(< \varepsilon)$ -balls with centers in \mathfrak{M} . Something approximating the converse is true as well, but we won't need it. Also note that a zeroset F is algebraic if and only if it is $(< \varepsilon)$ -algebraic for every $\varepsilon > 0$.

Note that the set of non- $(< \varepsilon)$ -algebraic types is always closed, and if \mathfrak{M} is a model, then the set of non- $(< \varepsilon)$ -algebraic types in $S_1(\mathfrak{M})$ is precisely $S_1(\mathfrak{M}) \setminus \mathfrak{M}^{<\varepsilon}$, so in particular a type $p \in S_n(A)$ is strongly minimal if for all sufficiently small $\varepsilon > 0$, p has a unique global extension in the set $S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$ (where \mathfrak{C} is the monster model), and furthermore that extension is relatively d -atomic in $S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$.

Compare this with the following definition of strongly minimal types in discrete logic: A type $p \in S_n(A)$ is strongly minimal if it has a unique global extension in the set $S_n(\mathfrak{C}) \setminus \mathfrak{C}$, and furthermore that extension is relatively isolated in $S_n(\mathfrak{C}) \setminus \mathfrak{C}$.

Note that the definition of strongly minimal type implies that if p is a strongly minimal type and q is a global extension of it that is not equal to p 's special extension, then $q \in \mathfrak{C}^{<\varepsilon}$ for all sufficiently small $\varepsilon > 0$, so in particular q is realized and therefore algebraic. This implies that over any parameter set containing the domain of p , p has a unique non-algebraic extension.

The following is a characterization of strongly minimal types in terms of Morley rank and degree, which is developed in the context of continuous logic in [BY08c].

very special). Calling such types 'weakly minimal,' while arguably sensible, would invite confusion. 'Locally minimal' is another option, but this is unfortunate when talking about 'locally minimal global types' and erroneously suggests a direct relationship with local types.

The notion of Morley rank used here is the one corresponding to the (f, ε) -Cantor-Bendixson derivative defined in that paper, but the statement of this proposition is not sensitive to the particular kind of Morley rank used, since they are all asymptotically equivalent as $\varepsilon \rightarrow 0$ and since in a totally transcendental theory the statement that $Md_\varepsilon(p) = 1$ for all sufficiently small $\varepsilon > 0$ is precisely the same as the statement that p is a stationary type. The machinery of Morley rank is not used anywhere else in this section, and this proposition is included only for comparison to the classical statement that a type p is strongly minimal if and only if $M(R, d)(p) = (1, 1)$, where $M(R, d)(p) = (MR(p), Md(p))$, so we will omit actually defining Morley rank here.

Proposition 4.1.9. *A type $p \in S_n(A)$ is strongly minimal if and only if $M(R, d)_\varepsilon(p) = (1, 1)$ for all sufficiently small $\varepsilon > 0$.*

Proof. (\Rightarrow). Assume that $p \in S_n(A)$ is strongly minimal. Let q be the unique pre-minimal global extension, so that in particular for all sufficiently small $\varepsilon > 0$, q is relatively d -atomic in $S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$. $(S_n(\mathfrak{C}))'_\varepsilon = S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$, so q is relatively d -atomic in $(S_n(\mathfrak{C}))'_\varepsilon$. Therefore $q \notin (S_n(\mathfrak{C}))''_\varepsilon$ and q is contained in open subsets of $(S_n(\mathfrak{C}))'_\varepsilon$ of arbitrarily small diameter, so $M(R, d)_\varepsilon(q) = (1, 1)$. q is the unique extension of p to $(S_n(\mathfrak{C}))'_\varepsilon$, so $M(R, d)_\varepsilon(p) = (1, 1)$ as well.

(\Leftarrow). Since p has ordinal Morley ranks it has global non-forking extensions. Since $Md_\varepsilon(p) = 1$ for all sufficiently small $\varepsilon > 0$, it has a unique global non-forking extension. Let q be its unique global non-forking extension. For each sufficiently small $\varepsilon > 0$, we have that p is contained in a relatively open subset of $(S_n(\mathfrak{C}))'_\varepsilon$ of diameter $\leq 2\varepsilon$. For any $0 < \delta < \varepsilon$, we have that $(S_n(\mathfrak{C}))'_\delta \supseteq (S_n(\mathfrak{C}))'_\varepsilon$, therefore q has open neighborhoods of arbitrarily small diameter in $(S_n(\mathfrak{C}))'_\varepsilon = S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$. And thus q is relatively d -atomic

in $S_n(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$, and q is a pre-minimal global extension of p .

We still need to show that q is the unique pre-minimal global extension of p . Assume that p has another pre-minimal global extension r . By the (\Rightarrow) direction we'd have that $MR_\varepsilon(r) = 1$ for all sufficiently small $\varepsilon > 0$, but that would contradict that $Md_\varepsilon(p) = 1$ for all sufficiently small $\varepsilon > 0$. \square

There is one additional hiccup in the development of strongly minimal sets in continuous logic. Although we typically think of strongly minimal types as coming from strongly minimal sets, it is not hard to show that the following is true in discrete logic:

If $p \in S_1(A)$ is a type whose unique non-algebraic extension to some parameter set $B \supseteq A$ is contained in a strongly minimal set D , definable over B , then there is a strongly minimal set E definable over A containing p as its unique non-algebraic type.

This pleasant fact fails in continuous logic (see Counterexample C.2.2). Assuming we are working over models, the relationship between strongly minimal types and strongly minimal sets is precisely as it is in discrete logic. Curiously, this arguably relies on the same kind of behavior that gives us the mild pathology of Counterexample C.2.1.

Proposition 4.1.10. *If $p \in S_1(\mathfrak{M})$ is a strongly minimal (resp. pre-minimal) type with \mathfrak{M} a model, then there is an \mathfrak{M} -definable strongly minimal (resp. minimal) set D containing p (with no assumptions on $S_1(\mathfrak{M})$ or T).*

Proof. For each $k < \omega$, let $F_k = S_1(\mathfrak{M}) \setminus \mathfrak{M}^{<2^{-k}}$. Let $\ell < \omega$ be large enough that $p \in F_k$ for any $k \geq \ell$. By Lemma 2.4.5 part (i), we can find a closed set $A \subseteq S_1(\mathfrak{M})$ and a formula $\varphi : S_1(\mathfrak{M}) \rightarrow [0, 1]$ such that

- $p \in \llbracket \varphi = 0 \rrbracket \subseteq A$,
- $p \in \llbracket \varphi = 0 \rrbracket \subseteq \text{int } A^{<\varepsilon}$ for every $\varepsilon > 0$, and
- $\llbracket \varphi \leq 2^{-k-1} \rrbracket \cap A \cap F_k = \{p\}$ for every $\ell \leq k < \omega$.

Note that while $\llbracket \varphi = 0 \rrbracket$ may contain types other than p , if $q \in \llbracket \varphi = 0 \rrbracket \setminus \{p\}$, then $q \notin F_k$ for every $\ell \leq k < \omega$, so in particular q is algebraic and realized in \mathfrak{M} .

For each k with $\ell \leq k < \omega$, let $G_k = A \cap \llbracket 2^{-k-2} \leq \varphi \leq 2^{-k-1} \rrbracket$. Note that $G_k \cap F_k = \emptyset$, since at most it could contain $\{p\}$, but it cannot contain p . This implies that $G_k \subseteq \mathfrak{M}^{<2^{-k}} = \bigcup_{a \in \mathfrak{M}} B_{<2^{-k}}(a)$, so by compactness there is a finite set $M_k \subseteq \mathfrak{M}$ such that $G_k \subseteq M_k^{<2^{-k}}$. By restricting to a subset if necessary we may assume that $M_k \subseteq G_k^{<2^{-k}}$ as well, i.e. $d_H(G_k, M_k) < 2^{-k}$. Note that as a finite union of elements of \mathfrak{M} , each M_k is an \mathfrak{M} -definable set.

Let $D = \llbracket \varphi = 0 \rrbracket \cup \bigcup_{\ell \leq k < \omega} M_k$. To see that D is closed, note that $\llbracket \varphi = 0 \rrbracket$ is clearly closed so we only need to argue that the accumulation points of $\bigcup_{\ell \leq k < \omega} M_k$ are all contained in D . Let q be an accumulation point of $\bigcup_{\ell \leq k < \omega} M_k$, and assume that $\varphi(q) > 0$. Find U , an open neighborhood of q , such that \bar{U} is disjoint from $\llbracket \varphi = 0 \rrbracket$. Find $\varepsilon > 0$ small enough that $\bar{U}^{\leq \varepsilon}$ is disjoint from $\llbracket \varphi = 0 \rrbracket$. By compactness there is a $k < \omega$ large enough that $(A \cap \llbracket \varphi \leq 2^{-k} \rrbracket)^{\leq 2^{-k}}$ is disjoint from $\bar{U}^{\leq \varepsilon}$. This implies that $D \cap \bar{U}^{\leq \varepsilon}$ is a finite set, so q must be in D . If $\varphi(q) = 0$ then $q \in \llbracket \varphi = 0 \rrbracket \subseteq D$, therefore D is closed.

To see that D is definable, pick $\varepsilon > 0$, and note that $D^{<\varepsilon/2} \supseteq \llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k$ for some $m < \omega$ (since $d_H(G_m, M_m) < 2^{-m}$ for every $m < \omega$), so we have

$$D^{<\varepsilon} \supseteq \left(\llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k \right)^{<\varepsilon/2} \cup \bigcup_{\ell \leq k < \omega} M_k^{<\varepsilon}.$$

So we just have to argue that $\llbracket \varphi = 0 \rrbracket \subseteq \text{int} \left(\llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k \right)^{< \varepsilon/2}$. Find $\delta > 0$ small enough that $\delta < \frac{\varepsilon}{2}$ and $\llbracket \varphi \leq \delta \rrbracket^{\leq \delta} \cap A \subseteq \llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k$, which must be possible by compactness. Let $q \in \llbracket \varphi < \delta \rrbracket \cap \text{int} A^{< \delta}$. We have that $d_{\text{inf}}(q, A) < \delta$, so let $r \in A$ such that $d(q, r) < \delta$. This implies that $r \in \llbracket \varphi \leq \delta \rrbracket^{\leq \delta}$, so in particular $r \in \llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k$. Therefore

$$\begin{aligned} \llbracket \varphi \leq \delta \rrbracket \cap \text{int} A^{< \delta} &\subseteq \text{int} \left(\llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k \right)^{< \delta} \\ &\subseteq \text{int} \left(\llbracket \varphi = 0 \rrbracket \cup \bigcup_{m \leq k < \omega} G_k \right)^{< \varepsilon/2}, \end{aligned}$$

as required. So $D \subseteq \text{int} D^{< \varepsilon}$ for every $\varepsilon > 0$, and thus D is definable. By construction every type $q \in D \setminus \{p\}$ is algebraic (and realized in \mathfrak{M}), so D is a minimal set. If the unique non-algebraic type in D is strongly minimal then D is a strongly minimal set as well, by Proposition 4.1.12 (the proof of that proposition does not rely on this proposition). \square

Really if A is any parameter set, $p \in S_1(A)$ is a strongly minimal type, D is any definable set containing p , and $\mathfrak{M} \supseteq A$ is a model, then we get an $A \cup D(\mathfrak{M})$ -definable strongly minimal set containing p . And we also have the same for any open or open-in-definable set containing p .

4.1.4 Approximately (Strongly) Minimal Pairs

We can recover a fact analogous to Proposition 4.1.10 without assuming anything about the parameter set, but we need a dictionary theory and a slight weakening of the notion of strongly minimal set.

Definition 4.1.11.

- (i) A pair (D, φ) of a definable set $D \subseteq S_n(A)$ and an A -formula φ is *approximately minimal* if the zeroset $\llbracket \varphi = 0 \rrbracket \subseteq D$ is non-algebraic and for every pair $F, G \subseteq D$ of disjoint A -zerosets, every model $\mathfrak{M} \supseteq A$, and every $\varepsilon > 0$, at least one of $F \cap \llbracket \varphi \leq \varepsilon \rrbracket$ and $G \cap \llbracket \varphi \leq \varepsilon \rrbracket$ is $(< \varepsilon)$ -algebraic.
- (ii) A pair (D, φ) is *approximately strongly minimal* if it is approximately minimal over every set of parameters over which it is definable.
- (iii) A definable set D is approximately (strongly) minimal if there is some formula φ such that (D, φ) is approximately (strongly) minimal. \triangleleft

Obviously a (strongly) minimal set is approximately (strongly) minimal if we just let φ be the distance predicate of the set. Note that for any kind of minimality it is sufficient to check sets of the form $\llbracket \varphi \leq r \rrbracket$ with φ restricted and r rational.

Proposition 4.1.12.

- (i) If (D, φ) is approximately minimal (over the set A) then there is a unique non-algebraic A -type $p \in \llbracket \varphi = 0 \rrbracket \subseteq D$. We say that p is the generic type of (D, φ) .
- (ii) If (D, φ) is approximately strongly minimal (resp. approximately minimal) then its generic type is strongly minimal (resp. pre-minimal).
- (iii) If (D, φ) is (approximately) minimal and its generic type is strongly minimal, then (D, φ) is (approximately) strongly minimal (where a pair (D, φ) is strongly minimal if D is strongly minimal).

Proof. (i) Assume that there are two distinct non-algebraic types, q_0 and q_1 contained in $\llbracket \varphi = 0 \rrbracket \subseteq S_n(A)$. Let $F_0, F_1 \subseteq \llbracket \varphi = 0 \rrbracket$ be disjoint A -zerosets such that $p_i \in F_i$ for both $i < 2$. Since for each $\varepsilon > 0$, at least one of F_0 and F_1 must be $(< \varepsilon)$ -algebraic, it must be that for some $i < 2$, $F_i(\mathfrak{C}) \subseteq \mathfrak{C}^{< \varepsilon}$ for arbitrarily small $\varepsilon > 0$, so in particular $F_i(\mathfrak{C}) \subseteq \mathfrak{C}$, implying that F_i is algebraic. This is a contradiction, therefore there cannot be two non-algebraic types in $\llbracket \varphi = 0 \rrbracket \subseteq S_n(A)$, but $\llbracket \varphi = 0 \rrbracket$ is non-algebraic so there must be at least one.

(ii) Let $X = A$ if (D, φ) is approximately minimal over A , and let $X = \mathfrak{C}$ if (D, φ) is approximately strongly minimal. We need to show that p is d -atomic in the set of non- $(< \varepsilon)$ -algebraic types in $S_n(X)$. For any $\varepsilon > 0$ let F_ε denote the (closed) set of non- $(< \varepsilon)$ -algebraic types in $S_n(X)$. Fix $\varepsilon > 0$ small enough that $p \in F_\varepsilon$. Find $\delta > 0$ small enough that $\delta < \frac{1}{2}\varepsilon$ and $p \notin (D \cap \llbracket \varphi \geq \frac{1}{2}\varepsilon \rrbracket)^{\leq \delta}$, and then consider $U = (D^{< \delta} \setminus (D \cap \llbracket \varphi \geq \frac{1}{2}\varepsilon \rrbracket)^{\leq \delta}) \cap F_\varepsilon$. Note that $p \in U$ and that U is relatively open in F_ε . We want to show that $U \subseteq B_{\leq \delta}(p)$. Let $q \in U$. By construction this implies that $d(q, D) < \delta$ so there is some $r \in D$ such that $d(q, r) < \delta$. Also by construction r cannot be in $D \cap \llbracket \varphi \geq \frac{1}{2}\varepsilon \rrbracket$. Assume that $r \in D \cap \llbracket 0 < \varphi < \frac{1}{2}\varepsilon \rrbracket$. There must be some $0 < \gamma < \sigma < \frac{1}{2}\varepsilon$ such that $r \in D \cap \llbracket \gamma \leq \varphi \leq \sigma \rrbracket$, but $D \cap \llbracket \gamma \leq \varphi \leq \sigma \rrbracket$ is $(< \frac{1}{2}\varepsilon)$ -algebraic, so in particular $r \notin F_{\varepsilon/2}$. But this is a contradiction since for any model $\mathfrak{M} \supseteq X$ (i.e. $\mathfrak{M} \supseteq A$ or $\mathfrak{M} = \mathfrak{C}$) and extension q' to $S_n(\mathfrak{M})$, there is an extension r' to $S_n(\mathfrak{M})$ such that $d(q', \mathfrak{M}) \leq d(q', r') + d(r', \mathfrak{M}) < \delta + \frac{1}{2}\varepsilon < \varepsilon$. But $q \in F_\varepsilon$, so q has an extension q'' with $d(q'', \mathfrak{M}) \geq \varepsilon$.

Therefore r must be in $\llbracket \varphi = 0 \rrbracket$. Assume that $r \neq p$. This implies that r is algebraic so that in particular $r \in \mathfrak{M}$ for any $\mathfrak{M} \supseteq X$, which is again a contradiction since for any extension q' of q there is an extension r' of r such that $d(q', \mathfrak{M}) \leq d(q', r') + d(r', \mathfrak{M}) <$

$\delta + 0 < \frac{1}{2}\varepsilon < \varepsilon$. Hence it must be the case that $r = p$. Since this is true for any $q \in U$, this implies that $U \subseteq B_{<\delta}(p) \subseteq B_{\leq\delta}(p)$. Since we can do this for any sufficiently small $\delta > 0$, we have that p is relatively d -atomic in F_ε and the same is true for any sufficiently small $\varepsilon > 0$.

(iii) If D is approximately minimal but not approximately strongly minimal then its generic type has two distinct non-algebraic extensions to some set of parameters, so it is not strongly minimal. The only thing to prove is that if D is minimal and its generic type is strongly minimal then D is strongly minimal (and not just approximately strongly minimal). This follows from the fact that if p is the global strongly minimal type in $D \subseteq S_n(\mathfrak{C})$, then any $q \in D \setminus \{p\}$ must be an extension of some type in $D \subseteq S_n(A)$. If it is an extension of $p \upharpoonright A$, then it is algebraic, and if it is an extension of some $r \neq p \upharpoonright A$ then it also must be algebraic. \square

It is easy to come up with an example of a definable set D such that (D, φ) and (D, ψ) are approximately strongly minimal but have different generic types—the union of two disjoint strongly minimal sets for instance.

Corollary 4.1.13.

- (i) *If $D \subseteq S_n(A)$ is minimal then there is a unique non-algebraic type $p \in D$.*
- (ii) *If D is strongly minimal then there is a unique non-algebraic type $p \in D \subseteq S_n(A)$ for any set of parameters A over which D is definable.*

The strongly minimal type whose existence is guaranteed by Corollary 4.1.13 is also referred to as the *generic type* of the corresponding set.

Finally we come to the advantage we gain by passing to this weaker notion, as promised at the beginning of this section. Given a strongly minimal type in a dictionary

theory we can always find an approximately strongly minimal pair pointing to that type and definable over the same set that the type is over, and likewise with pre-minimal types and minimal sets.

Proposition 4.1.14.

- (i) *If $S_n(A)$ is dictionary and $p \in S_n(A)$ is a pre-minimal type, then there is an A -definable approximately minimal pair (D, P) pointing to p .*
- (ii) *If $D \subseteq S_n(A)$ is approximately minimal over a model \mathfrak{M} , then there is a $(D(\mathfrak{M}) \cup A)$ -definable minimal set $E \subseteq D$ pointing to the same type.*
- (iii) *If $D \subseteq S_n(A)$ is approximately strongly minimal (as part of the pair (D, P)), then for any model $\mathfrak{M} \supseteq A$, there is a $(D(\mathfrak{M}) \cup A)$ -definable strongly minimal set $E \subseteq D$ pointing to the same type. Furthermore, for any model $\mathfrak{N} \succ \mathfrak{M}$, $E(\mathfrak{N}) \setminus \llbracket P(\mathfrak{N}) = 0 \rrbracket \subseteq D(\mathfrak{M})$.*

Proof. (i) This follows from applying Lemma 2.4.5 to the closed set $\{p\}$ which is relatively definable in the set $F_i \subseteq S_n(A)$, where F_i is the set of all non- $(< 2^{-i-k})$ -algebraic types and $k > 0$ is chosen so that $p \in F_i$ for every $i < \omega$. The lemma gives us a definable set D and a formula φ such that for each $i < \omega$, $D \cap F_i \cap \llbracket \varphi \leq 2^{-i} \rrbracket = p$. To see that (D, φ) is an approximately minimal pair, pick $\varepsilon > 0$. Find $i < \omega$ such that $2^{-i-1} < \varepsilon \leq 2^{-i}$, and let $G, H \subseteq D$ be disjoint zerosets. At most one of G or H can contain p , so assume without loss of generality that $p \notin G$. We have that $G \cap F_i \cap \llbracket \varphi \leq 2^{-i} \rrbracket \subseteq p$, so $G \cap F_i \cap \llbracket \varphi \leq \varepsilon \rrbracket \subseteq p$ as well, but $p \notin G$, so $G \cap F_i \cap \llbracket \varphi \leq \varepsilon \rrbracket$. This implies that $G \cap \llbracket \varphi \leq \varepsilon \rrbracket$ is contained in the set of $(< 2^{-i-k})$ -algebraic types so in particular since $2^{-i-k} \leq 2^{-i-1} < \varepsilon$, we have that every type in $G \cap \llbracket \varphi \leq \varepsilon \rrbracket$ is $(< \varepsilon)$ -algebraic. Therefore (D, φ) is an approximately minimal pair.

(ii) This follows from the comment after Proposition 4.1.10.

(iii) Most of this follows from part (ii) and Proposition 4.1.12. The only thing we need to verify is the last sentence, which follows from the fact that anything in $D(\mathfrak{M})$ not realizing the strongly minimal type over \mathfrak{M} must be algebraic over \mathfrak{M} in the first place and so realized in it. The strongly minimal type is contained in $\llbracket \varphi = 0 \rrbracket$, so the result follows. \square

Note that parts (ii) and (iii) do not require any assumptions about the type space. The following example shows that the dictionariness stipulation in part (i) cannot be removed.

Example 4.1.15. *A non-dictionaric superstable theory with a strongly minimal type over \emptyset but no approximately strongly minimal sets over \emptyset .*

Verification. Let $\mathcal{L} = \{P_0, P_1\}$ be a language with two unary 1-Lipschitz $[0, 1]$ -valued predicates, and let \mathfrak{M} be an \mathcal{L} -structure whose universe is $\omega \times [0, 1]$ and whose metric is given by $d((n, x), (m, y)) = 1$ if $n \neq m$ and $d((n, x), (n, y)) = 2^{-n} + d(x, y)$ if $x \neq y$. Let $P_0((n, x)) = 2^{-n}$ and $P_1((n, x)) = x$. Finally let $T = \text{Th}(\mathfrak{M})$.

The type space $S_1(\emptyset)$ is homeomorphic to $(\omega + 1) \times [0, 1]$.

Note that if a definable set has non-empty intersection with one of the sets of types of the form $\{n\} \times [0, 1]$ for $n < \omega$, then it must contain all of it, because this set is metrically isolated from the rest of the type space and is topologically connected but uniformly metrically discrete. So to show that none of the types in $\{\omega\} \times [0, 1]$ are pointed to by an approximately strongly minimal set, all we need to do is show that if a definable set contains one such type then it must contain some type in $\{n\} \times [0, 1]$ for some $n < \omega$. This follows immediately because if $p \in \{\omega\} \times [0, 1]$ then it is the limit of

types in $\{n\} \times [0, 1]$ for $n < \omega$ that are uniformly metrically separated, so if F is a closed set whose intersection with $\{n\} \times [0, 1]$ for $n < \omega$ is empty and whose intersection with $\{\omega\} \times [0, 1]$ is precisely p , then $p \notin \text{int } F^{<\varepsilon}$ for any $0 < \varepsilon < 1$, and so F is not definable.

So if D is a definable set containing some $p \in \{\omega\} \times [0, 1]$, then D must contain all of $\{n\} \times [0, 1]$ for sufficiently large $n < \omega$, so since D is closed it must contain all of $\{\omega\} \times [0, 1]$. If we let φ be a formula such that $\varphi(p) = 0$, then for any $\varepsilon > 0$, $D \cap \llbracket \varphi \leq \varepsilon \rrbracket$ contains some $q \in \{\omega\} \times [0, 1]$ with $q \neq p$. \square

In all of the examples of approximately minimal sets we know of there is an obvious pseudo-metric ρ on D such that D/ρ is a minimal set whose generic type corresponds exactly to the generic type in the original set. For instance, in Counterexample C.2.2 there is a definable pseudo-metric ρ such that $\rho(x, y) = 0$ if and only if either x and y are both in the same Hilbert space sphere, or they are not in a Hilbert space sphere and $x = y$. It is not clear if this is always possible. In the case of discrete strongly minimal types, however, it is so.

Definition 4.1.16. A type p is *discrete* if there is an $\varepsilon > 0$ such that if $a, b \models p$ and $d(a, b) < \varepsilon$, then $a = b$. \triangleleft

Proposition 4.1.17. *If (D, φ) is an approximately minimal pair over some parameter set A pointing to a discrete pre-minimal type p in a dictionary type space, then there is an A -definable set $E \subseteq D$ containing p and an A -definable equivalence relation ρ on E such that E/ρ is minimal and such that the quotient map is a bijection when restricted to the pre-minimal type. (In particular this means that if p is strongly minimal, then the corresponding type in E/ρ is strongly minimal as well.)*

Proof. Let $\varepsilon > 0$ be such that if $a, b \models p$ and $d(a, b) < \varepsilon$ then $d(a, b) = 0$.

By compactness there must be a $\delta > 0$ such that for any $a, b \in D(\mathfrak{C}) \cap \llbracket \varphi(\mathfrak{C}) \leq \delta \rrbracket$, if $d(a, b) < \frac{3}{4}\varepsilon$, then $d(a, b) < \frac{1}{4}\varepsilon$. This implies that in the zeroset $D \cap \llbracket \varphi \leq \delta \rrbracket$, the formula $\rho(x, y) = \frac{4}{3\varepsilon}(d(x, y) \div \frac{1}{4}\varepsilon) \downarrow 1$ is a $\{0, 1\}$ -valued equivalence relation which is equality on the set of realizations of p . Let E be a definable subset of D such that $\llbracket \varphi = 0 \rrbracket \subseteq E \subseteq \llbracket \varphi < \delta \rrbracket$. Clearly ρ is the required equivalence relation on E . \square

This leaves the question in general.

Question 4.1.18. *If $D \subseteq S_n(A)$ is approximately minimal with generic type p then, does there always exist an A -definable pseudo-metric ρ on D such that D/ρ is minimal with generic type q and the quotient map $D \rightarrow D/\rho$ restricted to the set of realizations of p is a bijection with the set of realizations of q ?*

4.2 Categoricity

Here we recall a few definitions and facts that will be relevant to the rest of this chapter.

Definition 4.2.1. A theory T is κ -categorical if any two models $\mathfrak{M}, \mathfrak{N} \models T$ with $\#^{\text{dc}}\mathfrak{M} = \#^{\text{dc}}\mathfrak{N} = \kappa$ are isomorphic. \triangleleft

Fact 4.2.2 (Morley's Theorem for Continuous Logic, [BY05, SU11]). *If a countable theory is κ -categorical for some $\kappa \geq \aleph_1$, then it is λ -categorical for all $\lambda \geq \aleph_1$.*

In light of this fact we have the following definition.

Definition 4.2.3. A theory is *inseparably categorical* if it is κ -categorical for some (or equivalently all) $\kappa \geq \aleph_1$. \triangleleft

Fact 4.2.4 ([BY05]). *An inseparably categorical theory is ω -stable, in the sense of Definition 2.5.11.*

Fact 4.2.5. *Any totally transcendental theory is atomic (i.e. for any set of parameters A , d -atomic types are dense in $S_1(A)$), so in particular atomic models exist over any set of parameters.*

Just as in discrete logic, if a model \mathfrak{M} is atomic over A , then it is prime over A .

4.3 Vaughtian Pairs

Definition 4.3.1. If $X \subseteq S_1(A)$ is a definable (resp. open, open-in-definable, or locatable) set containing a non-algebraic type, with A countable, and if $\mathfrak{M} \succ \mathfrak{N} \supseteq A$ is a proper elementary pair such that $X(\mathfrak{M}) = X(\mathfrak{N})$, then we say that $(\mathfrak{M}, \mathfrak{N})$ is a *definable* (resp. *open*, *open-in-definable*, or *locatable*) *Vaughtian pair* (with regards to X). A Vaughtian pair with no qualifier is a definable Vaughtian pair. \triangleleft

Note that if a theory has no open-in-definable Vaughtian pairs then it has no open Vaughtian pairs and no definable Vaughtian pairs, since open sets and definable sets are special cases of open-in-definable sets. In discrete logic these three notions are essentially the same, however in continuous logic they are all distinct (see Counterexample C.2.4).⁴

Proposition 4.3.2. *Suppose that T is a countable theory, $U \subseteq D \subseteq S_n(B)$ is an open-in-definable set, and $B \subseteq \mathfrak{N} \prec \mathfrak{M}$ is a Vaughtian pair over U , then there exists a model \mathfrak{A} such that for some open $V \subseteq D$ containing a non-algebraic type, $\#^{\text{dc}}\mathfrak{A} = \aleph_1$ but $\#^{\text{dc}}V(\mathfrak{A}) = \aleph_0$.*

⁴Although, at the moment it is unknown, but unlikely, that no open-in-definable Vaughtian pairs implies no locatable Vaughtian pairs.

Proof. Find $p \in U$ that is non-algebraic, and find φ such that $p \in \llbracket \varphi < \frac{1}{2} \rrbracket \subseteq U$. Since φ is restricted and D is definable, they are definable over some countable set $B_0 \subseteq B$. Now it is clear that $(\mathfrak{M}, \mathfrak{N})$ is still an open-in-definable Vaughtian pair with regards to $V = \llbracket \varphi < \frac{1}{2} \rrbracket$.

Let $(\mathfrak{D}_1, \mathfrak{D}_0)$ be a countable elementary sub-pre-structure of $(\mathfrak{M}, \mathfrak{N})$ such that $\mathfrak{D}_0 \supseteq B_0$. By the same argument as in discrete logic, we can find countable pre-structures $(\mathfrak{A}_1, \mathfrak{A}_0) \succeq (\mathfrak{D}_1, \mathfrak{D}_0)$ such that \mathfrak{A}_1 and \mathfrak{A}_0 realize the same types, and are both (exactly) \aleph_0 -homogeneous (as pre-structures, i.e. only over parameters that are actually in them), so that in particular $\mathfrak{A}_1 \cong \mathfrak{A}_0$. We can also ensure that \mathfrak{A}_1 and \mathfrak{A}_0 both realize dense subsets of D , which implies that \mathfrak{A}_0 realizes a metrically dense subset of $U(\mathfrak{A}_0) = U(\mathfrak{A}_1)$, since U is open-in- D .

We can run the same elementary chain argument as in discrete logic, the argument that \mathfrak{A}_i for limit i is still isomorphic (as a countable pre-structure) to \mathfrak{A}_0 still works, but we need to argue that at each stage i the set $U(\mathfrak{A}_0)$ is still metrically dense in $U(\mathfrak{A}_i)$. Obviously, by construction, $U(\mathfrak{A}_0)$ is metrically dense in $U(\mathfrak{A}_1)$. Suppose that for some $i < \omega_1$, $U(\mathfrak{A}_0)$ is metrically dense in $U(\mathfrak{A}_i)$ and $(\mathfrak{A}_{i+1}, \mathfrak{A}_i) \cong (\mathfrak{A}_1, \mathfrak{A}_0)$. We get from this that $U(\mathfrak{A}_i)$ is metrically dense in $U(\mathfrak{A}_{i+1})$, implying that $U(\mathfrak{A}_0)$ is metrically dense in $U(\mathfrak{A}_{i+1})$ as well. Suppose that for some limit $i < \omega_1$ we have that for all $j < i$, $U(\mathfrak{A}_0)$ is metrically dense in $U(\mathfrak{A}_j)$. Then any $a \in U(\mathfrak{A}_i)$ is actually in $U(\mathfrak{A}_j)$ for some $j < i$, so it is still in the metric closure of $U(\mathfrak{A}_0)$, therefore $U(\mathfrak{A}_0)$ is still metrically dense in $U(\mathfrak{A}_i)$.

After taking the union and metric completion to get $\mathfrak{A} = \overline{\bigcup_{i < \omega_1} \mathfrak{A}_i}$, we need to argue that $U(\mathfrak{A}_0)$ is still metrically dense in $U(\mathfrak{A})$. Suppose that $a \in U(\mathfrak{A})$. Since it is a relatively open set in a definable set, it must be the metric limit of some sequence

$a_k \in \bigcup_{i < \omega_1} U(\mathfrak{A}_i)$. Therefore since $U(\mathfrak{A}_0)$ is dense in every $U(\mathfrak{A}_i)$, it must be in the metric closure of $U(\mathfrak{A}_0)$ as well, so in particular $\#^{\text{dc}}U(\mathfrak{A}) = \aleph_0$.

Finally we need to show that $\#^{\text{dc}}\mathfrak{A} = \aleph_1$. Let $\varepsilon > 0$ be such that for some $a \in \mathfrak{A}_1$, $d(a, \mathfrak{A}_0) > \varepsilon$. By isomorphism, for each $i < \omega_1$, we can find an $a_i \in \mathfrak{A}_{i+1} \setminus \mathfrak{A}_i^{<\varepsilon}$, this is a $(> \varepsilon)$ -separated set of size \aleph_1 , so $\#^{\text{dc}}\mathfrak{A} \geq \aleph_1$. On the other hand \mathfrak{A} clearly has a dense subset of size \aleph_1 , so $\#^{\text{dc}}\mathfrak{A} = \aleph_1$, as required. \square

Corollary 4.3.3. *If T is a countable \aleph_1 -categorical theory, then T has no open-in-definable Vaughtian pairs.*

Proof. Suppose that T is a countable theory with an open-in-definable Vaughtian pair. By Proposition 4.3.2, T has a model \mathfrak{A} with $\#^{\text{dc}}\mathfrak{A} = \aleph_1$ but for which some non-algebraic type over a countable set p satisfies $\#^{\text{dc}}p(\mathfrak{A}) \leq \omega$. This implies that \mathfrak{A} is not \aleph_1 -saturated, implying that T is not \aleph_1 -categorical. \square

Corollary 4.3.4. *If T is a countable \aleph_1 -categorical theory, then T has no open Vaughtian pairs in any imaginaries. In particular it has no imaginary Vaughtian pairs.*

Proof. An imaginary expansion of a κ -categorical theory is still κ -categorical. \square

Given the existence of Counterexample C.2.10, we arrive at a natural question:

Question 4.3.5. *If T is a countable \aleph_1 -categorical theory, does it follow that T has no locatable Vaughtian pairs?*

4.4 Theories with Strongly Minimal Sets over the Prime Model

4.4.1 Main Theorem

Assuming T is a theory with strongly minimal sets, part of the Baldwin-Lachlan characterization goes through exactly. This statement is analogous to the discrete statement ‘For a countable theory T , if T has a prime model and a minimal set definable over it, then for any $\kappa \geq \aleph_1$, T is κ -categorical if and only if T has no Vaughtian pairs.’ Our continuous generalization of this statement is made more complicated by a few factors. We strengthen the result by using the weakening of strongly minimal set given in Definition 4.1.11. We also need one of two strengthenings of no Vaughtian pairs, either of which is sufficient. And, given the presence of certain counterexamples in continuous logic (such as Counterexample C.2.10), we would like to state the result both for definable sets in the home sort and for arbitrary imaginaries.

Theorem 4.4.1. *Let T be a countable complete theory with non-compact models, and let κ be any uncountable cardinal.*

(i) *If T has a prime model and an approximately minimal set definable over it, then the following are equivalent.*

(a) *T is κ -categorical.*

(b) *T is dictionary and has no Vaughtian pairs.*

(c) *T has no open-in-definable Vaughtian pairs.*

(ii) If T has a prime model and an approximately minimal imaginary definable over it, then the following are equivalent.

(a) T is κ -categorical.

(b) T is dictionary, and T^{eq} has no Vaughtian pairs.

(c) T^{eq} has no open Vaughtian pairs.

Proof. Since T has an approximately minimal set definable over its prime model, by Proposition 4.1.14 we have that there is a minimal set definable over its prime model.

(i). If T is κ -categorical, then T is ω -stable, and therefore dictionary, and also has no open-in-definable Vaughtian pairs, and therefore no Vaughtian pairs.

If either (b) or (c) is true, then by Proposition 4.1.4, the definable minimal set is strongly minimal. Let D be this strongly minimal set. Let A be a basis in D of cardinality κ , and let \mathfrak{A} be prime over $\text{acl}(A)$. If \mathfrak{B} is a model of density character κ , then since T has no Vaughtian pairs, $D(\mathfrak{B})$ has density character κ , and so we can find a basis B of $D(\mathfrak{B})$ of cardinality κ . Therefore we can find an isomorphism $D(\mathfrak{A}) \cong D(\mathfrak{B})$. Since \mathfrak{A} is prime over $D(\mathfrak{A})$ we can extend this isomorphism to an embedding $\mathfrak{A} \preceq \mathfrak{B}$, but since T has no Vaughtian pairs this must be an isomorphism. Therefore all models of density character κ are isomorphic to \mathfrak{A} .

(ii). The proof here is the same as the proof of (i) with the following notes: If T is dictionary, then T^{eq} is dictionary over models, which is enough to show that minimal sets are strongly minimal in case (b). In case (c), if T has no open Vaughtian pairs in imaginaries, then it has no open-in-definable Vaughtian pairs in imaginaries, since definable subsets of imaginaries are imaginaries. \square

Note that the conclusion of Theorem B.3.2 also holds if we know that T has an

approximately minimal set (or imaginary) definable over every model, as then we can show that T_A for some countable set of constants is κ -categorical, and therefore ω -stable, implying that T has a prime model with an approximately minimal set (or imaginary) definable over it.

4.4.2 Strongly Minimal Sets over the Prime Model

One might hope that if an inseparably categorical theory has a strongly minimal set over some model, then it has one over its prime model, but this is not true (see Theorem C.2.9). What is unclear at the moment is the possibility of an inseparably categorical theory that has a strongly minimal imaginary over some models but not others. What we do have is that this cannot happen if the strongly minimal set is discrete.

Proposition 4.4.2. *If T is a dictionary theory with no imaginary Vaughtian pairs such that for some model \mathfrak{M} there is an infinite discrete \mathfrak{M} -definable imaginary, then for every model \mathfrak{N} there is an infinite discrete \mathfrak{N} -definable imaginary.*

Proof. Assume that T is countable. By the imaginary normal form lemma, we may assume that the infinite discrete imaginary over \mathfrak{M} is a definable subset $D(x, \bar{a})$ of some \emptyset -definable imaginary I . Assume without loss that the metric on $D(\cdot, \bar{a})$ is $\{0, 1\}$ -valued.

Since $D(x, \bar{a})$ is a distance predicate, part of $\text{tp}(\bar{a})$ says that

$$\chi_0(\bar{a}) = \sup_x \inf_y D(y, \bar{a}) \uparrow |D(x, \bar{a}) - d(x, y)|$$

and

$$\chi_1(\bar{a}) = \sup_x |D(x, \bar{a}) - \inf_y ((D(y, \bar{a}) + d(x, y)) \downarrow 1)|$$

both vanish. (These are the axioms for distance predicates given in [BYBHU08] right before Theorem 9.12.) $\text{tp}(\bar{a})$ also says that

$$\eta(\bar{a}) = \sup_{x,y} (d(x,y) \downarrow (1 \div d(x,y))) \div 4(D(x,\bar{a}) + D(y,\bar{a}))$$

vanishes. (This is just $\sup_{x,y \in D(\cdot, \bar{a})} (d(x,y) \downarrow (1 \div d(x,y)))$ expanded out.)

Fix \mathfrak{N} another model of \mathfrak{M} , and let $\{\bar{b}_i\}_{i < \omega}$ be a sequence of elements of \mathfrak{N} such that $\text{tp}(\bar{b}_i) \rightarrow \text{tp}(\bar{a})$ (we can choose it to be a sequence rather than a net since T is countable), such that in particular $\chi_0(\bar{b}_i), \chi_1(\bar{b}_i), \eta(\bar{b}_i) < 2^{-i}$ for every $i < \omega$. For each $i < \omega$ let E_i be a \bar{b}_i -definable set such that $[D(\cdot, \bar{b}_i) < 2^{-i+1}] \supseteq E_i \supseteq \text{int} E_i \supseteq [D(\cdot, \bar{b}_i) \leq 2^{-i}]$.

For each $i < \omega$, we have that

$$\mathfrak{A} \models \sup_{x,y \in E_i} (d(x,y) \downarrow (1 \div d(x,y))) \div 8 \cdot 2^{-i+1}.$$

So in particular for sufficiently large $i < \omega$, for any $x, y \in E_i$, either $d(x,y) < \frac{1}{4}$ or $d(x,y) > \frac{3}{4}$, implying that we can define a $\{0, 1\}$ -valued equivalence relation on E_i by $\rho(x,y) = 2(d(x,y) \div \frac{1}{4}) \downarrow 1$ for sufficiently large $i < \omega$.

Now assume that for all $i < \omega$ for which ρ defines an equivalence relation on E_i , there are finitely many ρ -equivalence classes in E_i . Let $\mathfrak{A} \succ \mathfrak{N}$ be a proper elementary extension. Fix a non-principal ultrafilter \mathcal{U} on ω , and consider the ultraproduct $(\mathfrak{B}, \mathfrak{N}', \bar{c}) = \prod_{i < \omega} (\mathfrak{A}, \mathfrak{N}, \bar{b}_i) / \mathcal{U}$, where $(\mathfrak{A}, \mathfrak{N}, \bar{b}_i)$ is a structure whose universe is \mathfrak{A} with a distance predicate for \mathfrak{N} and constants for \bar{b}_i . Note that \mathfrak{B} is a proper elementary extension of \mathfrak{N}' , because \mathfrak{A} is uniformly a proper elementary extension of \mathfrak{N} in the product. Clearly we have that $\text{tp}(\bar{c}) = \text{tp}(\bar{a})$, so $D(\cdot, \bar{c})$ is a definable set which is infinite and has

a $\{0, 1\}$ -valued metric.

We want to argue that $\mathfrak{B} \succ \mathfrak{N}'$ is a Vaughtian pair with regards to the set $D(\cdot, \bar{c})$. Let $\alpha(i)$ be a sequence such that $\alpha(i) \in (\mathfrak{A}, \mathfrak{N}, \bar{b}_i)$ and such that the limit α/\mathcal{U} is in $D(\cdot, \bar{c})$. This means that we must have $\lim_{\mathcal{U}} D(\alpha(i), \bar{b}_i) = 0$. Pick $\varepsilon > 0$. For each $\alpha(i)$ we have that there is a $\beta(i)$ such that $D(\beta(i), \bar{b}_i) < 2^{-i}$ and $|D(\alpha(i), \bar{b}_i) - d(\alpha(i), \beta(i))| < 2^{-i}$, so in particular $\beta(i) \in E_i(\mathfrak{A})$. Since E_i only has finitely many ρ -classes, there must be a $\gamma(i) \in E_i(\mathfrak{N})$ such that $d(\beta(i), \gamma(i)) < \frac{1}{4}$, implying that $d(\alpha(i), \gamma(i)) < \frac{1}{4} + 2^{-i+1}$ and in particular $d(\alpha(i), \mathfrak{N}) < \frac{1}{4} + 2^{-i+1}$. So then in the limit we have $d(\alpha/\mathcal{U}, \mathfrak{N}) \leq \frac{1}{4}$, but since $D(\cdot, \bar{c})$ has a $\{0, 1\}$ -valued metric this implies that $\alpha/\mathcal{U} \in D(\mathfrak{N}', \bar{c})$, so $(\mathfrak{B}, \mathfrak{N}')$ is a Vaughtian pair.

Since T has no imaginary Vaughtian pairs this is a contradiction and some E_i must have infinitely many ρ -classes, so then E_i/ρ is an infinite discrete imaginary over \mathfrak{N} . \square

Corollary 4.4.3. *If T is an inseparably categorical theory with a discrete strongly minimal imaginary over some set, then it has a discrete strongly minimal imaginary over the prime model.*

Proof. By the previous proposition there is an infinite discrete imaginary over the prime model, maybe not strongly minimal, but by ω -stability we can find a minimal set in it which must by no imaginary Vaughtian pairs be strongly minimal. \square

4.4.3 The Number of Separable Models

We immediately get the following, generalizing a result originally due to Morley [Mor70].

Proposition 4.4.4. *If T is an inseparably categorical theory with a minimal set or imaginary definable over the prime model then T has at most countably many separable*

models.

Proof. By Proposition 3.2.13 we may assume that D is a strongly minimal set (rather than a strongly minimal imaginary) by appending a \emptyset -definable imaginary sort. Let $D(x; \bar{a})$ be a minimal set definable over the prime model. Since T is dictionaric and has no Vaughtian pairs, this is a strongly minimal set. Let \mathfrak{A} be any separable model of T . The prime model \mathfrak{M} embeds into \mathfrak{A} so we can find $\bar{b} \in A$ with $\bar{b} \equiv \bar{a}$, and we get that $D(x; \bar{b})$ is a strongly minimal set. It must be the case that $\dim(D(\mathfrak{A}; \bar{b})) \leq \omega$, since \mathfrak{A} is separable. Assume that we chose \bar{b} so that $\dim(D(\mathfrak{A}; \bar{b})) \leq \omega$ among parameters with the same type as \bar{b} . If \mathfrak{B} is any other separable model we may find $\bar{c} \in \mathfrak{B}$ with $\bar{c} \equiv \bar{a}$ and $\dim(D(\mathfrak{B}; \bar{c}))$ minimal. If $\dim(D(\mathfrak{A}; \bar{b})) = \dim(D(\mathfrak{B}; \bar{c}))$, then we get an elementary map $f : D(\mathfrak{A}; \bar{b}) \cup \bar{b} \rightarrow D(\mathfrak{B}; \bar{c}) \cup \bar{c}$, which extends to an isomorphism between \mathfrak{A} and \mathfrak{B} , since T has no (imaginary) Vaughtian pairs. Therefore since there are only countably many possible dimensions for separable models we get that there are at most countably many separable models. \square

Lemma 4.4.5. *If T is an inseparably categorical theory and (D, φ) is a \emptyset -definable approximately minimal pair, then every model \mathfrak{M} of T is prime over $\{\varphi(\mathfrak{M}) = 0\}$.*

Proof. Assume that some model \mathfrak{A} is not prime over $\{\varphi(\mathfrak{A}) = 0\}$. Let $\mathfrak{B} \prec \mathfrak{A}$ be prime over $\{\varphi(\mathfrak{A}) = 0\}$, so in particular $\{\varphi(\mathfrak{A}) = 0\} = \{\varphi(\mathfrak{B}) = 0\}$. Let \mathfrak{M} be the prime model of T , and find some embedding $\mathfrak{M} \preceq \mathfrak{B}$. Since D is \emptyset -definable, we can find a strongly minimal $E \subseteq M$ definable over \mathfrak{M} . Note that every element of $E(\mathfrak{C}) \setminus \{\varphi(\mathfrak{C}) = 0\}$, where \mathfrak{C} is the monster model, is algebraic over \mathfrak{M} . This implies that $E(\mathfrak{B}) = E(\mathfrak{A})$ as well, but then this is a Vaughtian pair, which is a contradiction. Therefore every model \mathfrak{M} of T is prime over \mathfrak{M} . \square

The difficulty in characterizing the number of separable models of an inseparably categorical theory seems to be related to the phenomenon of non- d -finite types, identified by Ben Yaacov and Usvyatsov in [BYU07]. Finitary types in continuous logic can behave analogously to ω -types in discrete logic. Ben Yaacov and Usvyatsov identified a class of types they call d -finite which behave like discrete finitary types. In their paper they were able to prove that a superstable theory with ‘enough uniformly d -finite types’ (where uniformly d -finite is a technical strengthening of d -finite) has either 1 or infinitely many separable models, whereas in general an ω -stable continuous theory can have any finite number of separable models, including 2 (they also showed that in the presence of ‘enough d -finite types,’ a continuous theory cannot have exactly 2 separable models).

So in some easy cases we get the full Baldwin-Lachlan theorem on the number of separable models of an inseparably categorical theory:

Theorem 4.4.6. *If T is an inseparably categorical theory and any of the following occur, then T has either 1 or ω separable models.*

- *T has a \emptyset -definable approximately minimal pair (possibly in an imaginary).*
- *T is ultrametric or has totally disconnected type spaces.*
- *T has enough uniformly d -finite types.*

Proof. The only difficult case is the first one, but this is covered by Lemma 4.4.5. The second case follows from the fact that such theories are interdefinable with many-sorted discrete theories, and the third case is a direct corollary of Theorem 4.7 in [BYU07]. \square

Counterexample C.2.6 shows that we cannot hope to show that every model of an inseparably categorical theory is exactly homogeneous, although at the moment they

seem to always be approximately homogeneous. The problem in Counterexample C.2.6 is that the strongly minimal set requires non- d -finite parameters. We can get some traction by assuming that we have a strongly minimal set definable over a d -finite tuple of parameters, but adapting the proof of the Baldwin-Lachlan theorem any further than this is unclear at the moment.

Proposition 4.4.7. *Suppose that $p(x, \bar{a})$ is a strongly minimal type, $\text{tp}(\bar{a})$ is d -finite, and T is a dictionary theory. If \mathfrak{M} is a model containing \bar{a} and there is $n < \omega$ such that for every $\varepsilon > 0$, there is $\bar{b} \in \mathfrak{M}$ with $d(\bar{a}, \bar{b}) < \varepsilon$ and $\dim(p(\mathfrak{M}, \bar{b})) \geq n$, then $\dim(p(\mathfrak{M}, \bar{a})) \geq n$.*

Proof. Let $q(\bar{y}, \bar{a})$ be the type of an n -element independent sequence in $p(x, \bar{a})$. Find $\varepsilon > 0$ small enough that $B_{\leq 2\varepsilon}(c) \cap p(\mathfrak{C}, \bar{a}) \subseteq \text{acl}(c\bar{a})$ for every $c \models p(x, \bar{a})$. (Note that this works because we can define an approximately strongly minimal pair over \bar{a} pointing at $p(x, \bar{a})$. By then examining the definition of approximately strongly minimal pair, we get that for any $c \models p(x, \bar{a})$ there is some $\varepsilon > 0$ such that the inclusion $B_{\leq 2\varepsilon}(c) \cap p(\mathfrak{C}, \bar{a}) \subseteq \text{acl}(c\bar{a})$ holds, but then by automorphisms of \mathfrak{C} this works for every $c \models p(x, \bar{a})$.) Fix an approximately strongly minimal pair $(D(x, \bar{a}), \varphi(x, \bar{a}))$ pointing at $p(x, \bar{a})$, and find a $\gamma > 0$ small enough that if $\bar{a} \equiv \bar{b}$ and $d(\bar{a}, \bar{b}) < \gamma$, then $d_H(D(\mathfrak{C}, \bar{a}), D(\mathfrak{C}, \bar{b})) < \varepsilon$.

By d -finiteness (of $\text{tp}(\bar{b})$), we can find a $\delta > 0$ such that for any \bar{b} with $d(\bar{a}, \bar{b}) < \delta$ and any $\bar{c} \models q(\bar{y}, \bar{b})$, we can find \bar{e} such that $\bar{a}\bar{e} \equiv \bar{b}\bar{c}$ and $d(\bar{e}, \bar{c}) < \varepsilon$.

Let $\bar{b} \in \mathfrak{M}$ be such that $\bar{a} \equiv \bar{b}$ and $d(\bar{a}, \bar{b}) < \delta \downarrow \gamma$. By construction we have that $d_H(D(\mathfrak{M}, \bar{a}), D(\mathfrak{M}, \bar{b})) < \varepsilon$. Let \bar{c} be an independent tuple of length n of realizations of $p(x, \bar{b})$. Let \bar{e}^0 be a tuple of elements of $D(\mathfrak{M}, \bar{a})$ such that $d(\bar{c}, \bar{e}^0) < \varepsilon$. By construction we also have that there is a tuple \bar{e} (in the monster model) such that $\bar{a}\bar{e} \equiv \bar{b}\bar{c}$ and

$d(\bar{e}, \bar{c}) < \varepsilon$. For each $i < n$, we have that $d(e_i^0, e_i) \leq d(e_i^0, c_i) + d(c_i, e_i) < \varepsilon + \varepsilon$. Therefore $e_i \in B_{\leq 2\varepsilon}(e_i^0) \cap p(\mathfrak{C}, \bar{a})$ for each i , but this implies by construction that $e_i \in \text{acl}(e_i^0 \bar{a})$, so in particular $e_i \in \mathfrak{M}$. Therefore $\bar{e} \in \mathfrak{M}$, and we have that $\dim(p(\mathfrak{M}, \bar{a})) \geq n$. \square

Corollary 4.4.8. *If $\text{tp}(\bar{a})$ is d -finite, $p(x, \bar{a})$ is strongly minimal, and T is dictionaric, then in any approximately homogeneous model \mathfrak{M} , for any $\bar{a}, \bar{b} \in \mathfrak{M}$ with $\bar{a} \equiv \bar{b}$, we have that $\dim(p(\mathfrak{M}, \bar{a})) = \dim(p(\mathfrak{M}, \bar{b}))$.*

For the following recall that in continuous logic there is an ω -stable theory with precisely 2 separable models. Also note that we do not know whether or not inseparably categorical theories have enough d -finite types (in the technical sense of [BYU07]).

Corollary 4.4.9. *If T is an inseparably categorical theory with a strongly minimal type definable over a d -finite tuple in the prime model, then T does not have precisely 2 separable models.*

Proof. Assume that T is not \aleph_0 -categorical. The prime model and the approximately \aleph_0 -saturated model are both approximately \aleph_0 -homogeneous by Proposition 2.5.7 and Fact 1.5 in [BYU07], respectively, so if $D(x; \bar{a})$ is a strongly minimal set with $\text{tp}(\bar{a})$ atomic and d -finite, then for any $\bar{b} \equiv \bar{a}$, $\dim(D(\mathfrak{M}; \bar{a})) = \dim(D(\mathfrak{M}; \bar{b}))$ where \mathfrak{M} is either prime or approximately \aleph_0 -saturated. We know that the approximately \aleph_0 -saturated model must have $\dim(D(\mathfrak{M}; \bar{a})) = \omega$, so we must have that $\dim(D(\mathfrak{N}; \bar{a})) = n < \omega$ where \mathfrak{N} is the prime model (otherwise T would be \aleph_0 -categorical).

All we need to do is argue that there is a model $\mathfrak{A} \succ \mathfrak{N}$ which is neither prime nor approximately \aleph_0 -saturated. Let b realize the non-algebraic type over \mathfrak{N} in $D(x; \bar{a})$, and let \mathfrak{A} be atomic over $\mathfrak{N}b$. Clearly $\dim(D(\mathfrak{A}; \bar{a})) > n$. We need to argue that it is $n + 1$.

By Lemma 4.5 in [BYU07] if $\dim(D(\mathfrak{A}; \bar{a})) > n + 1$, then there is some c realizing the non-algebraic type in $D(x; \bar{a})$ over $\mathfrak{A}b$, implying that $\text{tp}(c/\mathfrak{A})$ forks over $\mathfrak{A}b$, but this can only happen if $c \in \text{acl}(\mathfrak{A}b)$, contradicting that c realizes the strongly minimal type over $\mathfrak{A}b$. Therefore $\dim(D(\mathfrak{A}; \bar{a})) = n + 1$ and \mathfrak{A} is neither prime nor approximately \aleph_0 -saturated and T has at least 3 separable models. \square

4.5 Generalizations of Strongly Minimal Sets

Here we will explore two generalizations of strongly minimal sets motivated by trying to characterize inseparably categorical theories.

4.5.1 Orthonormalizable Sets

While IHS is not strongly minimal, it still behaves very similarly to a strongly minimal set in some important ways. The following definition is motivated by this more general behavior.

Definition 4.5.1. A definable set $D(\bar{x})$ is *orthonormalizable* if it is non-algebraic and for any set of parameters A and any two non-algebraic types $p, q \in \llbracket D \rrbracket_A \subseteq S_{\bar{x}}(A)$, any realization of p is inter-algebraic with a realization of q over A .

A theory is *orthonormalizable* if $(d(x, x) = 0)$ is orthonormalizable. \triangleleft

Note that a strongly minimal set is clearly orthonormalizable. That said, orthonormalizable is strictly weaker, even in discrete logic. Consider, for example, a theory consisting of two infinite definable sets with a bijection between them. This theory is orthonormalizable, but not strongly minimal. Counterexample C.2.5 gives an example

of a theory in which the home sort is orthonormalizable and has no strongly minimal sets, but which also has a strongly minimal imaginary.

Proposition 4.5.2. *Any countable orthonormalizable theory is ω -stable.*

Proof. Fix a countable parameter set A . Let $p \in S_1(A)$ be a non-algebraic type. Let b be some realization of p , and consider $\text{acl}(Ab)$. This set is metrically separable. By assumption, every type in $S_1(A)$ is realized in $\text{acl}(Ab)$, so $S_1(A)$ must be metrically separable as well, since the map $c \mapsto \text{tp}(c/A)$ is 1-Lipschitz. \square

More generally, one can show that an orthonormalizable theory (in any size language) is always totally transcendental.

While strongly minimal sets allow for a generalization of the process of finding a basis of a vector space, orthonormalizable sets allow for a generalization of Gram–Schmidt orthonormalization.

Proposition 4.5.3. *If T is an orthonormalizable theory with prime model \mathfrak{M} , then for any non-algebraic type $p \in S_1(\mathfrak{M})$ and any $\mathfrak{N} \succeq \mathfrak{M}$, there is a Morley sequence $\{a_i\}_{i < \beta} \subseteq \mathfrak{N}$ in p over \mathfrak{M} such that $\mathfrak{N} = \text{acl}(\mathfrak{M}a_{<\beta})$.*

Proof. Pick some non-algebraic type $p_0 \in S_1(\mathfrak{M})$, and let p be the unique global non-forking extension of it. For each ordinal i , find $b_i \in \mathfrak{N}$ such that $b_i \notin \text{acl}(\mathfrak{M}a_{<i})$, if this is possible, otherwise stop. By orthonormalizability there is a_i which is inter-algebraic with b_i over $\mathfrak{M}a_{<i}$ and which realizes $p \upharpoonright \mathfrak{M}a_{<i}$.

This process must stop at some β . By construction $\mathfrak{N} = \text{acl}(\mathfrak{M}a_{<\beta})$ and $\{a_i\}_{i < \omega}$ is a Morley sequence over \mathfrak{M} . \square

Corollary 4.5.4. *If T is a countable orthonormalizable \mathcal{L} -theory then it is κ -categorical for every $\kappa > |\mathcal{L}|$.*

Proof. By ω -stability, indiscernible sequences are indiscernible sets so all we need to do is show that if $\#^{\text{dc}}\mathfrak{M} = \#^{\text{dc}}\mathfrak{N} > \aleph_0$, then the Morley sequences constructed in Proposition 4.5.3 are the same cardinality. But this follows from the fact that for any set B in any structure \mathfrak{A} , $\#^{\text{dc}}\text{acl}_{\mathfrak{A}}(B) \leq \aleph_0 + |\mathcal{L}| + \#^{\text{dc}}B$ and the fact that a non-trivial indiscernible sequence is $(> \varepsilon)$ -separated, so if I is a non-trivial indiscernible sequence in \mathfrak{F} , then $\#^{\text{dc}}\mathfrak{F} \geq |I|$. \square

Proposition 4.5.5. *If T is a countable orthonormalizable theory, then acl restricted to D is the closure operator of a pregeometry with countable character.*

Proof. acl automatically obeys reflexivity, monotonicity, transitivity, and countable character, so all we need to verify is exchange.

Let \mathfrak{M} be a model, and fix a and b such that $a \notin \text{acl}(\mathfrak{M}) = \mathfrak{M}$ and $b \notin \text{acl}(\mathfrak{M}a)$. Since \mathfrak{M} is a model, $\text{tp}(a/\mathfrak{M})$ is stationary. Let p be its unique global non-forking extension. By orthonormalizability, there is b' realizing $p \upharpoonright \mathfrak{M}a$ such that b and b' are inter-algebraic over $\mathfrak{M}a$. By construction, ab' is a two element Morley sequence in p . If we let $\{c_i\}_{i \leq \kappa}$ with $\kappa = (\#^{\text{dc}}\mathfrak{M} + 2^{\aleph_0})^+$ be a Morley sequence in p over \mathfrak{M} , then we have that for any $i < \kappa$, $ab' \equiv_{\mathfrak{M}} c_i c_{\kappa}$. This implies that $\text{tp}(a/\mathfrak{M}b')$ can have too many realizations to be algebraic, so $a \notin \text{acl}(\mathfrak{M}b')$, but now if we assume that $a \in \text{acl}(\mathfrak{M}b)$, then $\text{acl}(\mathfrak{M}b) = \text{acl}(\mathfrak{M}b')$ (since b and b' are inter-algebraic over $\mathfrak{M}a$), which is a contradiction. Therefore $a \notin \text{acl}(\mathfrak{M}b)$. So for any model \mathfrak{M} and a and b , we have that $b \in \text{acl}(\mathfrak{M}a)$ if and only if $a \in \text{acl}(\mathfrak{M}b)$.

Now for the general case, let E be a set of parameters, and assume that $a \notin \text{acl}(E)$ and $b \notin \text{acl}(Ea)$ but that $a \in \text{acl}(b)$. Since b is in $\text{acl}(a)$, we can find a set $\{c_i\}_{i < \lambda}$ of distinct realizations of $\text{tp}(b/a)$, with $\lambda = (\#^{\text{dc}}E + 2^{\aleph_0})^+$. For any c_i we have that

$a \in \text{acl}(c_i)$. Now, since $\text{tp}(a)$ is not algebraic, if we fix a model $\mathfrak{M} \supseteq E$ with $\#^{\text{dc}}\mathfrak{M} = \aleph_0 + \#^{\text{dc}}E$, we can find $a' \equiv a$ such that $a' \notin \mathfrak{M}$. Let σ be an automorphism of the monster taking a to a' , and for each $i < \lambda$, let $c'_i = \sigma c_i$. By cardinality considerations, there must be λ many c'_i 's that are not in $\text{acl}(\mathfrak{M}a')$, and furthermore, by the infinitary pigeonhole principle, there must be some type $p(x, a') \in S_x(\mathfrak{M}a')$ such that λ many c'_i 's that are not in \mathfrak{M} realize $p(x, a')$. This implies that $p(x, a')$ is not algebraic over $\mathfrak{M}a'$. Let c_i be such that c'_i realized $p(x, a')$. Now we have that $a' \notin \text{acl}(\mathfrak{M})$ and $c'_i \notin \text{acl}(\mathfrak{M}a')$, but $a' \in \text{acl}(\mathfrak{M}c'_i)$, contradicting the conclusion of the earlier part of this proof. Therefore we must have that if $a \notin \text{acl}(E)$ and $b \notin \text{acl}(Ea)$, then $b \notin \text{acl}(Ea)$, and so acl satisfies exchange. \square

Question 4.5.6. *Is there a characterization of orthonormalizable sets on the level of formulas or definable sets, analogous to the definition of a strongly minimal set?*

Unfortunately Counterexample C.2.7 shows that there are inseparably categorical theories that do not have any orthonormalizable sets in their home sort. Just as with strongly minimal sets, the natural question is then whether or not they are always present in some imaginary.

Question 4.5.7. *Does every inseparably categorical theory have an orthonormalizable imaginary?*

4.5.2 Peripheral Types

The unique orthonormal type in Hilbert space has a special property that can be seen as a finite ε generalization of minimality, specifically if $o \in S_1(\mathfrak{H})$ is the unique type of a norm 1 element orthogonal to the Hilbert space \mathfrak{H} , then for any $p \in S_1(\mathfrak{H})$ we have




Figure 10: $d(p, o) \leq \sqrt{2(1 - d_{\text{inf}}(p, \mathfrak{H}))}$ in $S_1(\mathfrak{H})$

that $d(p, o) \leq \sqrt{2(1 - d_{\text{inf}}(p, \mathfrak{H}))}$ (see Figure 10). (Actually, every non-algebraic type in IHS has something like this property.) Peripheral sets and types are defined as a generalization of this property.

Definition 4.5.8. Given a model \mathfrak{M} , for any $p \in S_x(\mathfrak{M})$, the *degree of non-algebraicity* of p , written $\text{nalg}(p)$, is $\sup\{\varepsilon > 0 : p \notin \mathfrak{M}^{<\varepsilon}\}$, where $\sup \emptyset = 0$. For a type $p \in S_x(A)$, with A not necessarily a model, we define $\text{nalg}(p)$ to be $\sup\{\text{nalg}(q) : q \in S_x(\mathfrak{M}), q \supseteq p, \mathfrak{M} \supseteq A\}$. This agrees with the other definition when A is a model.

A \mathfrak{M} -definable set $D(x)$ is ε -peripheral (over \mathfrak{M}) if the set $\llbracket D \rrbracket \setminus \mathfrak{M}^{<\varepsilon}$ is a singleton, $\{p\}$, such that there is a modulus $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ for which we have

$$d(p, q) \leq \alpha([\text{nalg}(p) \dot{-} \text{nalg}(q)] + D(q)).$$

In such a case, we say that D points to p .

An A -definable set $D(x)$ is *strongly ε -peripheral* if it is ε -peripheral over every model

$\mathfrak{M} \supseteq A$.

A type p is (*strongly*) ε -peripheral if it is pointed to by a (strongly) ε -peripheral set. \triangleleft

The concept of peripheral set can be formalized in the context of arbitrary sets of parameters, and, likewise, the concept of peripheral types can be formalized directly, without relying on an ambient definable set, but we have not bothered to formalized these here.

Note that unlike strongly minimal sets and orthonormalizable sets, there isn't a clear pregeometry on the set of realizations of a strongly peripheral type.

In discrete logic, a strongly minimal type is a non-algebraic set that is isolated by a clopen condition (some definable set) and a simple closed condition (non-algebraicity). A strongly peripheral type, likewise, is ' d -atomized' by a definable set and some simple semi-continuous information (degree of non-algebraicity). Note that if $p \in S_1(\mathfrak{M})$ is ε -peripheral, then it is d -atomic-in- $S_1(\mathfrak{M}) \setminus \mathfrak{M}^{<\varepsilon}$, but the converse does not hold (see Counterexample C.2.8).

Like minimal sets, peripheral sets always exist in ω -stable theories.

Proposition 4.5.9. *Let T be a countable ω -stable theory. For any separable model $\mathfrak{M} \models T$, any non-empty open $U \subseteq S_1(\mathfrak{M})$, and any interval (δ, γ) with $0 < \delta < \gamma$ such that $U \setminus \mathfrak{M}^{<\eta}$ is non-empty for any $\eta \in (\delta, \gamma)$, there exists an ε -peripheral type $p \in U$ such that $\varepsilon \in (\delta, \gamma)$.*

Proof. Let $f(p) = d_{\text{inf}}(p, \mathfrak{M})$. By Lemma 2.5.9, there exists some $\eta_0 \in (\delta, \frac{\delta+\gamma}{2})$ such that $\{f = \eta_0\} \subseteq \overline{\{f > \eta_0\}}$. Since T is ω -stable, d -atomic-in- $\{f \geq \eta_0\}$ types are dense in $\{f \geq \eta_0\}$. Let p_0 be a d -atomic-in- $\{f \geq \eta_0\}$ type in $\{f \geq \eta_0\} \cap U$ (which is non-empty

by assumption). Let $U_0 = U$.

At any stage $i < \omega$, given $\eta_i \in (\delta, \frac{\delta+\gamma}{2})$ such that $\{f = \eta_i\} \subseteq \overline{\{f > \eta_i\}}$ and $p_i \in U_i$, if $f(p_i) > \eta_i$, then stop and set $p = p_i$, otherwise we have that $p_i \in \{f = \eta_i\}$, and so in particular $d_{\text{inf}}(p_i, \{f > \eta_i\}) = 0$. Find $\zeta_i > 0$ small enough that $\zeta_i < 2^{-i}$ and $B_{\leq \zeta_i}(p_i) \subseteq U_i$. Since p_i is d -atomic-in- $\{f \geq \eta_i\}$, we have that $p_i \in \text{int}_{\{f \geq \eta_i\}} B_{\leq \zeta_i}(p_i)$. Let U_{i+1} be an open subset of $S_1(\mathfrak{M})$ such that $U_{i+1} \subseteq U_i$ and $U_{i+1} \cap \{f \geq \eta_i\} = \text{int}_{\{f \geq \eta_i\}} B_{\leq \zeta_i}(p_i)$. Since $p_i \in \overline{\{f > \eta_i\}}$, it must be the case that $\{f \geq \sigma\} \cap U_{i+1}$ is non-empty for some $\sigma > \eta_i$. Let η_{i+1} be some such σ with $\eta_{i+1} < \frac{\delta+\gamma}{2}$. Since T is ω -stable, d -atomic-in- $\{f > \eta_{i+1}\}$ types are dense in $\{f > \eta_{i+1}\}$, and we can find a d -atomic-in- $\{f > \eta_{i+1}\}$ type $p_{i+1} \in \{f > \eta_{i+1}\} \cap U_{i+1}$. Note that, by construction, $d(p_i, p_{i+1}) \leq \zeta_i < 2^{-i}$.

Case 1: The construction did not stop at any stage.

$\{p_i\}_{i < \omega}$ is a metric Cauchy sequence. Let p be its limit. By construction we have that $f(p) > \delta$ and $f(p) \leq \frac{\delta+\gamma}{2} < \gamma$, or in other words $f(p) \in (\delta, \gamma)$. Let $\varepsilon = f(p)$, and note that this is the limit of η_i as $i \rightarrow \infty$. Also note that for each $i < \omega$, we have that p is contained in an open subset of $\{f \geq \varepsilon\}$ with metric diameter less than 2^{-i} , namely $\{f \geq \varepsilon\} \cap U_{i+1}$. This implies that p is d -atomic-in- $\{f \geq \varepsilon\}$. Since T is ω -stable and therefore dictionaric, we can find a definable set D such that $D \cap \{f \geq \varepsilon\} = \{p\}$.

Now we need to argue that for every $\lambda > 0$, there exists a $\chi > 0$ such that if $q \in \llbracket D \rrbracket$ and $f(q) > f(p) - \chi$, then $d(p, q) < \lambda$. Find $i < \omega$ such that $2\zeta_i < \lambda$. Since $D \cap \bigcap_{j < \omega} \{f \geq \eta_j\} = \{p\} \subseteq U_{i+1}$, there must exist a $j < \omega$ such that $D \cap \{f \geq \eta_j\} \subseteq U_{i+1}$. We may assume that $j > i$. So since $\text{diam } U_{i+1} \leq 2\zeta_i < \lambda$, we have that if $f(q) > f(p) - \eta_j$, then $d(p, q) < \lambda$, so set $\chi = \eta_j$.

Case 2: The construction did stop at some stage.

p is d -atomic-in- $\{f \geq \eta_i\}$. Let $\varepsilon = f(p)$. Since T is ω -stable and therefore dictionaric,

we can find a definable set D such that $D \cap \{f \geq \eta_i\} = \{p\}$.

We need to argue that for every $\lambda > 0$, there exists a $\chi > 0$ such that if $q \in \llbracket D \rrbracket$ and $f(q) > f(p) - \chi$, then $d(p, q) < \lambda$. Set $\chi = \frac{\varepsilon - \eta_i}{2}$. Suppose that $q \in D$ and $f(q) > f(p) - \frac{\varepsilon - \eta_i}{2} = \frac{\varepsilon + \eta_i}{2} > \eta_i$. This implies that $q \in \{f \geq \eta_i\}$, but by construction this implies that $q = p$, and so $d(p, q) = 0 < \lambda$.

In either case, by Proposition A.1.8 part (ii), we can find a modulus $\beta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that for any $q \in \llbracket D \rrbracket$, we have $d(p, q) \leq \beta(f(p) \dot{-} f(q))$.

Now for any $q \in S_1(\mathfrak{M})$, there exists an $r \in \llbracket D \rrbracket$ such that $d(q, r) = d_{\text{inf}}(q, \llbracket D \rrbracket)$. This implies that $f(q) \geq f(r) - d_{\text{inf}}(q, \llbracket D \rrbracket)$, and so $f(p) \dot{-} f(q) \leq f(p) \dot{-} f(r) + d_{\text{inf}}(q, \llbracket D \rrbracket)$. So if we set $\alpha(x) = \beta(x) + x$, we get that $d(p, q) \leq d(p, r) + d(r, q) \leq \beta(f(p) \dot{-} f(r)) + D(q) \leq \alpha((f(p) \dot{-} f(q)) + D(q))$, so we have that D is a peripheral set pointing to p , as required. \square

Just as with minimal sets, over an approximately ω -saturated model peripheral sets are strongly peripheral.

Lemma 4.5.10. *For any closed formula $F(x)$ and any model \mathfrak{M} of a complete theory T , the set $\llbracket F \rrbracket_{\mathfrak{M}} \setminus \mathfrak{M}^{<\varepsilon}$ is non-empty if and only if $T \models \forall \bar{x} \exists y F(y) \wedge \bigwedge_{i < \omega} d(x_i, y) \geq \varepsilon$.*

Proof. (\Rightarrow) Since $\llbracket F \rrbracket_{\mathfrak{M}} \setminus \mathfrak{M}^{<\varepsilon}$ is non-empty, for any $\bar{a} \in \mathfrak{M}$, $\llbracket F \rrbracket_{\mathfrak{M}} \setminus \bar{a}^{<\varepsilon}$ is non-empty, which implies that $T \models \forall \bar{x} \exists y F(y) \wedge \bigwedge_{i < \omega} d(x_i, y) \geq \varepsilon$.

(\Leftarrow) If $T \models \forall \bar{x} \exists y F(y) \wedge \bigwedge_{i < \omega} d(x_i, y) \geq \varepsilon$, then by compactness $\llbracket F \rrbracket_{\mathfrak{M}} \setminus \mathfrak{M}^{<\varepsilon}$ is non-empty. \square

Lemma 4.5.11. *If \mathfrak{M} is an approximately ω -saturated structure and $D(x, \bar{a})$ is a definable set with \bar{a} some ω -tuple of parameters from \mathfrak{M} , then for any type $p(\bar{x}, \bar{a}) \in S_{\bar{x}}(\bar{a})$,*

with \bar{x} a tuple of length at most ω , and any $\varepsilon > 0$, there exists $\bar{b}\bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models p(\bar{c}, \bar{b})$ and $d_H(D(\bar{a}), D(\bar{b})) < \varepsilon$.

Proof. By uniform continuity, for any $\varepsilon > 0$, there exists an $n < \omega$ and a $\delta > 0$ such that if \bar{b} has $d(\bar{a} \upharpoonright n, \bar{b} \upharpoonright n) < \delta$ and $\bar{a} \equiv \bar{b}$, then $d_H(D(\bar{a}), D(\bar{b})) < \varepsilon$. The result follows from applying Proposition 2.5.1 part (ii) to the type $\text{tp}(\bar{b}_{\geq n}\bar{c}/\bar{b}_{< n})$. \square

Proposition 4.5.12. *If $D(x, \bar{a})$ is ε -peripheral over \mathfrak{M} and \mathfrak{M} is approximately ω -saturated, then $D(x, \bar{a})$ is strongly ε -peripheral.*

Proof. Suppose that there is some $\mathfrak{N} \succeq \mathfrak{M}$ such that D is not an ε -peripheral set. There are two ways this can happen, either $\llbracket D \rrbracket_{\mathfrak{N}} \setminus \mathfrak{N}^{< \varepsilon}$ is not a singleton (note that it must be non-empty), or it is a singleton $\{p\}$ and there is some $\delta > 0$ such that for every $\gamma > 0$, there exists a type $q \in \llbracket D \rrbracket_{\mathfrak{N}}$ with $d_{\text{inf}}(q, \mathfrak{N}) > d_{\text{inf}}(q, \mathfrak{M}) - \gamma$ but $d(p, q) > \delta$.

First assume that $\llbracket D \rrbracket_{\mathfrak{N}} \setminus \mathfrak{N}^{< \varepsilon}$ is not a singleton. Let p and q be two distinct types in it. Let $\varphi(x, \bar{b})$ be a restricted formula such that $\varphi(p, \bar{b}) = 0$ and $\varphi(q, \bar{b}) = 1$, with $\bar{b} \in \mathfrak{N}$. By Lemma 4.5.10, the type $\text{tp}(\bar{b}/\bar{a})$ contains the formula

$$\forall \bar{x} \exists y D(y, \bar{a}) \wedge \varphi(y, \bar{b}) = k \wedge \bigwedge_{i < \omega} d(x_i, y) \geq \varepsilon$$

for both $k < 2$. Let $r(\bar{b}, \bar{a})$ be $\text{tp}(\bar{b}/\bar{a})$.

Find $\delta > 0$ small enough that if $|\varphi(e, \bar{b}) - \varphi(f, \bar{b})| \geq 1$ then $d(e, f) > \delta$. Let α be the modulus witnessing that $D(x, \bar{a})$ is ε -peripheral. Find $\gamma > 0$ small enough that $\alpha(\gamma) < \frac{1}{2}\delta$. By Lemma 4.5.11, we can find $\bar{a}'\bar{c} \in \mathfrak{M}$ such that $\mathfrak{M} \models r(\bar{c}, \bar{a}')$ and $d_H(D(\bar{a}), D(\bar{a}')) < \gamma$. Since \bar{c} satisfies $r(\bar{x}, \bar{a}')$, by Lemma 4.5.10, we have that the sets $\llbracket D(x, \bar{a}') \rrbracket \setminus \mathfrak{M}^{< \varepsilon}$ are non-empty for both $k < 2$. Let s_k be a type in $\llbracket D(x, \bar{a}') \rrbracket \setminus \mathfrak{M}^{< \varepsilon}$

for both $k < 2$. Both s_k satisfy $\text{nalg}(s_k) \geq \varepsilon = \text{nalg}(t)$, where t is the unique type in $\llbracket D(x, \bar{a}) \rrbracket_{\mathfrak{M}} \setminus \mathfrak{M}^{<\varepsilon}$, so we have that $d(t, s_k) < \frac{1}{2}\delta$ for both $k < 2$, but by the triangle inequality this implies that $d(s_0, s_1) < \delta$, contradicting the fact that $\varphi(s_0, \bar{c}) = 0$ and $\varphi(s_1, \bar{c}) = 1$. Therefore $\llbracket D \rrbracket_{\mathfrak{N}} \setminus \mathfrak{N}^{<\varepsilon}$ must be a singleton for every $\mathfrak{N} \succeq \mathfrak{M}$.

Second assume that there is an elementary extension $\mathfrak{N} \succ \mathfrak{M}$ and a $\delta > 0$ such that for every $\gamma > 0$, there exists a types $q \in \llbracket D \rrbracket_{\mathfrak{M}}$ with $d_{\text{inf}}(q, \mathfrak{N}) > d_{\text{inf}}(q, \mathfrak{M}) - \gamma$ but $d(p, q) > \delta$. By the same argument as before this implies that there is some $\delta' > 0$ such that for every $\gamma > 0$, there are two types $s_0, s_1 \in \llbracket D \rrbracket_{\mathfrak{M}} \setminus \mathfrak{M}^{<\varepsilon-\gamma}$ with $d(s_0, s_1) > \delta'$, which also contradicts that $D(x, \bar{a})$ is ε -peripheral. \square

Corollary 4.5.13. *Every countable ω -stable theory has strongly peripheral types over some model.* \square

Although peripheral types were motivated by the specific case of types in IHS, at the moment is it unclear if this relationship generalizes to other inseparably categorical Banach theory.

Question 4.5.14. *If T is a Banach theory and p is a minimal wide type, does it follow that p is strongly peripheral?*

The construction of a peripheral type in Proposition 4.5.9 has two different outcomes. Either the construction continues for ω steps or at some point we find a type p and an $\varepsilon > 0$ such that $d_{\text{inf}}(p, \mathfrak{M}) > \varepsilon$ but p is d -atomic-in- $S_1(\mathfrak{M}) \setminus \mathfrak{M}^{<\varepsilon}$. A strongly peripheral type in this second category has the property that there is a $\delta > 0$ such that any forking extension q of p has $\text{nalg}(q) \leq \text{nalg}(p) - \delta$, i.e. there is a discrete jump in any forking extension. All know examples of types like this have strongly minimal quotients. If

this were always the case it would be helpful with the characterization of inseparably categorical theories.

Question 4.5.15. *If T is a countable ω -stable theory and p is a strongly peripheral type over T with the property that there is some $\delta > 0$ such that for any forking extension q of p , $\text{nalq}(q) \leq \text{nalq}(p) - \delta$, does there exist a quotient sort such that the image of p in this sort is strongly minimal?*

Chapter 5

Structures and Theories

In this chapter we will collect various facts about specific classes of structures and theories.

5.1 Ultrametric Structures

p -adic Banach spaces are a natural example of ultrametric structures (i.e. ultrametric metric structures). ℓ^∞ and c_0 spaces over p -adic fields are known to behave somewhat analogously to Hilbert spaces. In particular they have a good notion of orthonormal bases [Roo78, Thm. 5.16] and as such they are also inseparably categorical. This is not surprising as the unit balls of these spaces are in a precise sense the inverse limit of the sequence of structures $(\mathbb{Z}/p^n\mathbb{Z})^\omega$ as $n \rightarrow \omega$.

The relationship between ultrametric theories and theories with totally disconnected type spaces is summarized in the following theorem.

Theorem 5.1.1. *Let T be a countable theory, the following are equivalent:*

- (i) *For every $n < \omega$ and every parameter set A , $S_n(A)$ is totally disconnected (where we take a single point to be totally disconnected).*
- (ii) *For every finite parameter set \bar{a} , $S_1(\bar{a})$ is totally disconnected.*

(iii) For every $n < \omega$, $S_n(\emptyset)$ is totally disconnected.

(iv) The diagonal in $S_2(\emptyset)$ (i.e. $\llbracket d(x, y) = 0 \rrbracket$) has a basis of clopen neighborhoods.

(v) T is dictionaric, and there is a definable metric ρ , uniformly equivalent to d , such that the distance set, $\rho(T) = \{\rho(a, b) \mid a, b \in \mathfrak{M} \models T\}$, contains no neighborhood of 0.

(vi) T is dictionaric, and there is a definable ultrametric ρ , uniformly equivalent to d , such that $\rho(T) \subseteq \{0\} \cup \{2^{-i} \mid i < \omega\}$.

Proof. (i) \Rightarrow (iii). This is immediate.

(iii) \Rightarrow (i). Assume that some $S_n(A)$ has a nondegenerate continuum C . Let $p, q \in C$ be distinct types. By compactness there must exist a restricted formula $\varphi(\bar{x}; \bar{a})$ and a finite parameter set $\bar{a} \in A$ such that $p(\bar{x}) \models \varphi(\bar{x}; \bar{a}) = 0$ and $q(\bar{x}) \models \varphi(\bar{x}; \bar{a}) = 1$. This implies that $p \upharpoonright \bar{a}$ and $q \upharpoonright \bar{a}$ are still distinct types, so since the natural projection $\pi : S_n(A) \rightarrow S_n(\bar{a})$ is continuous, $\pi(C)$ must be a nondegenerate continuum, thus $S_n(\bar{a}) \subseteq S_{n+|\bar{a}|}(\emptyset)$ fails to be totally disconnected.

(iii) \Rightarrow (ii). Each $S_1(\bar{a})$ is a subspace of $S_{1+|\bar{a}|}(\emptyset)$, so this is immediate as well.

(ii) \Rightarrow (iii). If $S_1(\emptyset)$ is not totally disconnected, then we're done.

Let n be the first $n < \omega$ such that $S_n(\emptyset)$ is totally disconnected but $S_{n+1}(\emptyset)$ is not totally disconnected. Consider the projection $\pi : S_{n+1}(\emptyset) \rightarrow S_n(\emptyset)$. Since $S_n(\emptyset)$ is totally disconnected, any continuum in $S_{n+1}(\emptyset)$ must be contained in a single fiber. This fiber is isomorphic to $S_1(\bar{a})$ for some parameter set \bar{a} , so $S_1(\bar{a})$ is not totally disconnected.

(iii) \Rightarrow (iv). This is immediate.

(iv) \Rightarrow (vi). Let $\{\varphi_i\}_{i < \omega}$ be an enumeration of $\{0, 1\}$ -valued restricted formulas of two variables corresponding to indicator functions of clopen neighborhoods of the diagonal

in $S_2(\emptyset)$ (note that there are only countably many of them since every $\{0, 1\}$ -valued definable predicate is equal to some $\{0, 1\}$ -valued restricted formula). Define ρ by

$$\rho(x, y) = \sup_{i < \omega} 2^{-i} \sup_z |\varphi_i(x, z) - \varphi_i(y, z)|.$$

This is a definable predicate and so is continuous on $S_2(\emptyset)$. It is an ultra-pseudo-metric since it is the supremum of a family of ultra-pseudo-metrics. It also clearly only takes on values in the set $\{0\} \cup \{2^{-i} | i < \omega\}$.

Since it vanishes on the diagonal, it must be uniformly dominated by d . Let p be a 2-type not on the diagonal. There is some clopen neighborhood Q of the diagonal which does not contain p . Let φ_i be its indicator function. We then have $\rho^p \geq 2^{-i}$, so ρ does not vanish anywhere besides the diagonal, so by compactness it must be uniformly equivalent to d .

(vi) \Rightarrow (v). This is immediate.

(v) \Rightarrow (iii). For any $p \in U \subseteq S_n(\emptyset)$, let D be a definable set such that $p \in D \subseteq U$. Note that a set definable relative to d is definable relative to any metric uniformly equivalent to d . By compactness there is an $\varepsilon > 0$ small enough that $D^{\rho \leq \varepsilon} \subseteq U$. Since there are arbitrarily small gaps in the distance set we can find $0 < \delta < \varepsilon$ such that $\delta \notin \rho(T)$, and we have that $D^{\rho \leq \varepsilon} = D^{\rho < \varepsilon}$ is a clopen set. Therefore $S_n(\emptyset)$ has a basis of clopen sets and is totally disconnected. \square

Lemma 5.1.2. *If X is a metrically separable ultrametric space, then its distance set, $d(X) = \{d(x, y) | x, y \in X\}$, is countable.*

Proof. Let $\{x_i\}_{i < \omega}$ be a countable dense subset of X . There are only countably many distances $d(x_i, x_j)$. By the ultrametric inequality if we choose x_i close enough to y and

x_j close enough to z , $d(y, z) = d(x_i, x_j)$. \square

Lemma 5.1.3. *If T is an ultrametric theory, then the underlying metric of each $S_n(A)$ is an ultrametric. (Recall that the metric on n -tuples is defined as the maximum of the componentwise distances.)*

Proof. Let $p, q, r \in S_n(A)$. We can find a model \mathfrak{M} containing n -tuples $\bar{a} \models p$, $\bar{b} \models q$, and $\bar{c} \models r$ such that $d(\bar{a}, \bar{b}) = d(p, q)$ and $d(\bar{b}, \bar{c}) = d(q, r)$. Then $d(p, r) \leq d(\bar{a}, \bar{c}) \leq d(\bar{a}, \bar{b}) \uparrow d(\bar{b}, \bar{c}) = d(p, q) \uparrow d(q, r)$. \square

Note that Lemma 5.1.3 isn't entirely trivial in the sense that not all metric properties of models of theories transfer directly to the type spaces of that theory.

Corollary 5.1.4. *If T is an ultrametric theory whose distance set,*

$$d(T) := \{d(a, b) \mid a, b \in \mathfrak{M} \models T\},$$

is somewhere dense, then its type space $S_2(T)$ is not small (and in particular T is not ω -stable).

Proof. Assume that $d(T)$ is dense in some interval (r, s) . Assume without loss that $0 < r$. Since $d(T)$ is dense in (r, s) , for every $t \in (r, s)$ and every $\varepsilon > 0$ we can find $a, b \in \mathfrak{M}$ for some model such that $|t - d(a, b)| < \varepsilon$, so by compactness there exists $c, e \in \mathfrak{N}$ such that $d(c, e) = t$. This means that the set $[|d(x, y) - t| = 0]$ is non-empty in $S_2(\emptyset)$ for every $t \in (r, s)$. Let p and q be 2-types such that $d^p = u > v = d^q$. Let $ab \models p$ and $ce \models q$ in some model. If $d(ab, ce) < v$, then we have that $d(a, b) \leq d(a, c) \uparrow d(c, e) \uparrow d(e, b) \leq v$, which is a contradiction, so we have $d(ab, ce) \geq v$.

Since $r > 0$, if we pick a type $p_t \in [|d(x, y) - t| = 0]$ for each $t \in (r, s)$, the set $\{p_r | t \in (r, s)\}$ will be uncountable and $(> r)$ -separated, so $S_2(T)$ is not small. \square

It is possible for an ultrametric theory T with somewhere dense distance set to have $S_1(T)$ be a single point (even if we require that T be superstable). for T to be superstable.

Corollary 5.1.5. *Every totally transcendental ultrametric theory T has totally disconnected type spaces.*

Proof. Every countable reduct T_0 of T is ω -stable, so by Corollary 5.1.4, $d(T_0)$ is nowhere dense in $[0, 1]$, so Theorem 5.1.1 applies and T_0 has totally disconnected type spaces. Since this is true for every countable reduct, $\{0, 1\}$ -valued formulas are logically complete in the full theory and T has totally disconnected type spaces as well. \square

As a word of warning, not all ultrametric theories have totally disconnected or even dictionaric type spaces.

Corollary 5.1.6. *Every theory T with totally disconnected type spaces is bi-interpretable with a many-sorted discrete theory T_{dis} .*

Proof. Since T has totally disconnected type spaces, it is dictionaric. By Theorem 5.1.1, we can find an ultrametric ρ uniformly equivalent to d such that $\rho(T) \subseteq \{0\} \cup \{2^{-i} | i < \omega\}$. This means that we can define a sequence of $\{0, 1\}$ -valued equivalence relations $E_k(x, y) = 2^{k+1}(2^{-k} \dot{-} \rho(x, y)) \downarrow 1$.

Each imaginary $T_k = T/E_k$ is metrically discrete and dictionaric by Proposition 3.3.1, so for every parameter set A , $S_n((T/E_k)_A)$ is totally disconnected and metrically discrete (the distance set of a type space is always a subset of the distance set of the

theory). If we form a many-sorted theory $\bigsqcup_{k < \omega} T/E_k$ with all the definable relations between different imaginaries, all of the type spaces will still be totally disconnected and metrically discrete, because for any finite set $T/E_{k_0}, \dots, T/E_{k_{n-1}}$, the largest k_i theory has surjectively defined maps to the others, so the mixed-sort type spaces are discrete quotients of some $S_n(T/E_k)$ and are thus totally disconnected and metrically discrete.

To see that this is a bi-interpretation note that the discrete theory has an ω -ary imaginary consisting of the direct limit $\lim_{\rightarrow} T/E_k$; specifically we can consider the imaginary sort $\prod_{k < \omega} T/E_k$, and then note that the set

$$D = \{\alpha \in \prod_{k < \omega} T/E_k : (\forall k < \omega) \alpha(k) E_k \alpha(k+1)\}$$

is definable. To see that D is definable, note that for each $k < \omega$, the set

$$D_k = \{\alpha \in \prod_{k < \omega} T/E_\ell : (\forall \ell < k) \alpha(\ell) E_\ell \alpha(\ell+1)\}$$

is clopen and satisfies $D \subseteq D_k \subseteq D^{<2^{-k}}$. D is identical to the home sort of the original theory. \square

Now that we know that ω -stable ultrametric theories are discrete theories in disguise, a Baldwin-Lachlan characterization is immediate.

Corollary 5.1.7. *If T is an ultrametric theory or has totally disconnected type spaces, then T is inseparably categorical if and only if it is ω -stable and has no imaginary Vaughtian pairs.*

5.2 Locally Compact Structures

In this section we will show that ω -stable theories with a locally compact model have strongly minimal sets, and at the end we will also discuss some points related to d -finiteness and raise some related questions. We should note that having a locally compact model is not equivalent to having every model locally compact.

Proposition 5.2.1.

- (i) *If T is a theory with a non-compact, locally compact model \mathfrak{M} , and $S_1(\mathfrak{M})$ is CB-analyzable (in particular if T is ω -stable or $S_1(\mathfrak{M})$ is small), then pre-minimal types are dense in $S_1(\mathfrak{M}) \setminus \mathfrak{M}$ (which is closed).*
- (ii) *If T is a totally transcendental theory with non-compact models such that every model is locally compact, then strongly minimal global types are dense among non-algebraic global types.*

Proof. (i) First we need to see that $S_1(\mathfrak{M}) \setminus \mathfrak{M}$ is closed. For every $a \in \mathfrak{M}$ there is a $\varepsilon > 0$ such that $B_{\leq 2\varepsilon}(a)$ is compact. This implies that $B_{< \varepsilon}(a)$ is algebraic over a by Lemma 2.3.62. Therefore in particular $B_{< \varepsilon}(a) \cap \mathfrak{M} = B_{< \varepsilon}(a) \subseteq S_1(\mathfrak{M})$. Therefore \mathfrak{M} as a subset of $S_1(\mathfrak{M})$ is a union of open sets and is itself open. Since $S_1(\mathfrak{M})$ is CB-analyzable, d -atomic-in- $S_1(\mathfrak{M}) \setminus \mathfrak{M}$ types are dense in $S_1(\mathfrak{M}) \setminus \mathfrak{M}$. Let p be any such type. Since $S_1(\mathfrak{M})$ is CB-analyzable, it is dictionary, so let $D \subseteq S_1(\mathfrak{M})$ be a definable set such that $D \cap (S_1(\mathfrak{M}) \setminus \mathfrak{M}) = \{p\}$. Then D is a minimal set: If $F, G \subseteq D$ are two M -zerosets, at most one of them can contain p , so at most one of them can be non-algebraic. Since this is true of any type that is d -atomic-in- $S_1(\mathfrak{M}) \setminus \mathfrak{M}$, we have that pre-minimal types are dense in $S_1(\mathfrak{M}) \setminus \mathfrak{M}$.

(ii) This is immediate from the fact that pre-minimal types over \aleph_0 -saturated structures are strongly minimal. \square

Corollary 5.2.2. *If T is a theory with a locally compact model, then T is inseparably categorical if and only if it is ω -stable and has no Vaughtian pairs.*

Given where Section 4.4.3 leaves off in attempting to reproduce in continuous logic the Baldwin-Lachlan theorem on the number of separable models, an exceedingly natural question is to wonder when the assumption in Corollary 4.4.9 holds. [BYU07] does not establish any conditions under which d -finite types can be found in abundance.

It is even trivial to construct an example of an ω -stable theory in which d -finite types do not occur in the home sort (the theory of an ω -product sort, for instance). This case, however, is sometimes covered by passing to imaginaries (which is part of the definition of ‘enough d -finite types’ given in [BYU07]), but it is easy to blame the lack of local compactness for this kind of issue in this case.

Local compactness is not enough to ensure a good selection of d -finite types, even in superstable theories. The theory presented in Example 2.2.1, for instance, is a superstable theory whose models are all locally compact but which has non- d -finite¹ types. Over a model it has no non- d -finite types, but it is not too hard to modify the example to have non- d -finite types over any set of parameters. If we take infinitely many disjoint copies of a model of Example 2.2.1, for instance, with no relationship between them, then such a theory would likely fail to have enough d -finite types (in the technical sense of [BYU07, Def. 2.7]).

This leaves a very narrow window in which we might hope to always have many d -finite types at hand.

¹Or should it be d -infinite, or perhaps in- d -finite?

Question 5.2.3. *If T is an ω -stable theory whose models are all locally compact, is it true that every type $p \in S_n(A)$ is d -finite, for any finite n and any set of parameters A ? Uniformly d -finite? What if T has uniformly locally compact models?*

Another obvious, more pointed question is this:

Question 5.2.4. *Does an inseparably categorical theory with locally compact models always have a strongly minimal set over some finite set of parameters in its prime model?*

5.3 Strongly Minimal Theories

In this section we will characterize those strongly minimal theories that do not interpret an infinite discrete structure (the essentially continuous strongly minimal theories), and also present a general characterization of strongly minimal groups, identifying the essentially continuous among them.

5.3.1 Strongly Minimal Theories That Do Not Interpret Infinite Discrete Theories

Obviously if \mathfrak{M} is a discrete strongly minimal structure and X is some compact metric space with a transitive automorphism group, then $\mathfrak{M} \times X$ is strongly minimal, but this example is in some sense trivial, as \mathfrak{M} and $\mathfrak{M} \times X$ are bi-interpretable. In [Noq17], Noquez raised the question of whether or not there are any non-trivially continuous examples of strongly minimal theories. Here is a minimally non-trivial continuous example, giving a positive answer:

Example 5.3.1. Consider \mathbb{R} with the metric $d(x, y) = \frac{|x-y|}{1+|x-y|}$. $\text{Th}(\mathbb{R}, +)$ is strongly minimal but does not interpret any infinite discrete structure.

Verification. We will verify these statements in Theorems 5.3.3 and 5.3.22. \square

A similar argument shows that the same is true of $(\mathbb{R}^n, +)$ for any $n < \omega$. Similarly, one can show that the underlying additive group of a p -adic field, \mathbb{Q}_p , with the appropriate metric, is strongly minimal, although it does have a discrete strongly minimal imaginary.

It turns out that, in the context of strongly minimal theories, this condition of being unable to interpret an infinite discrete structure has a very tight topological characterization in terms of models and also is very constraining in the special case of strongly minimal groups.

Definition 5.3.2. A continuous theory T is *essentially continuous* if it does not interpret an infinite discrete structure.² \triangleleft

The goal of the rest of this subsection is to prove the following characterization of essentially continuous strongly minimal theories.

Theorem 5.3.3. *Let T be a continuous strongly minimal theory. T is essentially continuous if and only if some model of T has a non-compact connected component.*

Proof of \Rightarrow . We will prove the contrapositive. Assume that every model of T has compact connected components (although note for later that we only actually use the fact that connected components of generic elements are compact). For any $\varepsilon > 0$, we will let E_ε be the equivalence relation on models of T induced by $aE_\varepsilon b$ if and only if there

²This term is somewhat folkloric. The only reference containing it we could find is [EG12].

exists a sequence c_0, \dots, c_n with $c_0 = a$, $c_n = b$ and $d(c_i, c_{i+1}) < \varepsilon$ for every $i < n$. We want to show the following.

- (*) There is a $\xi > 0$ such that the equivalence relation E_ξ is definable (by a $\{0, 1\}$ -valued formula) and has infinitely many equivalence classes in any model of T .

From which it is immediate that T is not essentially continuous, as the imaginary quotient by this equivalence relation is infinite and discrete.

Let \mathfrak{M} be some model of T , and let p be the generic type over \mathfrak{M} .

Claim: There is some $\varepsilon > 0$ such that for any $a \in \mathfrak{C}$ with $a \models p$, the E_ε -equivalence class of a is compact.

Proof of claim. Fix $a \in \mathfrak{C}$ with $a \models p$. Let F be the connected component of a . By assumption F is compact, and by Lemma 4.1.7, there is some $\delta > 0$ such that every closed δ -ball in any model of this theory is compact. Let b_0, \dots, b_{n-1} be elements of F chosen so that $U = \bigcup_{i < n} B_{<\delta}(b_i)$ covers F . Since $\text{cl } U \subseteq \bigcup_{i < n} B_{\leq\delta}(b_i)$ is compact and F is a connected component of it (and in particular is closed), we have that there is some clopen-in- $\text{cl } U$ set $Q \subseteq U$ such that $F \subseteq Q$. Since Q is relatively closed in $\text{cl } U$, it is closed, and since Q is relatively open in U , it is open, so it is actually clopen. Moreover, there is some $\varepsilon > 0$ such that $Q^{\leq\varepsilon}$ is disjoint from the complement of Q , since $\text{cl } U$ is compact. Clearly we have that the E_ε -equivalence class of a is contained in Q . It's not hard to see that an E_ε -equivalence class must be closed (for any $\varepsilon > 0$), so we have that it is compact. □_{claim}

Fix $\varepsilon > 0$ as in the claim. Now note that for any $b \models p$, the E_ε -equivalence class of b must be isometric to the E_ε -equivalence class of a , since there is an automorphism of the monster model taking b to a . By the compactness of the equivalence class, there exists

a natural number n such that if $aE_\varepsilon c$, then there exists a sequence e_0, \dots, e_n such that $a = e_0$, $c = e_n$, and $d(e_i, e_{i+1}) < \varepsilon$ for each $i < n$ (even if we only need a shorter chain, we can just use repetitions). This fact is uniformly true for any $b \models p$. By compactness, there is a $\xi > 0$ with $\xi < \varepsilon$ such that this statement is still true and such that for realizations of p , the E_ξ -equivalence classes are the same as the E_ε -equivalence classes. By construction, we have that if $aE_\xi b$ and $\neg aE_\xi c$, then $d(b, c) \geq \varepsilon$.

Now consider the formulas

$$\begin{aligned}\lambda(x, y) &:= \inf_{z_1, \dots, z_n} \max_{i \leq n} d(z_i, z_{i+1}) \div \varepsilon, \\ \eta(x, y) &:= \sup_{z_1, \dots, z_n} \min_{i \leq n} \xi \div d(z_i, z_{i+1}), \text{ and} \\ \chi(x) &:= \inf_y \max\{\lambda(x, y), \eta(x, y)\},\end{aligned}$$

where z_0 is understood to mean x and z_{n+1} is understood to mean y (note the $+1$, but also note that the $+1$ is not actually essential, rather it just makes the formulas λ and η easier to write).

To unpack what these mean, given a and b , $\lambda(a, b) > 0$ implies that for any sequence c_0, \dots, c_{n+1} with $c_0 = a$ and $c_{n+1} = b$, for some $i \leq n$, we have $d(c_i, c_{i+1}) > \varepsilon$ (although note that $\lambda(a, b) > 0$ is actually slightly stronger than this statement in insufficiently saturated models). On the other hand, $\eta(a, b) > 0$ means that there is a chain c_0, \dots, c_{n+1} with $c_0 = a$, $c_{n+1} = b$, and $d(c_i, c_{i+1}) < \xi$ for every $i \leq n$. So, putting it together, $\chi(a) > 0$ implies that for every b , one of these two conditions holds. By the above statements regarding any a satisfying p , we have that $\chi(a) > 0$.

So we have that $\llbracket \chi(x) > 0 \rrbracket$ is an open neighborhood of p in $S_1(\mathfrak{M})$, so in particular,

$\llbracket \chi(x) = 0 \rrbracket$ is compact (note, also, that χ always takes on non-negative values).

Claim: There is a natural number k , such that for any a and b , if $aE_\xi b$, then this is witnessed by a chain of length no more than k .

Proof of claim. Let X be a finite subset of $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$ such that $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket \subseteq X^{<\xi}$. For any $a \in \llbracket \chi(\mathfrak{C}) = 0 \rrbracket$, let $F_a \subseteq \llbracket \chi(\mathfrak{C}) = 0 \rrbracket$ be the E_ξ -equivalence class of a computed inside $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$, i.e. the set of all elements b of $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$ satisfying $aE_\xi b$ with a witnessing chain contained entirely in $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$. Note that by construction, $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket = \bigcup_{x \in X} F_x$. Also note that each F_a is closed, and therefore compact (as a set of elements of \mathfrak{C}).

For each $x \in X$, we have that

$$F_x \subseteq B_{<\xi}(x) \cup (B_{<\xi}(x))^{<\xi} \cup ((B_{<\xi}(x))^{<\xi})^{<\xi} \cup \dots$$

By the compactness of F_x , some finite initial segment of this is sufficient to cover $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$. Let ℓ be the longest length of one of these initial segments. Let $k = \ell + 1 + n$.

To see that this k is sufficient, we need to consider three cases.

1. $\chi(a) > 0$. In this case, by our previous discussion of $\chi(x)$, we know that there is a chain of length no more than $n + 1$ witnessing $aE_\xi b$, which is no more than k .
2. $\chi(a) = 0$ and the chain witnessing $aE_\xi b$ is entirely contained in $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$. In this case, we have that $b \in F_a$. There is some $x \in X$ with $d(a, x) < \xi$, so by the definition of ℓ we get a chain of length $1 + \ell$ witnessing that $aE_\xi b$.
3. $\chi(a) = 0$ and the chain witnessing $aE_\xi b$ is not entirely contained in $\llbracket \chi(\mathfrak{C}) = 0 \rrbracket$. In this case, let c_0, \dots, c_m , with $c_0 = a$ and $c_m = b$, be the chain witnessing that

$aE_\xi b$. Let c_i be the first such that $\chi(c_i) = 0$ but $\chi(c_{i+1}) > 0$. By construction, we can find a chain of length no more than ℓ witnessing that $aE_\xi c_i$ and a chain of length no more than n witnessing that $c_{i+1}E_\xi b$. Concatenating these chains gives a witnessing chain of length no more than $\ell + 1 + n$, which is k . \square_{claim}

Now, since this is true for all a and b , by compactness, there must actually be some $\delta > 0$ with $\delta < \xi$ such that for any a and b , with $aE_\xi b$, there is a chain c_0, \dots, c_k , with $c_0 = a$, $c_k = b$, and $d(c_i, c_{i+1}) < \delta$ for all $i < k$. We have just established that for any a and b , if $aE_\xi b$, then $\eta(a, b) > \xi - \delta$, but if $\neg aE_\xi b$, then $\eta(x, y) = 0$. So consider the formula

$$E(x, y) := \min \left\{ \frac{(\xi - \delta) \div \eta(x, y)}{\xi - \delta}, 1 \right\}.$$

We now have that this $E(x, y)$ is $\{0, 1\}$ -valued and has $E(x, y) = 0$ if and only if $xE_\xi y$, as required. To see that this equivalence relation has infinitely many equivalence classes in every model, note that if it had finitely many equivalence classes then that model would itself be compact, since its equivalence classes are compact, which is a contradiction. \square

It is not hard to show that if some model of T has non-compact connected components, then every non-prime model has non-compact connected components. The set E in the example mentioned in the text after Counterexample C.2.1 (i.e. the definable subset of \mathbb{R}) shows that this can fail for the prime model.

In order to prove the other direction of Theorem 5.3.3 we will need to collect some lemmas.

Lemma 5.3.4. *If T is a strongly minimal theory, $\mathfrak{M} \models T$, and for any a realizing the generic type over \mathfrak{M} , the connected component of a (in the monster model) is compact, then no model of T has a non-compact connected component.*

Proof. Since the proof of Theorem 5.3.3 only uses that the connected components of generic elements are compact, we have that there is some $\xi > 0$ such that the equivalence relation E_ξ (as defined in that proof) has compact equivalence classes. For any a the connected component of a is contained in the E_ξ -equivalence class of a , and so since connected components are always closed, we have that the connected component of a is compact. \square

The following is a (weak) analog of a well known fact in discrete logic that given two strongly minimal sets in an uncountably categorical theory, there is a definable finite-to-finite correspondence between them.

Lemma 5.3.5. *Let T be a strongly minimal theory. Suppose that I is a strongly minimal imaginary. There exists a formula $\varphi(x, y)$ (using at most the parameters needed to define I), with x a variable in the sort I and y a variable in the home sort, such that for any generic $a \in I$, the zeroset of $\varphi(a, y)$ is algebraic and has each realization generic over \mathfrak{M} and for any generic b in the home sort, the zeroset of $\varphi(x, b)$ is algebraic and has each realization generic over \mathfrak{M} .*

Proof. Let A be the set of parameters needed to define I , and let \mathfrak{M} be a separable model of T containing A . Let c be an element of $I(\mathfrak{C})$ that is generic over \mathfrak{M} , and let $\mathfrak{N} \succ \mathfrak{M}$ be a model of T that is prime over $\mathfrak{M}b$. Because \mathfrak{N} is a proper elementary extension of \mathfrak{M} we must have that $H(\mathfrak{N})$ is strictly larger than $H(\mathfrak{M})$, where H is the home sort. Let b be some element of $H(\mathfrak{N}) \setminus H(\mathfrak{M})$. We necessarily have that b is generic over \mathfrak{M} , but $\text{tp}(c/\mathfrak{M}b)$ is d -atomic, which by strong minimality means that it must be algebraic. Let $\varphi_0(x, y)$ be a formula such that $\varphi_0(b, y)$ is the distance predicate of $\text{tp}(c/\mathfrak{M}b)$.

A symmetric argument gives that $\text{tp}(b/\mathfrak{M}c)$ is also algebraic and therefore d -atomic, so let $\varphi_1(x, y)$ be a formula such that $\varphi_1(x, c)$ is the distance predicate of $\text{tp}(b/\mathfrak{M}c)$. The formula $\varphi(x, y) = \max\{\varphi_0(x, y), \varphi_1(x, y)\}$ has the required properties. \square

Remark 5.3.6. As an aside, one might wonder whether or not two strongly minimal sets in an inseparably categorical theory always have a definable compact-to-compact correspondence. The answer is yes. In Lemma 5.3.5 (which, of course, generalizes to arbitrary inseparably categorical theories), the set $\llbracket \varphi(x, y) = 0 \rrbracket \cap F(x, y)$, where $F(x, y)$ is the partial type consisting of all types for which at least one of the projections is generic, is relatively definable in $F(x, y)$. By Proposition 2.4.4, we can find a definable set $D(x, y)$ such that $D(x, y) \cap F(x, y) = \llbracket \varphi(x, y) = 0 \rrbracket \cap F(x, y)$. By strong minimality, it is easy to show that $D(x, y)$ is a compact-to-compact relation between the two sets, but it may fail to be either total or surjective. By uniform local compactness, for sufficiently small $\varepsilon > 0$, $\llbracket D(x, y) \leq \varepsilon \rrbracket$ is still a compact-to-compact relation. By Lemma 2.5.9, we may choose ε so that $\llbracket D(x, y) \leq \varepsilon \rrbracket$ is definable. Since the interior of $\llbracket D(x, y) \leq \varepsilon \rrbracket$ contains the ‘rank 2’ type (type of a pair of mutually generic elements), we have (where A and B are the strongly minimal sets in question) that the sets $U = \{a \in A : \neg(\exists b \in B) D(a, b) \leq \varepsilon\}$ and $V = \{b \in B : \neg(\exists a \in A) D(a, b) \leq \varepsilon\}$ are both pre-compact, so we have that $\llbracket D(x, y) \leq \varepsilon \rrbracket \cup (\overline{A} \times \overline{B})$ is a definable compact-to-compact correspondence between A and B .

This is still true for any pair of non-orthogonal strongly minimal types in an arbitrary theory, but a minor technicality is that one would need to show that the type spaces corresponding to $A \times B$ are dictionaric.

The original correspondence is actually slightly nicer than merely definable. It has the property that for any generic $a \in A$, the set $\llbracket \varphi(a, y) = 0 \rrbracket$ is definable and likewise for any generic $b \in B$, as well as the property that the distance predicates of these definable sets are uniformly definable, i.e. there is a formula $\psi(x, y) = 0$ such that for any generic $a \in A$, $\psi(a, y) = 0$ is the distance predicate of $\llbracket \varphi(a, y) = 0 \rrbracket$ (with a similar statement for B). This additional niceness is not always possible. Consider \mathbb{R} as a metric space and the \mathbb{R} -definable subset $E = \{\pm \ln(1+n) : n < \omega\}$ (which is very similar to the definable set mentioned after Counterexample C.2.1). Although the strongly minimal types in the home sort H and E are non-orthogonal in that they are literally the same type, there is no compact-to-compact correspondence between H and E with this additional property. To see this, note that if a formula $\chi(a, y)$ is a distance predicate for any parameter a , then the function $a \mapsto \llbracket \chi(a, y) = 0 \rrbracket$ is continuous with regards to the Hausdorff metric. Assume that $\varphi(x, y) = 0$ is a definable compact-to-compact correspondence between H and E , where x is a variable of sort H and y is a variable of sort E . Then consider the set $\llbracket \varphi(x, 0) = 0 \rrbracket$. This must contain some $a \in H$. By continuity, we have that $\varphi(a', 0)$ must be 0 for all a' , but this implies that 0 is related to a non-compact set of a' 's in A , which is a contradiction.

In this example it can be arranged that $\llbracket \varphi(x, b) = 0 \rrbracket$ is uniformly definable as a function of b , but it is not too difficult to show that if we take $E \cup \mathbb{R}_{\geq 0}$ and $E \cup \mathbb{R}_{\leq 0}$ as our two strongly minimal sets, then neither direction can be taken to be uniformly definable.

Now we can finish the proof of the characterization.

Proof of Theorem 5.3.3 \Leftarrow . Let \mathfrak{M} be the prime model of T . Assume that some model $\mathfrak{N} \succeq \mathfrak{M}$ of T has a non-compact connected component C . If C does not contain an element that is generic over \mathfrak{M} , then by Lemma 5.3.4 there is a model of T that contains a non-compact connected component C' which contains an element that is generic over \mathfrak{M} , so we may assume without loss of generality that C contains an element generic over \mathfrak{M} . Let a be some such generic element.

Now assume that T is not essentially continuous. Let I be an imaginary sort of T that is infinite and discrete. Since T is inseparably categorical, there is a strongly minimal set $D \subseteq I$ (possibly over some larger set of parameters). Let J be the imaginary sort corresponding to D . Let $\mathfrak{N} \succeq \mathfrak{M}$ be a separable elementary extension containing the parameters necessary to define D . By Lemma 5.3.5, there is an \mathfrak{M} -formula $\varphi(x, y)$, with x a variable in J and y a variable in the home sort, such that for any a in $J(\mathfrak{C})$, generic over \mathfrak{N} , the set of b 's in the home sort for which $\varphi(a, b) = 0$ holds is compact and consists only of elements generic over \mathfrak{N} .

Fix $b \in H(\mathfrak{C})$ generic, and note that the connected component of b is non-compact. By compactness there must exist an $\varepsilon > 0$ such that for any generic a , if $\varphi(a, b) > 0$, then $\varphi(a, b) > \varepsilon$. So consider the formula $\psi(x, y) := \min\{\frac{1}{\varepsilon}\varphi(x, y), 1\}$. We have that for generic a , $\psi(a, b) \in \{0, 1\}$. This is part of the type of b , so for any generic b this must be true as well. By uniform continuity of $\psi(x, y)$, this implies that there is a $\delta > 0$ such that if generic $b, b' \in H(\mathfrak{C})$ satisfy $d(b, b') < \delta$, then $\psi(a, b) = \psi(a, b')$, but this implies that for any b' in the connected component of b , $\psi(a, b') = \varphi(a, b') = 0$, which is a contradiction. Therefore no such I can exist, and T is essentially continuous. \square

Corollary 5.3.7. *For any strongly minimal theory T , the following are equivalent:*

- T is essentially continuous.
- Some model of T has a non-compact connected component.
- Every generic element of a model of T has a non-compact connected component.
- T does not have a \emptyset -definable $\{0, 1\}$ -valued equivalence relation on its home sort with infinitely many equivalence classes.

The example given immediately after Counterexample C.2.1 (i.e. \mathbb{R} with the appropriate metric), as well as some of the examples given in Remark 5.3.6, show that the prime model may fail to have a non-compact connected component in an essentially continuous strongly minimal theory.

After all of this, it is natural to wonder about the kinds of (pre-)geometries that can occur in essentially continuous strongly minimal theories. There are a myriad of specific questions in this vein. It is possible that characterizing essentially continuous strongly minimal theories might be easier than characterizing discrete strongly minimal theories. A negative answer to either of the following two questions would be an indication that essentially continuous strongly minimal sets admit less complexity than discrete strongly minimal sets.

Question 5.3.8. *Is there an essentially continuous strongly minimal set with the geometry of an algebraically closed field?*

Question 5.3.9. *Is there an essentially continuous strongly minimal set with non-locally modular, flat pregeometry?*

A difficult-to-rigorize follow-up question would be this: Can a Hrushovski construction build an essentially continuous strongly minimal set?

Question 5.3.10. *Is there an essentially continuous strongly minimal set whose geometry is not isomorphic to the geometry of any discrete strongly minimal set?*

Is there a continuous, but not essentially continuous, strongly minimal theory with the same?

This last question becomes trivial on cardinality grounds if we consider pregeometries, rather than geometries.

5.3.2 Characterization of Strongly Minimal Groups

In this section we will give a topological group theoretic characterization of strongly minimal groups in continuous logic and identify which among them are essentially continuous.

Definition 5.3.11. A theory of groups is a theory T in a language containing a binary function symbol $x \cdot y$ (which we will freely write as concatenation), a unary function symbol x^{-1} , and a constant symbol e such that T implies that

- $\sup_{xyz} d(x(yz), (xy)z) = 0$,
- $\sup_x d(ex, x) = 0$, and
- $\sup_x d(x^{-1}x, e) = 0$.

A *group structure* is a model of a theory of groups. ◁

Note that despite appearances, **IHS** is not an example of a theory of groups. In fact, it is possible to show that **IHS** does not even interpret a non-compact group. If we had taken the whole Hilbert space as the structure, rather than just the unit ball, and modified the metric accordingly, then the result theory would be a theory of groups.

The following was originally shown in [BY10a, Prop. 3.13]. We have given a full proof here for completeness.

Proposition 5.3.12. *Let T be a theory of groups. There is a definable bi-invariant metric ρ that is uniformly equivalent to d (the original metric).*

Proof. Consider the formula $\rho(x, y) = \sup_{z,w} d(zxw, zyw)$. First note that the formula is clearly bi-invariant in the sense that $\rho(x, y) = \rho(uxv, uyv)$ for any x, y, u, v . Also note that it is clearly symmetric, non-negative, and satisfies $\rho(x, x) = 0$.

To verify the triangle inequality, consider a, b, c , and assume we are working a sufficiently saturated model (since the triangle inequality is first-order, this will imply the same for all models of the theory). For any f, g we have that $d(fag, fcg) \leq d(fag, fbg) + d(fbg, fcg) \leq \rho(a, b) + \rho(b, c)$, so the triangle inequality holds.

This establishes that $\rho(x, y)$ is a pseudo-metric. To establish that it is a metric we need to show that it is uniformly equivalent to d . Since it is a definable pseudo-metric, it is automatically uniformly dominated by d , so we only need to show that d is uniformly dominated by it. By compactness, it is enough to show that in any model \mathfrak{M} if $\rho(a, b) = 0$, then $d(a, b) = 0$. Let a and b in \mathfrak{M} have $\rho(a, b) = 0$. Pass to an elementary extension $\mathfrak{N} \succeq \mathfrak{M}$ containing c and f such that $d(caf, cbf) = 0$, so in particular $caf = cbf$. By the group axioms we have that $a = c^{-1}caf f^{-1} = c^{-1}cbf f^{-1} = b$, i.e. $d(a, b) = 0$. By elementarity, this is true in \mathfrak{M} as well, so we are done. \square

Proposition 5.3.12 seems to be in conflict with the well known fact that there are metrizable groups that do not admit a bi-invariant metric. There is a hidden assumption here, which is that the group operations are uniformly continuous, and not merely continuous.

Our proof that strongly minimal groups are Abelian in continuous logic is heavily based on the proof of the analogous statement in [Bue17, Cor. 3.5.5], originally proven in [Rei75].

One should think of the following condition as being a natural analog of having finite order.

Definition 5.3.13. Given a topological group G , we say that an element a has *pre-compact orbit* if the orbit of a , $a^{\mathbb{Z}}$, is pre-compact (i.e. has compact closure).³ \triangleleft

Lemma 5.3.14. *Let G be a topological group with a topology induced by the metric d . If $a \in G$ has pre-compact orbit, then for every $\varepsilon > 0$ there are arbitrarily large ℓ such that $d(a^\ell, e) < \varepsilon$.*

Proof. For any $\delta > 0$, the set $\bigcup_{n \in \mathbb{Z}} B_{<\delta}(a^n)$ covers the closure of the orbit of a . By compactness we can find a finite set $X \subset \mathbb{Z}$ such that $\bigcup_{n \in X} B_{<\delta}(a^n)$ covers the closure of the orbit of a , and therefore the orbit of a . This implies that for any $m < \omega$ there exists $n \in X$ and $k > n + m$ such that $d(a^n, a^k) < \delta$. Since we can do this for any $\delta > 0$ and by considering $a^k a^{-n}$, we have that there are arbitrarily large ℓ such that $d(a^\ell, e) < \varepsilon$. \square

Lemma 5.3.15. *Let G be a metric group with bi-invariant metric d . If all elements of G have pre-compact orbit and all elements of $G \setminus \{e\}$ are conjugate, then $|G| \leq 2$.*

Proof. We may assume that G has more than one element. Fix $a \in G \setminus \{e\}$. We need to show that $a^2 = e$. If $a^2 = e$, then we are done, so assume that $a^2 \neq e$.

³Elements of a topological group with this property are referred to in some literature as just ‘compact.’

For any $\varepsilon > 0$, find the smallest positive n_ε such that $d(a^{n_\varepsilon}, e) < \varepsilon$, which, note, is larger than 1 whenever $\varepsilon \leq d(a, e)$. Note that by the conjugacy condition and the bi-invariance of the metric, n_ε does not depend on the choice of a .

Claim: n_ε is either 1 or a prime number.

Proof of claim. Assume that n_ε is not 1 and is the composite number mk , with $m, k > 1$. By assumption, $a^m \neq e$, so since a^{n_ε} and a^m are conjugate, we can find b such that $b^{-1}a^{n_\varepsilon}b = a^m$. By the bi-invariance of the metric we have that $d(a^{n_\varepsilon}, e) = d(b^{-1}a^{n_\varepsilon}b, b^{-1}eb) = d(a^m, e) < \varepsilon$, contradicting the minimality of n_ε . \square_{claim}

Suppose that n_ε is not eventually 2 as $\varepsilon \rightarrow 0$. Since $a^{-1} \neq e$, we can find a $b \in G$ such that $b^{-1}ab = a^{-1}$. This implies that for any k , $b^{-k}ab^k = a^{(-1)^k}$. By our assumption, for any $\gamma > 0$, there is an $\varepsilon > 0$ with $\varepsilon < \delta$ such that n_ε is odd and such that if $d(c, e) < \varepsilon$, then $d(c^{-1}ac, a) < \gamma$. Now we have that $b^{-n_\varepsilon}ab^{n_\varepsilon} = a^{-1}$, since n_ε is odd, implying that $d(a, a^{-1}) < \gamma$. Since we can do this for any $\gamma > 0$, we have that $a = a^{-1}$, which is a contradiction. Therefore n_ε must be eventually 2 as $\varepsilon \rightarrow 0$, but this implies that in fact $a^2 = e$.

So, by the conjugacy condition, we have that every element a of G satisfies $a^2 = e$. It is not hard to show that this implies that G is Abelian, but this implies that every element's conjugacy class consists solely of itself, and hence $|G| = 2$, as required. \square

Lemma 5.3.16. *If G is a strongly minimal group and H is a type-definable, non-compact subgroup of G , then $H = G$.*

Proof. If H is a proper subgroup of G , then the cosets of H are also non-compact type-definable subsets of G . Furthermore, H is disjoint from its non-trivial cosets, which contradicts strong minimality. \square

Proposition 5.3.17. *Let T be a strongly minimal theory of groups. T is Abelian.*

Proof. For the sake of this proof we will write $A \setminus B$ to mean the collection of right cosets of A in the set B , i.e. $\{Ab : b \in B\}$, as opposed to the set theoretic relative complement, which is written $A \setminus B$ elsewhere in this thesis. The $A \setminus B$ notation does not occur in this proof, and the $A \setminus B$ notation does not occur outside of this proof.

Assume that T is not Abelian. By Proposition 5.3.12 we may assume that T has a bi-invariant metric. Let Z be the center of G (i.e. the set of all elements g satisfying $gh = hg$ for all $h \in G$). For any $g \in G$, let $C(g)$ be the centralizer of g (i.e. the set of all elements h satisfying $h^{-1}gh = g$). Note that the orbit of g is contained in $C(g)$. This is a type-definable group defined by the closed condition $d(h^{-1}gh, g) = 0$. For any $g \in G$ with $g \notin Z$, $C(g)$ is a type-definable proper subgroup of G , so by Lemma 5.3.16 it is compact, hence g has pre-compact orbit.

Claim: For any $g \in G$ with $g \notin Z$, the set of conjugates of g , $a^G := \{h^{-1}gh : h \in G\}$, is not pre-compact (i.e. has compact closure).

Proof of claim. We want to argue that there is a natural topological isomorphism between (the metric closure of) $C(g) \setminus G$ and (the metric closure of) a^G , where $C(g) \setminus G$ is made into a metric space by the Hausdorff metric on sets, which, by bi-invariance and the fact that $C(g) \setminus G$ is a set of cosets, is equal to $\rho(x, y) = \inf_{z \in C(g)} d(x, zy)$. Then we will show that $C(g) \setminus G$ is not pre-compact.

It is a basic algebraic fact that there is a natural bijection between $C(g) \setminus G$ and g^G . Explicitly, for any a and b , we have that if $a^{-1}ga = b^{-1}gb$, then $gab^{-1} = ab^{-1}g$, so $ab^{-1} \in C(g)$, implying that $C(g)b = C(g)ab^{-1}b = C(g)a$. This gives a natural function from g^G to $C(g) \setminus G$. To see that it is surjective, note that for any $C(g)a$, the element $a^{-1}ga$ maps to $C(g)a$. To see that it is injective, if $G(g)a = G(g)b$, then this implies

that $a = cb$ for some $c \in C(g)$, so we have that $a^{-1}ga = b^{-1}c^{-1}gcb = b^{-1}gb$, as required.

All we need to do is argue that this bijection is metrically uniformly continuous with metrically uniformly continuous inverse. I claim that for any $\varepsilon > 0$, there is a $\delta > 0$ such that if $d(ag, ga) < \delta$, then $d(a, C(g)) < \varepsilon$. This follows from compactness and the fact that $C(g)$ is an algebraic set. For any given $\varepsilon > 0$, find such a $\delta > 0$, and consider a and b such that $d(a^{-1}ga, b^{-1}gb) < \delta$. By bi-invariance, this implies that $d(gab^{-1}, ab^{-1}g) < \delta$, so we have that there is some $c \in C(g)$ with $d(ab^{-1}, c) < \varepsilon$. This implies that $d(a, cb) < \varepsilon$, so we have that the distance between $C(g)a$ and $C(g)b$ in the ρ -metric is less than ε . Since we can do this for any $\varepsilon > 0$, we have that the map is uniformly continuous.

Now we just need to show that the map has uniformly continuous inverse. For any $\varepsilon > 0$, find a $\delta > 0$ small enough that for any c if $d(a, b) < \delta$, then $d(a^{-1}ca, b^{-1}cb) < \varepsilon$. Assume that $C(g)a$ and $C(g)b$ have ρ -distance less than δ . By bi-invariance, this implies that there is some $c \in C(g)$ such that $d(a, cb) < \delta$. This implies, by the choice of ε , that $d(a^{-1}ga, b^{-1}c^{-1}gcb) < \varepsilon$, and so $d(a^{-1}ga, b^{-1}gb) < \varepsilon$. Since we can do this for any $\varepsilon > 0$, we have that the inverse of the bijection is uniformly continuous.

To show that $C(g)\backslash G$ is not pre-compact, first note that we have already established that $C(g)$ is compact. If $C(g)\backslash G$ were pre-compact, then this would imply that for any $\varepsilon > 0$ there is a finite set of elements $X \subseteq C(g)\backslash G$ such that $G \subseteq \bigcup_{x \in X} q^{-1}(x)^{<\varepsilon}$, where $q : G \rightarrow C(g)\backslash G$ is the natural quotient map. Since each pre-image $q^{-1}(x)$ is isometric to $C(g)$, this is enough to imply that G is compact, which contradicts our assumptions.

□_{claim}

Now we have that for any a and b not in Z , the metric closure of a^G is equal to the metric closure of b^G , because these sets are definable and non-compact, and so, by strong minimality, must have non-empty intersection in some elementary extension,

which implies in the elementary extension that there is some c such that $c^{-1}ac = b$. By passing to a sufficiently saturated elementary extension, we may assume that every pair a and b not in Z are conjugate. This implies that in G/Z any two non-identity elements are conjugate (because conjugacy commutes with group homomorphisms). We also still have that every element has pre-compact orbit (after endowing G/Z with the Hausdorff metric on cosets, as we did with $C(g)\backslash G$), so by Lemma 5.3.15 we have that $|G/Z| \leq 2$. This implies that Z must be non-compact (otherwise G would be compact), but then by Lemma 5.3.16 we have that $G = Z$, which contradicts our assumption that G was not Abelian. \square

The corresponding result in [Bue17]—namely, Proposition 3.5.2—is the statement that any infinite ω -stable group has an infinite definable Abelian subgroup. The proof as stated does not generalize to continuous ω -stable groups, as we now have many different Morley ranks rather than a single one.

Question 5.3.18. *To what extent does [Bue17, Prop. 3.5.2] generalize to arbitrary ω -stable groups in continuous logic? Is it true that every non-compact ω -stable group has a definable non-compact Abelian subgroup? Does it follow if we assume that models of the theory are locally compact?*

It is known that type-definable subgroups in ω -stable theories are definable [BY10a], so it is not preposterous to hope that there may be nice definable subgroups in arbitrary ω -stable continuous groups.

We would like to take a moment to negatively answer a question given in the introduction of [BY10a], which to our knowledge does not have an answer at least in the published literature. In [BY10a], Ben Yaacov asked whether or not type-definable

groups in stable theories are always intersections of some family of definable groups. It turns out that, unlike in discrete logic, superstability is not enough to ensure that type-definable groups are the intersections of families of definable groups in continuous logic. It is not hard to show that the structure $(\mathbb{Q}, +, \cos, \sin)^4$ with the discrete metric has a weakly minimal and therefore superstable theory but also has the property that the type-definable group given by $\{\cos(x) = 1, \sin(x) = 0\}$ is not, in sufficiently saturated elementary extensions, the intersection of any family of definable subgroups. In fact, the only definable subgroups in this theory are the trivial subgroup and the group itself. This is in line with the general phenomenon, discussed in Chapter 2, that, while ω -stable theories are well behaved with regards to definable sets, even superstable theories are not.

For a full characterization of essentially continuous strongly minimal groups, we will need the following significant fact from the theory of topological groups.

Fact 5.3.19 ([MHS77, Thm. 25]). *Every locally compact Hausdorff Abelian group has an open subgroup topologically isomorphic to $\mathbb{R}^n \times K$ for some compact group K and some non-negative integer n .*

We will also need the following well known algebraic fact regarding Abelian groups.

Fact 5.3.20. *An Abelian group G is divisible if and only if it is an injective object in the category of Abelian groups, i.e. if and only if for any group H with subgroup $F \subseteq H$ and any homomorphism $f : F \rightarrow G$, there exists a homomorphism $g : H \rightarrow G$ extending f .*

Facts 5.3.19 and 5.3.20 imply that every locally compact Hausdorff Abelian group can be written in the form $\mathbb{R}^n \times G$, with n a non-negative integer and G a totally disconnected

⁴One should think of this as $(\mathbb{Q}, +)$ with a non-trivial homomorphism to the circle group S^1 .

locally compact Hausdorff group. (To see this note that any homomorphism extending the inclusion map of \mathbb{R}^n into $\mathbb{R}^n \times K$ to the entire group must be continuous, and G can be taken to be the kernel of this extension.)

Lemma 5.3.21. *If G is a group structure with a bi-invariant metric and H is a definable subgroup of G , then the imaginary sort given by quotienting by the definable pseudo-metric $\rho(x, y) = \inf_{z \in H} d(x, yz)$ corresponds to the set G/H of left cosets of H . (And a similar metric gives the set of right cosets of H .)*

If H is a normal subgroup, then the natural group structure on G/H is definable on G/ρ .

Proof. The fact that ρ is a pseudo-metric follows from the bi-invariance of the metric (really we only need right invariance). The easiest way to see this is that $\rho(x, y)$ is equal to $d_H(xH, yH)$, where d_H is the Hausdorff metric on sets.

We clearly have that if a and b are conjugate by an element of H , then $\rho(a, b) = 0$. Conversely, assume that $\rho(a, b) = 0$. This implies that for any $\varepsilon > 0$ we can find a c such that $d(a, bc) < \varepsilon$. By bi-invariance, this is equivalent to $d(e, a^{-1}bc) < \varepsilon$. So we have that for every $\varepsilon > 0$, $d_{\text{inf}}(a^{-1}b, H) < \varepsilon$, which implies that $a^{-1}b \in H$, since H is definable and therefore closed.

Now we just need to show that if H is a normal subgroup, then the group structure on G/H is definable on G/ρ . Let $q : G \rightarrow G/\rho$ be the quotient map. For any $a, b, c \in G$, we have that $\rho(q(a)q(b), q(c))$ is equal to $\rho(ab, c)$, since H is a normal subgroup. This is therefore a ρ -invariant formula and corresponds to a formula on the imaginary sort. \square

Theorem 5.3.22. *A completely metrizable topological group G can be given a metric under which it is a strongly minimal group if and only if it is non-compact and*

- there is a prime p such that every element of G has order p or
- G is divisible and can be written as $\mathbb{Q}^\kappa \oplus \mathbb{R}^n \oplus H$, where H has a compact subgroup K such that $H/K = \bigoplus_p (\mathbb{Z}/p^\infty \mathbb{Z})^{\alpha_p}$, with n and each α_p finite and κ arbitrary, where $\mathbb{Z}/p^\infty \mathbb{Z}$ is the p -Prüfer group.⁵

Furthermore, the resulting theory is essentially continuous if and only if G is of the second form with $n > 0$.

Proof. (\Rightarrow). Suppose that G is a strongly minimal group structure. By Proposition 5.3.17, G is Abelian.

For any prime p , let \overline{pG} denote the metric closure of $\{px : x \in G\}$, which is a subgroup. It is definable by $d(x, \overline{pG}) = \inf_y d(x, py)$. Suppose that for some p , \overline{pG} is a proper subgroup of G . By Lemma 5.3.16, this implies that \overline{pG} is compact. Consider the type-definable subgroup $\{x \in G : px = e\}$. If this is compact, then we have that G is homeomorphic to a product of two compact sets, which is a contradiction, therefore $\{x \in G : px = e\}$ must be all of G , by Lemma 5.3.16, and so G falls under the first bullet point.

Now assume that for every p , \overline{pG} is all of G . Let G' be an ω -saturated elementary extension of G . We now have that G' is a divisible group. By Fact 5.3.19, G' can be written as $\mathbb{R}^n \oplus H$, where H has a compact subgroup K such that H/K is a discrete Abelian group. n must be finite by local compactness. By replacing K with $G \cap K$, we may assume that K is a subgroup of G . Since G' is divisible, we must have that H is divisible as well. By the characterization of divisible Abelian groups, H/K can be written as $\mathbb{Q}^\kappa \oplus \bigoplus_p (\mathbb{Z}/p^\infty \mathbb{Z})^{\alpha_p}$. By Fact 5.3.20, we have that H can be written as

⁵The p -Prüfer group is isomorphic to the group of p^n th roots of unity under multiplication.

$\mathbb{Q}^k \oplus L$, where K is a subgroup of L and L/K is isomorphic to $\bigoplus_p (\mathbb{Z}/p^\infty \mathbb{Z})^{\alpha_p}$. For each prime p , the type-definable subgroup $\{x \in G' : px = e\}$ cannot be all of G' and so by Lemma 5.3.16 is compact. This implies that α_p must be finite. Since we can do this for every prime p , we have that G' falls under the case in the second bullet point.

G is an open subgroup of G' . This implies that $\mathbb{R}^n \oplus \{e\} \subseteq G'$ is also a subgroup of G . This implies that we can write G as $\mathbb{R}^n \oplus L$ with $K \subseteq L$ and L/K a discrete group. Each subgroup $\{x \in G : px = e\}$ is algebraic, so we must have that $\{x \in G : px = e\} = \{x \in G' : px = e\}$. Thinking of L/K as a subgroup of H/K , this implies that each factor of the form $\mathbb{Z}/p^\infty \mathbb{Z}$ is contained in L/K . Therefore L/K must be of the form $\mathbb{Q}^\lambda \oplus \bigoplus_p (\mathbb{Z}/p^\infty \mathbb{Z})^{\alpha_p}$ for some $\lambda \leq \kappa$, and G falls under the case in the second bullet point.

(\Leftarrow). Let G be a completely metrizable group such that every element of G has order p . Let K be a compact open subgroup of G such that G/K is discrete. By standard group theory, G and K are, algebraically speaking, vector spaces over \mathbb{F}_p . Let d^K be an arbitrary bi-invariant metric on K inducing the topology with diameter at most $\frac{1}{2}$. Let d be a metric on all of G defined by $d(a, b) = d^K(ab^{-1}, e)$ if $ab^{-1} \in K$ and $d(a, b) = 1$ otherwise. Because d^K is bi-invariant, this is a metric. Under this metric, K is contained in the $\frac{2}{3}$ -ball of e and so is in the algebraic closure of \emptyset and is in every model of $\text{Th}(G)$. If G' is any model of $\text{Th}(G)$ and G'' is any elementary extension of G' , it is not hard to show that any two elements of $G'' \setminus G'$ are automorphic, so $\text{Th}(G)$ is strongly minimal.

Let G be a completely metrizable group falling under the case in the second bullet point. Let K be the compact subgroup, and let d^K be an arbitrary bi-invariant metric giving the topology on K with diameter at most $\frac{1}{2}$. Given $a, b \in G$, write them as (a_0, a_1) and (b_0, b_1) , where $a_0, b_0 \in \mathbb{R}^n$ and $a_1, b_1 \in \mathbb{Q}^\kappa \oplus H$. Put a metric on G by $d(a, b) = 1$

if $a_1 b_1^{-1} \notin K$ and $d(a, b) = \max \left\{ \frac{\|a_0 - b_0\|_\infty}{1 + \|a_0 - b_0\|_\infty}, d^K(a_1 b_1^{-1}, e) \right\}$ otherwise. This is a metric which induces the topology on G . Note that $\frac{\|a_0 - b_0\|_\infty}{1 + \|a_0 - b_0\|_\infty}$ is a metric on \mathbb{R}^n , so d can be written as the maximum of a metric on \mathbb{R}^n and a metric on $\mathbb{Q}^\kappa \oplus H$ and therefore is a metric inducing the product topology. Note that since G is non-compact, the diameter of G with regards to d is 1.

Let G' be an elementary extension of (G, d) . The theory of (G, d) entails that the metric on G' is $[0, 1]$ valued, but also entails that for any $\varepsilon > 0$ with $\varepsilon < 1$, the closed ε -ball of any element is compact (and isometric to the corresponding closed ball around e).

Claim: For any prime p and any positive integer n , the type-definable set $\{x : p^n x = e\}$ is definable and algebraic.

Proof of claim. Each of these sets is compact in G , so we only need to show that they are definable. Fix $\varepsilon > 0$, and find $\delta > 0$ small enough that for $a \in K$ if $d^K(p^n a, e) < \delta$, then there exists a $b \in K$ with $d(a, b) < \varepsilon$ such that $p^n b = e$ (this always exists by compactness). We may take δ to be less than $\frac{1}{2}$ and less than ε .

For $a \in G$, suppose that $d(p^n a, e) < \delta$. If we write a as (a_0, a_1) with $a_0 \in \mathbb{R}^n$ and $a_1 \in \mathbb{Q}^\kappa \oplus H$, then this implies that $\frac{\|p^n a_0\|_\infty}{1 + \|p^n a_0\|_\infty} < p^n \varepsilon$ and also that $p^n a_1 \in K$ with $d(p^n a_1, e) < \delta$. By our choice of δ , this implies that there is some $b_1 \in K$ such that $d(a_1, b_1) < \varepsilon$ and $p^n b_1 = 0$. This implies that $d((a_0, a_1), (0, b_1)) < \max\{\varepsilon, \delta\} < \varepsilon$. Therefore, we have that the relevant set is definable. \square_{claim}

The claim implies that any $a \in G' \setminus G$ is divisible and torsion free. By the classification of divisible Abelian groups and the fact that G'/K is locally homeomorphic to G/K , we have that G' is topologically isomorphic to $G \oplus \mathbb{Q}^\lambda$ for some cardinal λ . Furthermore, we have that for any $a, b \in G'$, if $d(a, b) < 1$, then $b - a$ is in G . This implies

that G' can be realized as $G \oplus \mathbb{Q}^\lambda$ with the discrete metric on \mathbb{Q}^λ and the max metric on the product. From this we get that any two elements of $G' \setminus G$ are automorphic, and the same argument will work for any elementary extension G'' of G' , so we have that T is strongly minimal.

The ‘furthermore’ statement follows directly from Theorem 5.3.3. □

It is possible that G in the statement of Theorem 5.3.22 is not of the form $K \times G/K$. An example of this is the additive group of the p -adic numbers with the appropriate metric, as discussed after Example 5.3.1. K will be the subgroup of p -adic integers or some scaling of them, and the p -adic numbers do not decompose as a direct sum with any of these groups as a factor.

It is also possible to give a bi-invariant metric to a group of one of the forms given in Theorem 5.3.22 which will make it fail to be strongly minimal. This is very easy when the group is discrete— \mathbb{Q} with the metric $d(x, y)$ which is 1 when $x - y \in \mathbb{Z} \setminus \{0\}$ and 2 when $x - y \notin \mathbb{Z}$ —but it is also possible when the group is of a form that would result in an essentially continuous theory— \mathbb{R} with the metric $d(x, y) = \min\{|x - y|, 1\} + d(x, y + \mathbb{Z})$, which is a metric as it is the sum of a metric and a pseudo-metric.

Characterizing the metrics which make the groups identified in Theorem 5.3.22 strongly minimal seems difficult. We were unable to answer what ought to be the easiest question related to this issue.

Question 5.3.23. *If $(G, d, +)$ is a group structure such that $(G, d, +)$ is topologically one of the groups specified in Theorem 5.3.22, d is a bi-invariant metric, and (G, d) is a strongly minimal metric space, does it follow that $(G, d, +)$ is strongly minimal?*

Another direction for future study would be to replicate the very tight characterization of transitive faithful ω -stable group actions on strongly minimal sets:

Fact 5.3.24 ([Bue17, Thm. 3.5.2]). *If (G, X) is a (discrete) ω -stable transitive faithful group action with X strongly minimal, then $\text{MR}(G) \leq 3$ and*

1. *if $\text{MR}(G) = 1$, then G has a definable finite index subgroup H which acts regularly on X and*
2. *if $\text{MR}(G) \geq 2$, then there is a field K definable on X or $X \setminus \{a\}$ for some point a .*

Recall that a group action is *transitive* if the orbit of every element is the entire set, it is *faithful* if $\text{Stab}(X)$ is trivial, and it is *regular* if for any $a \in X$, the function $g \mapsto ga$ is a bijection. Note that in the first case, H is strongly minimal.

A completely reckless conjecture would be that only the first case of Fact 5.3.24 can occur in an essentially continuous theory (this would be related to a negative answer to Question 5.3.8).

Question 5.3.25. *Is it possible to have a locally compact metric group G with bi-invariant metric acting faithfully and transitively on a metric space X such that $\text{Th}(G, X)$ is essentially continuous, ω -stable, and has X strongly minimal, without G having a definable strongly minimal group with compact index that acts regularly on X ?*

Local compactness also adds a new parameter to these questions regarding these kinds of characterizations.

Question 5.3.26. *If G is a (not necessarily locally compact) metric group with bi-invariant metric acting faithfully and transitively on a metric space X such that $\text{Th}(G, X)$*

is ω -stable and has X strongly minimal, is G necessarily locally compact?

5.4 Banach and Hilbert Structures

In this section we will consider expansions of Banach spaces. We introduce the notion of an indiscernible subspace. An indiscernible subspace is a subspace in which types of tuples of elements only depend on their quantifier-free types in the reduct consisting of only the metric and the constant $\mathbf{0}$. Similarly to indiscernible sequences, indiscernible subspaces are always consistent with a Banach theory (with no stability assumption, see Theorem 5.4.9) but are not always present in every model. We will show that an indiscernible subspace always takes the form of an isometrically embedded real Hilbert space wherein the type of any tuple only depends on its quantifier-free type in the Hilbert space. The notion of an indiscernible subspace is of independent interest in the model theory of Banach and Hilbert structures, and in particular here we use it to improve the results of Shelah and Usvyatsov in the context of types in the full language (as opposed to Δ -types). Specifically, in this context we give a shorter proof of Shelah and Usvyatsov’s main result [SU19, Prop. 4.13], we improve their result on the strong uniqueness of Morley sequences in minimal wide types [SU19, Prop. 4.12], and we expand on their commentary on the “induced structure” of the span of a Morley sequence in a minimal wide type [SU19, Rem. 5.6]. This more restricted case is what is relevant to inseparably categorical Banach theories, so our work is applicable to the problem of characterizing such theories.

5.4.1 Banach Theory Background

For $K \in \{\mathbb{R}, \mathbb{C}\}$, we think of a K -Banach space X as being a metric structure \mathfrak{X} whose underlying set is the closed unit ball $B(X)$ of X with metric $d(x, y) = \|x - y\|$.⁶ This structure is taken to have for each tuple $\bar{a} \in K$ an $|\bar{a}|$ -ary predicate $s_{\bar{a}}(\bar{x}) = \left\| \sum_{i < |\bar{a}|} a_i x_i \right\|$, although we will always write this in the more standard form. Note that we evaluate this in X even if $\sum_{i < |\bar{a}|} a_i x_i$ is not actually an element of the structure \mathfrak{X} . For convenience, we will also have a constant for the zero vector, $\mathbf{0}$, and an n -ary function $\sigma_{\bar{a}}(\bar{x})$ such that $\sigma_{\bar{a}}(\bar{x}) = \sum_{i < |\bar{a}|} a_i x_i$ if it is in $B(X)$ and $\sigma_{\bar{a}}(\bar{x}) = \frac{\sum_{i < |\bar{a}|} a_i x_i}{\left\| \sum_{i < |\bar{a}|} a_i x_i \right\|}$ otherwise. If $|a| \leq 1$, we will write ax for $\sigma_a(x)$. Note that while this is an uncountable language, it is interdefinable with a countable reduct of it (restricting attention to rational elements of K). These structures capture the typical meaning of the ultraproduct of Banach spaces. We will often conflate X and the metric structure \mathfrak{X} in which we have encoded X .

Definition 5.4.1. A *Banach (or Hilbert) structure* is a metric structure which is the expansion of a Banach (or Hilbert) space. A *Banach (or Hilbert) theory* is the theory of such a structure. The adjectives *real* and *complex* refer to the scalar field K . \triangleleft

C^* - and other Banach algebras are commonly studied examples of Banach structures that are not just Banach spaces. The theory of an infinite dimensional Hilbert space is written IHS.

A central problem in continuous logic is the characterization of inseparably categorical countable theories, that is to say countable theories with a unique model in each uncountable density character. The analog of Morley's theorem was shown in continuous logic via related formalisms [BY05, SU11], but no satisfactory analog of the

⁶For another equivalent approach, see [BYBHU08], which encodes Banach structures as many-sorted metric structures with balls of various radii as different sorts.

Baldwin-Lachlan theorem or its precise structural characterization of uncountably categorical discrete theories in terms of strongly minimal sets is known. Some progress in the specific case of Banach theories has been made in [SU19], in which Shelah and Usvyatsov introduce the notion of a wide type and the notion of a minimal wide type, which they argue is the correct analog of strongly minimal types in the context of inseparably categorical Banach theories.

Definition 5.4.2. A type p in a Banach theory is *wide* if its set of realizations consistently contain the unit sphere of an infinite dimensional real subspace.

A type is *minimal wide* if it is wide and has a unique wide extension to every set of parameters. ◁

In [SU19], Shelah and Usvyatsov were able to show that every Banach theory has wide complete types using the following classical concentration of measure results of Dvoretzky and Milman, which Shelah and Usvyatsov refer to as the Dvoretzky-Milman theorem.

Fact 5.4.3 (Dvoretzky-Milman theorem). *Let $(X, \|\cdot\|)$ be an infinite dimensional real Banach space with unit sphere S , and let $f : S \rightarrow \mathbb{R}$ be a uniformly continuous function. For any $k < \omega$ and $\varepsilon > 0$, there exists a k -dimensional subspace $Y \subset X$ and a Euclidean norm⁷ $\|\!\|\!\cdot\!\|$ on Y such that for any $a, b \in S \cap Y$, we have $\|\!\|a\!\| \leq \|a\| \leq (1 + \varepsilon)\|\!\|a\!\|$ and $|f(a) - f(b)| < \varepsilon$.⁸*

Shelah and Usvyatsov showed that in a stable Banach theory every wide type has a minimal wide extension (possibly over a larger set of parameters) and that every

⁷A norm $\|\!\|\!\cdot\!\|$ is *Euclidean* if it satisfies the parallelogram law, $2\|\!\|a\!\|^2 + 2\|\!\|b\!\|^2 = \|\!\|a + b\!\|^2 + \|\!\|a - b\!\|^2$, or, equivalently, if it is induced by an inner product.

⁸Fact 5.4.3 without f is (a form of) Dvoretzky's theorem.

Morley sequence in a minimal wide type is an orthonormal basis of a subspace isometric to a real Hilbert space. Furthermore, they showed that in an inseparably categorical Banach theory, every inseparable model is prime over a countable set of parameters and a Morley sequence in some minimal wide type, analogously to how a model of a discrete uncountably categorical theory is always prime over some finite set of parameters and a Morley sequence in some strongly minimal type.

The key ingredient to our present work is the following result, due to Milman. It extends the Dvoretzky-Milman theorem in a manner analogous to the extension of the pigeonhole principle by Ramsey's theorem.⁹

Definition 5.4.4. Let $(X, \|\cdot\|)$ be a Banach space. If a_0, a_1, \dots, a_{n-1} and b_0, b_1, \dots, b_{n-1} are ordered n -tuples of elements of X , we say that \bar{a} and \bar{b} are *congruent* if $\|a_i - a_j\| = \|b_i - b_j\|$ for all $i, j \leq n$, where we take $a_n = b_n = \mathbf{0}$. We will write this as $\bar{a} \cong \bar{b}$. \triangleleft

Fact 5.4.5 ([Mil71], Thm. 3). *Let S^∞ be the unit sphere of a separable infinite dimensional real Hilbert space H , and let $f : (S^\infty)^n \rightarrow \mathbb{R}$ be a uniformly continuous function. For any $\varepsilon > 0$ and any $k < \omega$ there exists a k -dimensional subspace V of H such that for any $a_0, a_1, \dots, a_{n-1}, b_0, b_1, \dots, b_{n-1} \in S^\infty$ with $\bar{a} \cong \bar{b}$, $|f(\bar{a}) - f(\bar{b})| < \varepsilon$.*

Note that the analogous result for inseparable Hilbert spaces follows immediately, by restricting attention to a separable infinite dimensional subspace. Also note that by using Dvoretzky's theorem and an easy compactness argument, Fact 5.4.5 can be generalized to arbitrary infinite dimensional Banach spaces.

⁹The original Dvoretzky-Milman result is often compared to Ramsey's theorem, such as when Gromov coined the term *the Ramsey-Dvoretzky-Milman phenomenon* [Gro83], but in the context of Fact 5.4.5 it is hard not to think of the $n = 1$ case as being analogous to the pigeonhole principle and the $n > 1$ cases as being analogous to Ramsey's theorem.

Connection to Extreme Amenability

A modern proof of Fact 5.4.5 would go through the extreme amenability of the unitary group of an infinite dimensional Hilbert space endowed with the strong operator topology, or in other words the fact that any continuous action of this group on a compact Hausdorff space has a fixed point, which was originally shown in [GM83]. This connection is unsurprising. It is well known that the extreme amenability of $\text{Aut}(\mathbb{Q})$ (endowed with the topology of pointwise convergence) can be understood as a restatement of Ramsey's theorem. It is possible to use this to give a high brow proof of the existence of indiscernible sequences in any first-order theory T :

Proof. Fix a first-order theory T . Let Q be a family of variables indexed by the rational numbers. The natural action of $\text{Aut}(\mathbb{Q})$ on $S_Q(T)$, the Stone space of types over T in the variables Q , is continuous and so by extreme amenability has a fixed point. A fixed point of this action is precisely the same thing as the type of a \mathbb{Q} -indexed indiscernible sequence over T , and so we get that there are models of T with indiscernible sequences. \square

A similar proof of the existence of indiscernible subspaces in Banach theories (Theorem 5.4.9) is possible, but requires an argument that the analog of $S_Q(T)$ is non-empty (which follows from Dvoretzky's theorem) and also requires more delicate bookkeeping to define the analog of $S_Q(T)$ and to show that the action of the unitary group of a separable Hilbert space is continuous. In the end this is more technical than a proof using Fact 5.4.5 directly.

5.4.2 Asymptotically Hilbertian Spaces Do Not Interpret a Strongly Minimal Set

For an overview of the properties of asymptotically Hilbertian spaces in the context of continuous logic, see [HR16].

Compare the following with the fact that every discrete theory interprets a strongly minimal set.

Proposition 5.4.6. *If T is the theory of an asymptotically Hilbertian space then it does not have any non-algebraic locally compact imaginaries. In particular it does not interpret a strongly minimal set.*

Proof. Assume that T has a non-algebraic locally compact imaginary. By the same argument as in Proposition 5.2.1 this implies that it has a strongly minimal imaginary. We may assume that this is definable over the unique approximately \aleph_0 -saturated model \mathfrak{M} . Since T is \aleph_1 -categorical this implies that every model is prime over this strongly minimal imaginary. Let $\mathfrak{A} \succ \mathfrak{B} \succ \mathfrak{M}$ be a pair of proper elementary extensions such that in each extension the dimension of the strongly minimal set increases by 1. Since $\mathfrak{B} \succ \mathfrak{M}$ is a minimal extension it must be the case that the vector space dimension of the home sort increases by precisely 1. Likewise since $\mathfrak{A} \succ \mathfrak{B}$ is a minimal extension it must be the case that vector space dimension of the home sort increases by precisely 1. Let $b \in \mathfrak{B} \setminus \mathfrak{M}$ realize the unique type of an element of norm 1 orthogonal to \mathfrak{M} . Likewise let $a \in \mathfrak{A} \setminus \mathfrak{B}$ realize the unique type of an element of norm 1 orthogonal to \mathfrak{B} . It must be the case that $\mathfrak{A} = \mathfrak{M} \oplus V$ where V is a 2-dimensional Hilbert space generated by the orthogonal basis $\{a, b\}$. Find $\varepsilon > 0$ small enough that $\inf\{d(x, y) : x \in I(\mathfrak{M}), y \in I(\mathfrak{B}) \setminus I(\mathfrak{M})\} > \varepsilon$ and $\inf\{d(x, y) : x \in I(\mathfrak{B}), y \in I(\mathfrak{A}) \setminus I(\mathfrak{B})\} > \varepsilon$, where I is the strongly minimal imaginary.

Such an ε must exist since I is strongly minimal. Now find $\delta > 0$ small enough that if σ is a home sort automorphism fixing \mathfrak{M} and satisfying $d(x, \sigma(x)) < \delta$ for every x , then the induced automorphism of I , σ_I , satisfies $d(x, \sigma_I(x)) < \varepsilon$ for every x (such a δ must exist). Now let σ be the automorphism of \mathfrak{A} that fixes \mathfrak{M} and rotates V by an angle $\frac{\pi}{2n}$ small enough that $d(x, \sigma(x)) < \delta$ for every $x \in \mathfrak{A}$. Assume without loss of generality that $\sigma^n(a) = b$ and $\sigma^n(b) = -a$. Now if we look at σ_I^n , by construction it must be the case that $\sigma_I^n(I(\mathfrak{B})) \subseteq I(\mathfrak{B})$ and $\sigma_I^n(I(\mathfrak{B})) \not\subseteq I(\mathfrak{M})$. Therefore if we look at the structure $\mathfrak{N} = \sigma_I^n(\mathfrak{B})$, it must be the case that $a \in \mathfrak{N}$ and $I(\mathfrak{B}) = I(\mathfrak{N})$, but since T has no imaginary Vaughtian pairs this implies that $\mathfrak{N} = \mathfrak{A}$, which is a contradiction. \square

5.4.3 Indiscernible Subspaces

Definition 5.4.7. Let T be a Banach theory. Let $\mathfrak{M} \models T$, and let $A \subseteq \mathfrak{M}$ be some set of parameters. An *indiscernible subspace over A* is a real subspace V of \mathfrak{M} such that for any $n < \omega$ and any n -tuples $\bar{b}, \bar{c} \in V$, $\bar{b} \equiv_A \bar{c}$ if and only if $\bar{b} \cong \bar{c}$.

If p is a type over A , then V is an *indiscernible subspace in p (over A)* if it is an indiscernible subspace over A and $b \models p$ for all $b \in V$ with $\|b\| = 1$. \triangleleft

Note that an indiscernible subspace is a real subspace even if T is a complex Banach theory. Also note that an indiscernible subspace in p is not literally contained in the realizations of p , but rather has its unit sphere contained in the realizations of p . It might be more accurate to talk about “indiscernible spheres,” but we find the subspace terminology more familiar.

Indiscernible subspaces are very metrically regular.

Proposition 5.4.8. *Suppose V is an indiscernible subspace in some Banach structure.*

Then V is isometric to a real Hilbert space.

In particular, a real subspace V of a Banach structure is indiscernible over A if and only if it is isometric to a real Hilbert space and for every $n < \omega$ and every pair of n -tuples $\bar{b}, \bar{c} \in V$, $\bar{b} \equiv_A \bar{c}$ if and only if for all $i, j < n$, $\langle b_i, b_j \rangle = \langle c_i, c_j \rangle$.

Proof. For any real Banach space W , if $\dim W \leq 1$, then W is necessarily isometric to a real Hilbert space. If $\dim V \geq 2$, let V_0 be a 2-dimensional subspace of V . A subspace of an indiscernible subspace is automatically an indiscernible subspace, so V_0 is indiscernible. For any two distinct unit vectors a and b , indiscernibility implies that for any $r, s \in \mathbb{R}$, $\|ra + sb\| = \|sa + rb\|$, hence the unique linear map that switches a and b fixes $\|\cdot\|$. This implies that the automorphism group of $(V_0, \|\cdot\|)$ is transitive on the $\|\cdot\|$ -unit circle. By John's theorem on maximal ellipsoids [Joh48], the unit ball of $\|\cdot\|$ must be an ellipse, so $\|\cdot\|$ is a Euclidean norm.

Thus every 2-dimensional real subspace of V is Euclidean and so $(V, \|\cdot\|)$ satisfies the parallelogram law and is therefore a real Hilbert space.

The 'in particular' statement follows from the fact that in a real Hilbert subspace of a Banach space, the polarization identity [BB02, Prop. 14.1.2] defines the inner product in terms of a particular quantifier-free formula:

$$\langle x, y \rangle = \frac{1}{4} (\|x + y\|^2 - \|x - y\|^2).^{10} \quad \square$$

As mentioned in [SU19, Cor. 3.9], it follows from Dvoretzky's theorem that if p is a

¹⁰There is also a polarization identity for the complex inner product:

$$\langle x, y \rangle_{\mathbb{C}} = \frac{1}{4} (\|x + y\|^2 - \|x - y\|^2 + i\|x - iy\|^2 - i\|x + iy\|^2).$$

wide type and \mathfrak{M} is a sufficiently saturated model, then $p(\mathfrak{M})$ contains the unit sphere of an infinite dimensional subspace isometric to a Hilbert space. We refine this by showing that, in fact, an indiscernible subspace can be found.

Theorem 5.4.9. *Let A be a set of parameters in a Banach theory T , and let p be a wide type over A . For any κ , there is $\mathfrak{M} \models T$ and a subspace $V \subseteq \mathfrak{M}$ of dimension κ such that V is an indiscernible subspace in p over A . In particular, any $\aleph_0 + \kappa + |A|$ -saturated \mathfrak{M} will have such a subspace.*

Proof. For any set Δ of A -formulas, call a subspace V of a model \mathfrak{N} of T_A Δ -indiscernible in p if every unit vector in V models p and for any $n < \omega$ and any formula $\varphi \in \Delta$ of arity n and any n -tuples $\bar{b}, \bar{c} \in V$ with $\bar{b} \cong \bar{c}$, we have $\mathfrak{N} \models \varphi(\bar{b}) = \varphi(\bar{c})$.

Since p is wide, there is a model $\mathfrak{N} \models T$ containing an infinite dimensional subspace W isometric to a real Hilbert space such that for all $b \in W$ with $\|b\| = 1$, $b \models p$. This is an infinite dimensional \emptyset -indiscernible subspace in p .

Now for any finite set of A -formulas Δ and formula φ , assume that we've shown that there is a model $\mathfrak{N} \models T$ containing an infinite dimensional Δ -indiscernible subspace V in p over A . We want to show that there is a $\Delta \cup \{\phi\}$ -indiscernible subspace in V . By Fact 5.4.5, for every $k < \omega$ there is a k -dimensional subspace $W_k \subseteq V$ such that for any unit vectors $b_0, \dots, b_{\ell-1}, c_0, \dots, c_{\ell-1}$ in W_k with $\bar{b} \cong \bar{c}$, we have that $|\varphi^{\mathfrak{N}}(\bar{b}) - \varphi^{\mathfrak{N}}(\bar{c})| < 2^{-k}$. If we let $\mathfrak{N}_k = (\mathfrak{N}_k, W_k)$ where we've expanded the language by a fresh predicate symbol D such that $D^{\mathfrak{N}_k}(x) = d(x, W_k)$, then an ultraproduct of the sequence \mathfrak{N}_k will be a structure $(\mathfrak{N}_\omega, W_\omega)$ in which W_ω is an infinite dimensional Hilbert space.

Claim: W_ω is $\Delta \cup \{\varphi\}$ -indiscernible in p .

Proof of claim. Fix an m -ary formula $\psi \in \Delta \cup \{\varphi\}$, and let $f(k) = 0$ if $\psi \in \Delta$ and

$f(k) = 2^{-k}$ if $\psi = \varphi$. For any $k \geq 2m$ and $b_0, \dots, b_{m-1}, c_0, \dots, c_{m-1}$ in the unit ball of W_k , there is a $2m$ dimensional subspace $W' \subseteq W_k$ containing \bar{b}, \bar{c} . By compactness of $B(W')^m$ (where $B(X)$ is the unit ball of X), we have that for any $\varepsilon > 0$ there is a $\delta(\varepsilon) > 0$ such that if $|\langle b_i, b_j \rangle - \langle c_i, c_j \rangle| < \delta(\varepsilon)$ for all $i, j < m$, then $|\psi^{\mathfrak{M}}(\bar{b}) - \psi^{\mathfrak{M}}(\bar{c})| \leq f(k) + \varepsilon$. Note that we can take the function δ to only depend on ψ , specifically its arity and modulus of continuity, and not on k , since $B(W')^m$ is always isometric to $B(\mathbb{R}^{2m})^m$. Therefore, in the ultraproduct we will have $(\forall i, j < m) |\langle b_i, b_j \rangle - \langle c_i, c_j \rangle| < \delta(\varepsilon) \Rightarrow |\psi^{\mathfrak{M}}(\bar{b}) - \psi^{\mathfrak{M}}(\bar{c})| \leq \varepsilon$ and thus $\bar{b} \cong \bar{c} \Rightarrow \psi^{\mathfrak{M}}(\bar{b}) = \psi^{\mathfrak{M}}(\bar{c})$, as required. \square_{Claim}

Now for each finite set of A -formulas we've shown that there's a structure $(\mathfrak{M}_\Delta, V_\Delta)$ (where, again, V_Δ is the set defined by the new predicate symbol D) such that $\mathfrak{M}_\Delta \models T_A$ and V_Δ is an infinite dimensional Δ -indiscernible subspace in p . By taking an ultraproduct with an appropriate ultrafilter we get a structure (\mathfrak{M}, V) where $\mathfrak{M} \models T_A$ and V is an infinite dimensional subspace. V is an indiscernible subspace in p over A by the same argument as in the claim.

Finally note that by compactness we can take V to have arbitrarily large dimension and that any subspace of an indiscernible subspace in p over A is an indiscernible subspace in p over A , so we get the required result. \square

Together with the fact that wide types always exist in Banach theories with infinite dimensional models [SU19, Thm. 3.7], we get a corollary.

Corollary 5.4.10. *Every Banach theory with infinite dimensional models has an infinite dimensional indiscernible subspace in some model. In particular, every such theory has an infinite indiscernible set, namely any orthonormal basis of an infinite dimensional indiscernible subspace.*

5.4.4 Minimal Wide Types

Characterization of Morley Sequences in Terms of Indiscernible Subspaces

Compare the following Theorem 5.4.11 with this fact in discrete logic: If p is a minimal type (i.e. p has a unique global non-algebraic extension), then an infinite sequence of realizations of p is a Morley sequence in p if and only if it is an indiscernible sequence.

Here we are using the definition of Morley sequence for (possibly unstable) A -invariant types: Let p be a global A -invariant type, and let $B \supseteq A$ be some set of parameters. A sequence $\{c_i\}_{i < \kappa}$ is a *Morley sequence in p over B* if for all $i < \kappa$, $\text{tp}(c_i/Bc_{<i}) = p \upharpoonright Bc_{<i}$. Note that this definition of Morley sequence agrees with the standard definition for types that are stable in the sense of Lascar and Poizat (as described in [SU19, Def. 4.1]).

Theorem 5.4.11. *Let p be a minimal wide type over the set A . For $\kappa \geq \aleph_0$, a set of realizations $\{b_i\}_{i < \kappa}$ of p is a Morley sequence in (the unique global minimal wide extension of) p if and only if it is an orthonormal basis of an indiscernible subspace in p over A .*

Proof. All we need to show is that an orthonormal basis of an indiscernible subspace in p over A is a Morley sequence in p . The converse will follow from the fact that all Morley sequences in a fixed invariant type of the same length have the same type along with the fact that minimal wide types have a unique global wide extension, which is therefore invariant.

Let V be an indiscernible subspace in p over A . Let $\{e_i\}_{i < \kappa}$ be an orthonormal basis of V . By construction, $\text{tp}(e_0/A) = p$. Let q be the global minimal wide extension of p . Assume that for some $j < \kappa$ we've shown for all $i < j$ that $\text{tp}(e_i/Ae_{<i}) = q \upharpoonright Ae_{<i}$. Let

$W = \overline{\text{span}}(e_{\geq j})$. Since V is an indiscernible subspace over A , for all unit norm $b, c \in W$, $b \equiv_{Ae_{< j}} c$, so in particular $\text{tp}(b/Ae_{< j})$ is wide. Since p is minimal wide we must have $\text{tp}(b/Ae_{< j}) = q \upharpoonright Ae_{< j}$. Therefore $\{e_i\}_{i < \kappa}$ is a Morley sequence. \square

What is unclear at the moment is the answer to this question:

Question 5.4.12. *If p is a minimal wide type over the set A , is it stable in the sense of [SU19, Def. 4.1]? In other words, is every type q extending p over a model $\mathfrak{M} \supseteq A$ a definable type?*

Strongly Minimal Wide Types

At the moment the contents of this section are little more than an observation, but hopefully in the future it may be a fruitful one.

In [SU19], Shelah and Usvyatsov construct minimal wide types in an arbitrary stable theory.¹¹ This is analogous to the construction of minimal types in discrete stable theories (i.e. fork until you do not have a non-algebraic forking extension), but just as with that construction, the method in [SU19] does not give precise control over the resulting type.

There is a natural analog of strongly minimal types in the context of wide types. The relevant notion to generalize is Definition 4.1.8. This gives the following.

Definition 5.4.13. In a Banach theory T , a global type $p \in S_1(\mathfrak{C})$ is *strongly minimal wide* if it is d -atomic in the set of wide global types.

An arbitrary type is *strongly minimal wide* if it has a unique wide global extension and that extension is strongly minimal wide. \triangleleft

¹¹More specifically, they showed that any stable wide type has a minimal wide extension.

Just as with minimal and strongly minimal types in discrete logic, in a stable, non- ω -stable Banach theory (with infinite dimensional models) there may fail to be any strongly minimal wide types, although there are always minimal wide types. Even in an ω -stable theory, there may be minimal wide types that are not strongly minimal wide. Nevertheless, we do have following.

Proposition 5.4.14. *Let T be an ω -stable Banach theory. Any open subset U of $S_1(\mathfrak{C})$ (the space of global types) containing a wide type contains a strongly minimal wide type.*

Proof. It is not hard to see that the stated proposition is equivalent to saying that strongly minimal wide types are dense in the set of wide types. This follows immediately from Proposition 3.7 in [BY08c]. \square

Given a strongly minimal wide type p , by Proposition 2.4.4 we can find a definable set $D \subseteq S_1(A)$ such that $D \cap W = \{p\}$, where W is the (closed) set of wide types in $S_1(A)$. Since the set of norm 1 types is always definable (by the formula $1 - \|x\|$), we can require that $D(x) \models \|x\| = 1$ as well.

The issue with continuing the analogy with strongly minimal types is that while there is an easy characterization of definable sets containing a unique non-algebraic type which happens to be strongly minimal, there is not a clear analogous characterization of definable sets containing a unique wide type which happens to be strongly minimal wide.

Without the requirement that $D(x) \models \|x\| = 1$, we have a counterexample.

Proposition 5.4.15. *Let T be a Banach theory. For any zero set $F(x) \models \|x\| = 1$, there is a definable set $D(x)$ such that $D(x) \cap \llbracket \|x\| = 1 \rrbracket = F(x)$.*

Proof. Let $\varphi(x)$ be a $[0, 1]$ -valued formula such that $\llbracket \varphi(x) = 1 \rrbracket = F(x)$.

We need to show that the set

$$D := \{\mathbf{0}\} \cup \left\{ p \in S_1(A) : p(x) \models 0 < \|x\| \leq \varphi\left(\frac{x}{\|x\|}\right) \right\}$$

is definable.¹²

First we need to show that D is closed. Let $\{q_i\}_{i \in I}$ be a net limiting to some type in D . There are two cases. Either $\lim_{i \in I} \|q_i\|$ is 0 or it is strictly positive. The first case is covered by the fact that $\mathbf{0} \in D$, so assume that the second case holds. There must be some $\varepsilon > 0$ and some $i \in I$ such that for all $j \geq i$, $\|q_j\| > \varepsilon$. So now we have that $\lim_{i \in I} q_i \in D$ if and only if $\lim_{i \in I} \|q_i\| \leq \lim_{i \in I} \varphi\left(\frac{x}{\|x\|}\right)$ by the continuity of $\varphi\left(\frac{x}{\|x\|}\right)$ away from $\mathbf{0}$.

To establish that D is definable we need to show that if $\{q_i\}_{i \in I}$ limits to a type in D , then $\lim_{i \in I} d(q_i, D) = 0$. If $\{q_i\}_{i \in I}$ limits to $\mathbf{0}$, then $d(q_i, D) \leq d(q_i, \mathbf{0}) = \|q_i\| \rightarrow 0$ as well. If $\{q_i\}$ is limiting to a point other than $\mathbf{0}$, then $\{\|q_i\|\}_{i \in I}$ is eventually uniformly positive. So there is an $i \in I$ such that for all $j \geq i$,

$$d(q_j, D) \leq d\left(q_j, q_j \min\left\{1, \|q_j\|^{-1} \varphi\left(\frac{q_j}{\|q_j\|}\right)\right\}\right) = \|q_j\| \div \varphi\left(\frac{q_j}{\|q_j\|}\right).$$

We have already established that this quantity must go to 0 if $\{q_i\}_{i \in I}$ is limiting to $D \setminus \{\mathbf{0}\}$, so we have that D is definable.

Now by construction we have that $D(x) \cap \llbracket \varphi(x) = 0 \rrbracket = F(x)$. □

This implies that given a minimal wide type p we can always find a definable set D

¹²Note that while $\varphi\left(\frac{x}{\|x\|}\right)$ is not technically a formula, D is nevertheless well defined.

such that $D \cap \llbracket \|x\| = 1 \rrbracket = \{p\}$.

A natural attempt at a characterization would be definable sets which are wide but for which

for every formula $\varphi(x)$ and every $\varepsilon > 0$, there is an $n < \omega$ and a $\delta > 0$ such that if a finite dimensional subspace V with dimension at least n has $d_{\text{inf}}(a, D) < \delta$ for all $a \in V$ of norm 1, then for any $a, b \in V$ with norm 1, we have that $|\varphi(a) - \varphi(b)| \leq \varepsilon$.

But it is only clear that this establishes that D contains a unique wide type (which is therefore minimal wide).

The fundamental problem is that in general if D is a definable set and F is a closed set, then $D \cap F$ may fail to be relatively definable in F , in the sense of Definition 2.3.13. This is of course in opposition to the behavior of relative definability in discrete logic.

Even beyond this, what we really want is D to be a nice definable subspace containing a unique wide type that happens to be strongly minimal wide, but it is still entirely unclear that this is always possible.

This all leaves the following questions.

Question 5.4.16. *If $D(x)$ is a definable set such that $D(x) \models \|x\| = 1$ and $D(x)$ contains a unique global wide type p , is p strongly minimal wide?*

Question 5.4.17. *Is there a nice characterization of those definable sets which contain a unique wide type which happens to be strongly minimal wide?*

Question 5.4.18. *Is every strongly minimal wide type contained in a definable subspace in which it is the unique wide type?*

5.5 Randomizations

In this section we will present a proof that randomizations of discrete theories are dictionaric and a proof that randomizations of arbitrary continuous theories are dictionaric. We have included both proofs because the proof in the discrete case gives more detailed information about definable sets in the randomization of a discrete theory. We should note, though, that the second proof is considerably simpler.

See [BYJK09, AK15] for an introduction to randomizations in continuous logic and the relevant definitions. However, to avoid overloading the notation $\llbracket \cdot \rrbracket$, we will use some slightly different notation.

Notation 5.5.1. For a discrete formula φ , we will represent the element of the corresponding probability algebra in a randomization as $[\varphi]$. We will write comparisons between elements of the probability algebra with \sqsubseteq , and we will write \sqcap , \sqcup , and \neg for the Boolean algebra operations.

For any discrete formula φ , $\mathbb{E}\varphi$ is the continuous formula in the corresponding randomization that gives the expected value of φ (i.e. the probability that φ is true). \triangleleft

Lemma 5.5.2. *Given a and b in a model of a randomization of a discrete theory and c in the probability algebra sort, there is an e such that $[a = e] \sqsupseteq c$ and $[b = e] \sqsupseteq \neg c$.*

Proof. This follows from Facts 2.5 and 2.6 in [AK15]. \square

Proposition 5.5.3. *The randomization T^R of any discrete theory T is dictionaric.*

In particular, for any finite list $\varphi_0(\bar{x}, \bar{y}), \dots, \varphi_{n-1}(\bar{x}, \bar{y})$ of discrete formulas and any closed set $F \subseteq [0, 1]^n$, the closed formula

$$G(\bar{x}, \bar{y}) = ((\mathbb{E}\varphi_0(\bar{x}, \bar{y}), \dots, \mathbb{E}\varphi_{n-1}(\bar{x}, \bar{y})) \in F)$$

is \bar{y} -pointwise definable over T^R (i.e. for any fixed tuple of parameters \bar{a} , $G(\bar{x}, \bar{a})$ is a definable set).

Proof. Quantifier elimination tells us that every neighborhood U of a type $p(\bar{x}) \in S_n(A)$ has a sub-neighborhood of the form $\bigwedge_{i < k} |\mathbb{E}\varphi_i(\bar{x}, \bar{a}) - \mathbb{E}\varphi_i(p, \bar{a})| < \varepsilon$ for some sequence of restricted formulas $\{\varphi_i\}_{i < k}$, some finite list of parameters \bar{a} , and some $\varepsilon > 0$.

For each $\bar{\gamma} \in 2^k$, let $\psi_{\bar{\gamma}}(\bar{x}, \bar{y}) = \bigwedge_{i < k} \neg^{\gamma_i} \varphi_i(\bar{x}, \bar{y})$ for $\bar{\gamma} \in 2^k$ (where \neg^0 is ignored and \neg^1 is \neg). We are really going to show that the set $\bigwedge_{\bar{\gamma} \in 2^k} \mathbb{E}\psi_{\bar{\gamma}}(\bar{x}, \bar{a}) = \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a})$ is definable.

For each pair $(\bar{\gamma}, A)$, with $\bar{\gamma} \in A \subseteq 2^k$, let

$$\xi_{\bar{\gamma}, A}(\bar{x}, \bar{y}) = \psi_{\bar{\gamma}}(\bar{x}, \bar{y}) \wedge \bigwedge_{\bar{\eta} \in 2^k} \begin{cases} \exists \bar{z} \psi_{\bar{\eta}}(\bar{z}, \bar{y}) & \bar{\eta} \in A \\ \neg \exists \bar{z} \psi_{\bar{\eta}}(\bar{z}, \bar{y}) & \bar{\eta} \notin A \end{cases}.$$

Let Q be the set of all pairs $(\bar{\gamma}, A)$ with $\bar{\gamma} \in A \subseteq 2^k$, and note that for any give \bar{b} and \bar{c} (in a model in the original theory), $\xi_{\bar{\gamma}, A}(\bar{b}, \bar{c})$ is true for precisely one $(\bar{\gamma}, A) \in Q$.

We may assume that \bar{a} occurs in a model consisting of measurable maps from a given measure space into some structure (not necessarily finitely or countably valued maps).

Let the relevant measure space be Ω .

For each non-empty $A \subseteq 2^k$, let

$$X_A = \{\omega \in \Omega : \{\bar{\gamma} \in 2^k : \exists \bar{x} \psi_{\bar{\gamma}}(\bar{x}, \bar{a}(\omega))\} = A\},$$

and let $r_A = \mu(X_A)$. Note that the X_A form a partition of Ω .

Let Z be the set of $A \subseteq 2^k$ such that $r_A > 0$. For each $A \in Z$, let $C_A \subseteq [0, 1]^{2^k}$ be

the set of all elements of $[0, 1]^{2^k}$ which sum to 1 and whose support is contained in A . This is clearly a compact set.

Let

$$f : \prod_{A \in \mathcal{Z}} C_A \rightarrow [0, 1]^{2^k}$$

be the function defined by

$$f(\tau) = \sum_{A \in \mathcal{Z}} r_A \tau(A),$$

where $\tau(A)$ is understood to be an element of $C_A \subseteq [0, 1]^{2^k}$.

Let P be the element of $[0, 1]^{2^k}$ given by $P(\bar{\gamma}) = \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a})$. Note that since p is a consistent type, P is in the range of f .

By compactness, for every $\varepsilon > 0$, there is a $\delta > 0$ such that if $d(f(R), P) < \delta$, then $d(R, f^{-1}(P)) < \varepsilon$, where d is the max metric (with the metric on each C_A also the max metric).

Fix $\sigma > 0$. Let $\varepsilon = \frac{\sigma}{2^k}$ find the corresponding $\delta > 0$. Let \bar{b} be an element of the randomization such that $|\mathbb{E}\psi_{\bar{\gamma}}(\bar{b}, \bar{a}) - \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a})| < \delta$ for every $\bar{\gamma} \in 2^k$.

If we let B be the element of $\prod_{A \in \mathcal{Z}} C_A$ given by $B(A)(\bar{\gamma}) = \frac{1}{r_A} \mathbb{E}\xi_{\bar{\gamma}, A}(\bar{b}, \bar{a})$, then by the choice of \bar{b} we have that $d(f(B), P) < \delta$, and so by the choice of δ , we have that $d(B, f^{-1}(P)) < \varepsilon$. Let B' be some element of $f^{-1}(P)$ with $d(B, B') < \varepsilon$.

For each $(\bar{\gamma}, A) \in Q$, let $Y_{\bar{\gamma}, A} = \{\omega \in \Omega : \xi_{\bar{\gamma}, A}(\bar{b}(\omega), \bar{a}(\omega))\}$. Note that $Y_{\bar{\gamma}, A}$ is automatically a subset of X_A .

What we have now is that for each $(\bar{\gamma}, A) \in Q$, $|\mu(Y_{\bar{\gamma}, A}) - B'(A)(\bar{\gamma})| < r_A \varepsilon = \frac{r_A \sigma}{2^k}$. (Note that $\mu(Y_{\bar{\gamma}, A}) = \mathbb{E}\xi_{\bar{\gamma}, A}(\bar{b}, \bar{a})$.)

We are going to construct an element \bar{c} by changing the values of \bar{b} on a small subset of the corresponding measure space, using Lemma 5.5.2 and the fact that in any (metrically

complete) atomless probability algebra, for any element u and any $0 \leq r \leq \mu(u)$, there exists disjoint v_0, v_1 with $v_0 \sqcup v_1 = u$ and $\mu(v_0) = r$.

Now for each $A \in Z$, we can change the value of \bar{b} on X_A on a set of measure at most $r_A \sigma$ so that the measures of each “new $Y_{\bar{\gamma}, A}$ ” is now $B'(A)(\bar{\gamma})$. This is always possible (possibly passing to a larger model) because of the definition of $\xi_{\bar{\gamma}, A}(\bar{x}, \bar{y})$.

The overall change is therefore limited to be on a set of measure at most σ . So we have that there is an element of $\llbracket \bigwedge_{\bar{\gamma} \in 2^k} \mathbb{E}\psi_{\bar{\gamma}}(\bar{x}, \bar{a}) = \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a}) \rrbracket$ that is distance no more than σ away from \bar{b} . Since we can do this for every \bar{b} satisfying that $|\mathbb{E}\psi_{\bar{\gamma}}(\bar{b}, \bar{a}) - \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a})| < \delta$ for every $\bar{\gamma} \in 2^k$, and since we can do this for arbitrarily small $\sigma > 0$, we have that $\llbracket \bigwedge_{\bar{\gamma} \in 2^k} \mathbb{E}\psi_{\bar{\gamma}}(\bar{x}, \bar{a}) = \mathbb{E}\psi_{\bar{\gamma}}(p, \bar{a}) \rrbracket$ is definable.

So $S_n(\bar{a})$ has a network of definable sets and is therefore dictionaric.

The second part of the proposition follows from the fact that any set of the form given in the statement of the proposition—once particular parameters have been chosen—is a union of the kind of definable sets we have constructed here, and is therefore definable, since it is closed. \square

Since T^{eq} of a dictionaric theory T is not necessarily itself dictionaric, we need to be slightly careful about our statements here, as we have not proven anything about the probability algebra sort in a randomization of a dictionaric theory. We can work around this by noting that the above proof automatically generalized to many-sorted theories and that the probability algebra sort is clearly equivalent to the randomization of the sort **2**.

Note that the formulas $G(\bar{x}, \bar{y})$ in Proposition 5.5.3 are certainly not uniformly \bar{y} -definable in general. To see this, note that the type space $S_{\bar{y}}(A)$ is always connected,

but there are formulas $G(\bar{x}, \bar{y})$ of this form such that $G(\bar{x}, \bar{a})$ is empty for some choices of \bar{a} and non-empty for other choices, which by Proposition 2.3.18 implies that $G(\bar{x}, \bar{y})$ cannot be \bar{y} -uniformly definable.

By quantifier elimination and compactness, we get the following characterization of the definable sets in a randomization of a discrete theory.

Corollary 5.5.4. *If T is a discrete theory, then for any set of parameters A in the randomization T^R , any definable set $D \subseteq S_n(A)$ is a Hausdorff limit of definable sets of the form*

$$G(\bar{x}, \bar{a}) = ((\mathbb{E}\varphi_0(\bar{x}, \bar{a}), \dots, \mathbb{E}\varphi_{k-1}(\bar{x}, \bar{a})) \in F),$$

where \bar{a} is a tuple of parameters from A , $F \subseteq [0, 1]^k$ is some closed set, and $\varphi_0, \dots, \varphi_{k-1}$ is some tuple of discrete formulas.

A natural question arising from this corollary is the following.

Question 5.5.5. *What are the uniformly definable families in a randomization of a discrete theory?*

For continuous theories, we will use $\mathbb{E}\varphi$ to mean the expected value of the continuous formula φ .

Proposition 5.5.6. *For any continuous theory T , the randomization T^R is dictionaric.*

Proof. Let A be a set of parameters in T^R , and let $S_{\bar{x}}(A)$ be a type space over these parameters. Fix a finite list of $[0, 1]$ -valued A -formulas $\varphi_0(\bar{x}, \bar{a}), \dots, \varphi_{k-1}(\bar{x}, \bar{a})$, where \bar{a} is some tuple of parameters from A .

Let $F \subseteq [0, 1]^k$ be the set of values (r_0, \dots, r_k) such that $\bigwedge_{i < k} \mathbb{E}\varphi_i(\bar{x}, \bar{a}) = r_i$ is consistent.

Claim: F is a convex set.

Proof of claim. Let r and s be any points in F . Let \bar{b} and \bar{c} be elements of some model \mathfrak{M} containing A such that \bar{b} ‘satisfies’ r and \bar{c} ‘satisfies’ s in the obvious way (i.e. they take on those values on those formulas). We can find an elementary extension $\mathfrak{N} \succ \mathfrak{M}$ which contains, for any $t \in (0, 1)$, an element e_t in its probability algebra sort with the property that for any f in the probability algebra of \mathfrak{M} , $\mu(e_t \sqcap f) = t\mu(f)$. We can therefore construct (possibly in a larger elementary extension) a tuple \bar{g} with the property that $[\bar{g} = \bar{b}] \sqsupseteq e_t$ and $[\bar{g} = \bar{c}] \sqsupseteq \neg e_t$.¹³ Now we have that $\mathbb{E}(\varphi_i(\bar{g}, \bar{a})) = \mathbb{E}(\varphi_i(\bar{b}, \bar{a}) \cdot e_t) + \mathbb{E}(\varphi_i(\bar{c}, \bar{a}) \cdot \neg e_t)$, where we are thinking of e_t as a $\{0, 1\}$ -valued formula on the probability algebra. By the defining characteristic of e_t , we have that $\mathbb{E}(\varphi_i(\bar{b}, \bar{a}) \cdot e_t) = \mathbb{E}(\varphi_i(\bar{b}, \bar{a}))\mathbb{E}(e_t) = \mathbb{E}(\varphi_i(\bar{b}, \bar{a}))t$ and likewise for \bar{c} and $\neg e_t$, giving the required result.

□_{claim}

Let $\iota : S_{\bar{x}}(A) \rightarrow F$ be the obvious projection map.

Claim: For any type $p \in S_{\bar{x}}(A)$, if $G \subseteq F$ is a closed convex set such that $\iota(p) \in \text{int}_F G$, then $\iota^{-1}(G)$ is a definable set containing p in its interior.

Proof of claim. Let \mathfrak{M} be an $|A|^+$ -saturated model containing A . Let \bar{b} be a fixed realization of p in \mathfrak{M} . We need to show that for every $\varepsilon > 0$, there is a $\delta > 0$ such that for every tuple \bar{c} , if $\iota(\text{tp}(\bar{c}/A))$ is distance less than δ from G in F , then \bar{c} is distance less than ε away from $\llbracket (\iota^{-1}(G))(\mathfrak{M}) \rrbracket$ in \mathfrak{M} .

Fix $\varepsilon > 0$, and suppose that $\iota(\text{tp}(\bar{c}/A))$ is close enough to G that for some $t < \frac{\varepsilon}{\text{diam}(T)}$ (where $\text{diam}(T)$ is the diameter of models of T), $t\iota(p) + (1-t)\iota(\text{tp}(\bar{c}/A))$ is in G . There is a Euclidean distance δ which is small enough that this is possible whenever

¹³Note that there is a subtlety here which is that the $[-]$ function operating on closed formulas is not actually part of the language of randomizations of continuous structures. This statement nevertheless makes sense.

$d(\iota(\text{tp}(\bar{c}/A)), G) < \delta$. Also, note that this is always possible because $\iota(p)$ is in the interior of G . This can in general fail to happen if $\iota(p)$ is not in the interior.

Find an elementary submodel $\mathfrak{N} \succ \mathfrak{M}$ of density character $|A|$ containing $A\bar{b}\bar{c}$. By the same argument as in the previous claim and by $|A|^+$ -saturation, there is an element e of the probability algebra of \mathfrak{M} such that for every element f of the probability algebra of \mathfrak{N} , $\mu(e \sqcap f) = t\mu(f)$, and there is a tuple \bar{g} of elements of \mathfrak{M} such that $[\bar{g} = \bar{b}] \sqsupseteq e$ and $[\bar{g} = \bar{c}] \sqsupseteq \neg e$. From this we get that $\iota(\text{tp}(\bar{g}/A)) = t\iota(\text{tp}(\bar{b}/A)) + (1-t)\iota(\text{tp}(\bar{c}/A)) = t\iota(p) + (1-t)\iota(\text{tp}(\bar{c}/A))$ and therefore also that $\text{tp}(\bar{g}/A) \in \iota^{-1}(G)$. We also get that $d(\bar{c}, \bar{g}) \leq t \text{diam}(T) < \varepsilon$, by the choice of t and the fact that $\mathbb{E}(\bar{g} \neq \bar{c}) = 1 - \mathbb{E}(\bar{g} = \bar{c}) \leq 1 - (1-t) = t$.

Since we can do this for any $\varepsilon > 0$, we have that $\iota^{-1}(G)$ is definable. The fact that p is in the interior of $\iota^{-1}(G)$ is immediate. \square_{claim}

Since we can do this for any finite list of formulas and any arbitrarily small closed convex set surrounding $\iota(p)$, we have that for any type $p \in S_n(A)$ and any open neighborhood $U \ni p$, there is a definable set D such that $p \in D \subseteq U$, and so $S_n(A)$ is dictionaric. Since we can do this for any A and any n , we have that T^R is dictionaric as well. \square

Note that again, this argument clearly extends to many-sorted structures, and so by including the imaginary sort $\mathbf{2}$, the statement extends to formulas involving the probability algebra sort.

What is interesting is that the convex structure of the set F really does seem relevant. It is not too hard to construct examples in randomizations of continuous theories (such as DI^R) where not every closed subset of F corresponds to a definable set, in opposition

to the situation with randomizations of discrete theories. This raises a natural question.

Question 5.5.7. *When does a closed subset of F (as defined in the proof of Proposition 5.5.6) correspond to a definable set?*

It would be good to understand the structure of the definable sets in the randomization of a very simple¹⁴ theory with non-trivial definable sets.

Question 5.5.8. *What are the definable sets in $S_1(\text{PS}^R)$ (Definition 2.3.34)?*

¹⁴Simple in the colloquial sense, not the technical sense.

Chapter 6

Approximate Isomorphism and Approximate Categoricity

Introduction

There are many different notions of ‘approximate isomorphism’ in various branches of mathematics. The two that are best known are perhaps the Banach-Mazur distance between Banach spaces and the Gromov-Hausdorff distance between metric spaces. These two—as well as their lesser known cousins the Kadets distance between Banach spaces and the Lipschitz distance between metric spaces—seem to have fruitful interaction with continuous logic,¹ as explored in [BY08b], [Tel10], and [BDNT17]. This chapter is a synthesis of some ideas presented in [BY08b] and [BDNT17].

In [BY08b], Ben Yaacov introduces perturbation systems—a broad notion of approximate isomorphism—in order to generalize an unpublished result of Henson’s, specifically a Ryll-Nardzewski type characterization of Banach space theories that are ‘approximately separably categorical’ with regards to the Banach-Mazur distance. Ben Yaacov’s formalism requires that approximate isomorphisms be witnessed by uniformly continuous

¹There are also more specialized examples—namely the completely bounded Banach-Mazur distance in the context of operator spaces as well as another, unnamed distance (inducing what is referred to as the ‘weak topology’ on the class of operator spaces in question)—which have shown up in the context of model theory of C^* -algebras [GS15].

bijections with uniformly continuous inverses. As such, while it comfortably covers the Banach-Mazur and Lipschitz distances, it cannot accommodate the Gromov-Hausdorff or Kadets distances.

In [BDNT17], Ben Yaacov, Doucha, Nies, and Tsankov generalize Scott analysis to a family of continuous analogs of $\mathcal{L}_{\omega_1\omega}$. Among other results, they use this to exhibit Scott sentences that characterize separable metric structures not only up to isomorphism, but also up to Gromov-Hausdorff or Kadets distance 0. Their formalism does not seem to be able to capture the Banach-Mazur or Lipschitz distances, although the existence of Scott sentences capturing these was shown indirectly by their continuous Lopez-Escobar theorem and together with results in [CDK18]. In this chapter we will show that with a small modification the formalism of [BDNT17] can give Scott sentences for Banach-Mazur and Lipschitz distances, as well as other ‘well behaved’ notions of approximate isomorphism.

All four of the distances mentioned—Banach-Mazur, Kadets, Gromov-Hausdorff, and Lipschitz—can be expressed in terms of ‘correlations,’ i.e. total surjective relations between the structures in question (bijections being a special case of correlations) and some notion of ‘distortion’ that measures how good of an approximation of isomorphism the given correlation is. Our formalism will use this as a starting point, defining distortion in terms of certain appropriate designated collections of formulas, called ‘distortion systems.’ This is a more syntactic way of looking at something very similar to the objects studied in [BY08b], but without the requirement that the correlations in question be functions.

After giving three characterizations of our family of notions of approximate isomorphism and giving an explicit collection of formulas that captures the Banach-Mazur

distance, we will extend the result of [BDNT17] by giving Scott sentences for perturbation systems.

Despite the limitations of [BY08b] and [BDNT17] it seems that they each settled on fairly natural formalisms. As we will see there are two well-behaved classes of distortion systems. The first includes the Banach-Mazur and Lipschitz distances and is roughly the same thing as the class of perturbation systems. The second includes the Gromov-Hausdorff and Kadets distances and fits fairly well into the weak modulus formalism of [BDNT17].

Later we will present an extension of Ben Yaacov’s separable categoricity result to distortion systems in general and one direction of an approximate Morley’s theorem; specifically, we will define an appropriate notion of ‘ Δ -saturation’ for a distortion system Δ , show that if a theory is Δ - κ -categorical for some uncountable κ , then every model of it of density character κ is Δ -saturated, and then show that if every model of density character κ is Δ -saturated for some uncountable κ , then the same is true for any uncountable λ . The difficulty arises in trying to show that two inseparable Δ -saturated structures of the same density character are ‘almost Δ -approximately isomorphic’, where the ‘almost’ is a technical weakening that’s only non-trivial in certain poorly behaved ‘irregular’ distortion systems (all of the four motivating examples are regular). The best we seem to be able to get is that they are ‘potentially almost Δ -approximately isomorphic,’ i.e. almost Δ -approximately isomorphic in a forcing extension in which they are collapsed to being separable.

Finally, we will present some examples of theories with various combinations of exact and approximate categoricity, highlighting a gap in the currently known examples. Our explicit examples are in the context of pure metric spaces with Gromov-Hausdorff and

Lipschitz distances, but we should note that there is an earlier explicit construction due to Tellez of a Banach-Mazur- ω -categorical Banach space that is not ω -categorical [Tel10].

6.1 Approximate Isomorphism

Definition 6.1.1. Fix a language \mathcal{L} with sorts \mathcal{S} , \mathcal{L} -pre-structures \mathfrak{M} and \mathfrak{N} , and tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$ of the same length with elements in the same sorts.

- (i) The *sort-by-sort product of \mathfrak{M} and \mathfrak{N}* , written $\mathfrak{M} \times_{\mathcal{S}} \mathfrak{N}$, is the collection $\bigsqcup_{s \in \mathcal{S}} s(\mathfrak{M}) \times s(\mathfrak{N})$. If \mathcal{L} is single-sorted we will take $\times_{\mathcal{S}}$ to be the ordinary Cartesian product.
- (ii) A *correlation between \mathfrak{M} and \mathfrak{N}* is a set $R \subseteq \mathfrak{M} \times_{\mathcal{S}} \mathfrak{N}$ such that for each sort s , $R \upharpoonright s := R \upharpoonright s(\mathfrak{M}) \times s(\mathfrak{N})$ is a total surjective relation. We will write $\text{cor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ for the collection of correlations between \mathfrak{M} and \mathfrak{N} such that for each index i less than the length of \bar{m} , $(m_i, n_i) \in R$ (for any binary relation we will abbreviate this condition as $(\bar{m}, \bar{n}) \in R$). If \bar{m} and \bar{n} are empty we will write $\text{cor}(\mathfrak{M}, \mathfrak{N})$.
- (iii) An *almost correlation between \mathfrak{M} and \mathfrak{N}* is a correlation between dense sub-pre-structures of \mathfrak{M} and \mathfrak{N} . We will write $\text{acor}(\mathfrak{M}, \bar{n}; \mathfrak{N}, \bar{m})$ for the collection of almost correlations R between \mathfrak{M} and \mathfrak{N} such that $(\bar{m}, \bar{n}) \in R$. \triangleleft

Note that there is no requirement that correlations or almost correlations are closed. This will turn out to be inessential, but it is convenient for constructing them.

Almost correlations are natural to consider for two reasons. If $(\mathfrak{M}, \mathfrak{N}, R)$ is a metric structure in which R is a definable subset of $\mathfrak{M} \times \mathfrak{N}$, then there is a sentence that holds if and only if R is an almost correlation. There is no sentence that holds precisely when

R is a correlation. The other reason is that many constructions, such as back-and-forth constructions, naturally build an enumeration of a dense sub-pre-structure rather than an enumeration of the entire structure in question. This means that when one tries to build a correlation with some kind of iterative construction, one will often only build an almost correlation.

Definition 6.1.2. Let Δ be a collection of (finitary) \mathcal{L} -formulas, and let T be an \mathcal{L} -theory. Let $\mathfrak{M}, \mathfrak{N} \models T$ with $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$.

(i) For any relation $R \subseteq \mathfrak{M} \times_{\mathcal{S}} \mathfrak{N}$, we define the *distortion of R with respect to Δ* as

$$\text{dis}_{\Delta}(R) = \sup\{|\varphi^{\mathfrak{M}}(\bar{m}) - \varphi^{\mathfrak{N}}(\bar{n})| : \varphi \in \Delta, (\bar{m}, \bar{n}) \in R\}.$$

(ii) We define the Δ -*distance between (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n})* as

$$\rho_{\Delta}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = \inf\{\text{dis}_{\Delta}(R) : R \in \text{cor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})\}.$$

If \bar{m} and \bar{n} are empty we will just write $\rho_{\Delta}(\mathfrak{M}, \mathfrak{N})$.

(iii) We define the *almost Δ -similarity between (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n})* as

$$a_{\Delta}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = \inf\{\text{dis}_{\Delta}(R) : R \in \text{acor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})\}.$$

If \bar{m} and \bar{n} are empty we will just write $a_{\Delta}(\mathfrak{M}, \mathfrak{N})$.

(iv) We say that (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n}) are Δ -*approximately isomorphic*, written $(\mathfrak{M}, \bar{m}) \approx_{\Delta} (\mathfrak{N}, \bar{n})$, if $\rho_{\Delta}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = 0$.

(v) We say that (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n}) are *almost Δ -approximately isomorphic* if $a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = 0$. \triangleleft

Note that $a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = 0$ does not necessarily imply that there are dense sub-structures $\mathfrak{M}_0 \subseteq \mathfrak{M}$ and $\mathfrak{N}_0 \subseteq \mathfrak{N}$ such that $\mathfrak{M}_0 \approx_\Delta \mathfrak{N}_0$.

Given that the composition of correlations is a correlation, it is very easy to verify that ρ_Δ is a pseudo-metric and that \approx_Δ is an equivalence relation on the class of models of T . a_Δ in general is somewhat pathological. It can fail the triangle inequality and therefore in particular be different from ρ_Δ . It can even occur that $\rho_\Delta(\mathfrak{M}, \mathfrak{N}) = \infty$ while $a_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$. We will examine an example of this in Section C.4.1. Since it is likely (although currently unknown) that almost approximate isomorphism is not transitive, we won't give it a symbol that suggests an equivalence relation. These difficulties only occur with unfamiliar notions of approximate isomorphism. Later on we will identify two common conditions that each ensure $\rho_\Delta = a_\Delta$.

Here are some familiar examples:

- If Δ is the collection of all \mathcal{L} -formulas, then $\rho_\Delta(\mathfrak{M}, \mathfrak{N}) < \infty$ if and only if $\mathfrak{M} \cong \mathfrak{N}$.
- If T is the empty theory in the empty signature and $\Delta = \{\frac{1}{2}d(x, y)\}$, then $\rho_\Delta = d_{\text{GH}}$, the Gromov-Hausdorff distance.
- If T is the theory of (the unit balls of) Banach spaces, and Δ is the collection of all formulas of the form $\|\sum_i \lambda_i x_i\|$, with $\sum_i |\lambda_i| \leq 1$, then $\rho_\Delta = d_K$, the Kadets distance between \mathfrak{M} and \mathfrak{N} .
- If T is the empty theory in the empty signature and Δ is the collection of all formulas of the form $[\log d(x, y)]_{-r}^r$, then $\rho_\Delta = d_{\text{Lip}}$, the Lipschitz distance.

- If Δ is the collection of formulas that are 1-Lipschitz in any model of T , then

$$\rho_{\Delta}(\mathfrak{M}, \mathfrak{N}) = \inf\{d_H^{\mathfrak{C}}(f(\mathfrak{M}), g(\mathfrak{N})) \mid f : \mathfrak{M} \preceq \mathfrak{C}, g : \mathfrak{N} \preceq \mathfrak{C}\},$$

where d_H is the Hausdorff metric. This is a sort of elementary embedding variant of the Gromov-Hausdorff and Kadets distances.

The Banach-Mazur distance can also be formalized in this way but the specification of Δ is somewhat more complicated, so we will leave it to later.

Proposition 6.1.3. *Fix Δ , a collection of formulas.*

- (i) *For any relation $R \subseteq \mathfrak{M} \times_S \mathfrak{N}$, $\text{dis}_{\Delta}(R) = \text{dis}_{\Delta}(\overline{R})$, where \overline{R} is the metric closure of R (in each sort).*
- (ii) *There is a subset $\Delta_0 \subseteq \Delta$ with $|\Delta_0| = |\mathcal{L}|$, such that $\text{dis}_{\Delta} = \text{dis}_{\Delta_0}$.*

Proof. (i) This follows from the uniform continuity of each $\varphi \in \Delta$.

(ii) Choose a subset which is dense in Δ with regards to uniform convergence. Since the density character of the space of all \mathcal{L} -formulas under uniform convergence is $|\mathcal{L}|$, it is always possible to find such a Δ_0 . □

Given a collection Δ , we can often enlarge it in a natural way without changing the value of ρ_{Δ} .

Definition 6.1.4. If Δ is a collection of formulas, let $\overline{\Delta}$ be the closure of Δ under renaming variables, quantification, 1-Lipschitz connectives, logical equivalence modulo T , and uniformly convergent limits. ◁

Recall that 1-Lipschitz is meant in the sense of the max metric, i.e. a connective $F(\bar{x})$ is 1-Lipschitz if $|x_i - y_i| \leq \varepsilon$ for all i implies $|F(\bar{x}) - F(\bar{y})| \leq \varepsilon$.

Proposition 6.1.5. *For any $\mathfrak{M}, \mathfrak{N} \models T$ and $\bar{m} \in \mathfrak{M}, \bar{n} \in \mathfrak{N}$, $\text{dis}_\Delta(R) = \text{dis}_{\bar{\Delta}}(R)$, so in particular $\rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = \rho_{\bar{\Delta}}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.*

Proof. It's clear that if Δ is any set of formulas and Σ is Δ closed under renaming variables, 1-Lipschitz connectives, logical equivalence modulo T , and uniformly convergent limits, then $\text{dis}_\Sigma \leq \text{dis}_\Delta$. Since $\Sigma \supseteq \Delta$, we also clearly have $\text{dis}_\Delta \leq \text{dis}_\Sigma$, so $\text{dis}_\Sigma = \text{dis}_\Delta$.

Furthermore if $\{\Sigma_i\}_{i < \lambda}$ is some increasing chain of sets of formulas such that $\text{dis}_\Delta = \text{dis}_{\Sigma_i}$ for every $i < \lambda$, then we also have that $\text{dis}_\Delta = \text{dis}_{\Sigma_\lambda}$, where $\Sigma_\lambda = \bigcup_{i < \lambda} \Sigma_i$.

So the only difficulty is showing that quantification is safe. Suppose that $\varphi(\bar{x}, y) \in \Sigma$. We want to show that for any structures $\mathfrak{M}, \mathfrak{N} \models T$, $\bar{m} \in \mathfrak{M}, \bar{n} \in \mathfrak{N}$, and $R \in \text{cor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$, $|\inf_y \varphi^{\mathfrak{M}}(\bar{m}, y) - \inf_y \varphi^{\mathfrak{N}}(\bar{n}, y)| \leq \text{dis}_\Sigma(R)$.

For each $\varepsilon > 0$, find an $a \in \mathfrak{M}$ such that $\varphi^{\mathfrak{M}}(\bar{m}, a) < \inf_y \varphi^{\mathfrak{M}}(\bar{m}, y) + \varepsilon$. Since R is a correlation we can find a $b \in \mathfrak{N}$ such that $(a, b) \in R$, so we must have that

$$\begin{aligned} \inf_y \varphi^{\mathfrak{N}}(\bar{n}, y) &\leq \varphi^{\mathfrak{N}}(\bar{n}, b) < \varphi^{\mathfrak{M}}(\bar{m}, a) + \text{dis}_\Sigma(R) \text{ and} \\ \inf_y \varphi^{\mathfrak{N}}(\bar{n}, y) &< \inf_y \varphi^{\mathfrak{M}}(\bar{m}, a) + \text{dis}_\Sigma(R) + \varepsilon. \end{aligned}$$

By symmetry we have that $|\inf_y \varphi^{\mathfrak{M}}(\bar{m}, y) - \inf_y \varphi^{\mathfrak{N}}(\bar{n}, y)| \leq \text{dis}_\Sigma(R) + \varepsilon$, and since we can do this for any $\varepsilon > 0$ we have that $|\inf_y \varphi^{\mathfrak{M}}(\bar{m}, y) - \inf_y \varphi^{\mathfrak{N}}(\bar{n}, y)| \leq \text{dis}_\Sigma(R)$ as required. Since $x \mapsto -x$ is a 1-Lipschitz connective, we have this for sup as well.

So by iteratively alternating between closing under renaming variables, 1-Lipschitz connectives, logical equivalence modulo T , and uniformly convergent limits on the one

hand and closure under connectives on the other, we can form a chain $\{\Delta_i\}_{i < \omega_1}$ whose union is $\overline{\Delta}$ and which has the property that for every $i < \omega_1$, $\text{dis}_{\Delta_i}(R) = \text{dis}_{\Delta}(R)$. Therefore we have that $\text{dis}_{\Delta}(R) = \text{dis}_{\overline{\Delta}}(R)$, as required. \square

If we require one more thing from Δ we get something more, and we can justify the name ‘approximate isomorphism’ in that it is, naturally, approximately isomorphism:

Proposition 6.1.6. *Let $\Delta = \overline{\Delta}$ be logically complete. If (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n}) are almost Δ -approximately isomorphic (so in particular if $(\mathfrak{M}, \bar{m}) \approx_{\Delta} (\mathfrak{N}, \bar{n})$), then for any non-principal ultrafilter \mathcal{U} on ω , $(\mathfrak{M}, \bar{m})^{\mathcal{U}} \cong (\mathfrak{N}, \bar{n})^{\mathcal{U}}$, and therefore $(\mathfrak{M}, \bar{m}) \equiv (\mathfrak{N}, \bar{n})$.*

Proof. Let $\{R_i\}_{i < \omega}$ be a sequence of closed almost correlations between \mathfrak{M} and \mathfrak{N} such that $(\bar{m}, \bar{n}) \in R_i$ and $\text{dis}_{\Delta}(R_i) \leq 2^{-i}$. Let \mathcal{U} be a non-principal ultrafilter on ω . For each i , let $(\mathfrak{M}, \mathfrak{N}, R_i)$ be a metric structure containing \mathfrak{M} and \mathfrak{N} in different sorts and having distance predicates for the sets $R_i \upharpoonright s \subseteq s(\mathfrak{M}) \times s(\mathfrak{N})$ for each $s \in \mathcal{S}$. Consider the structure $(\mathfrak{M}', \mathfrak{N}', R') = \prod_{i < \omega} (\mathfrak{M}, \mathfrak{N}, R_i) / \mathcal{U}$. Clearly \mathfrak{M}' and \mathfrak{N}' are elementary extensions of \mathfrak{M} and \mathfrak{N} , respectively. Furthermore by \aleph_1 -saturation (in a countable reduct if the language is uncountable) we have that R' is a correlation, rather than just an almost correlation.

Finally for each formula $\varphi \in \Delta$, we have that

$$(\mathfrak{M}, \mathfrak{N}, R_i) \models \sup_{(\bar{x}, \bar{y}) \in R_i} |\varphi^{\mathfrak{M}}(\bar{x}, \bar{m}) - \varphi^{\mathfrak{N}}(\bar{y}, \bar{n})| \leq 2^{-i}.$$

This is expressible because R_i is a definable set. Therefore we have that

$$(\mathfrak{M}', \mathfrak{N}', R') \models \sup_{(\bar{x}, \bar{y}) \in R_i} |\varphi^{\mathfrak{M}'}(\bar{x}, \bar{m}') - \varphi^{\mathfrak{N}'}(\bar{y}, \bar{n}')| \leq 0$$

for each formula $\varphi \in \Delta$. By the logical completeness of Δ , this implies that R' is the graph of an isomorphism between \mathfrak{M}' and \mathfrak{N}' with $(\bar{m}', \bar{n}') \in R'$. \square

So we will give Δ 's with these properties a name.

Definition 6.1.7. A set of formulas Δ is a *distortion system* for T if it is logically complete and closed under renaming variables, quantification, 1-Lipschitz connectives, logical equivalence modulo T , and uniformly convergent limits. \triangleleft

In many of the motivating examples we aren't given a Δ that is already a distortion system and it's not immediately clear whether or not $\bar{\Delta}$ will be logically complete. There is an easy test, however.

Definition 6.1.8. A collection of \mathcal{L} -formulas Δ is *atomically complete* if, after closing under renaming of variables, any quantifier free type p is entirely determined by the values of $\varphi(p)$ for $\varphi \in \Delta$. \triangleleft

Proposition 6.1.9. *If Δ is atomically complete, then $\bar{\Delta}$ is a distortion system.*

Proof. Clearly we only need to show that $\bar{\Delta}$ is logically complete. Let φ be an atomic formula, and let $r < s$ be real numbers.

Claim: There is a $\bar{\Delta}$ -formula ψ and real numbers $u < v$ such that $\varphi \leq r \vdash \psi < u$ and $\varphi \geq s \vdash \psi > v$.

Proof of claim: Let p be a quantifier free type with $\varphi(p) \leq r$. By compactness there must exist a finite list χ_1, \dots, χ_k of Δ -formulas and an $\varepsilon_p > 0$ such that $\llbracket |\chi_1 - \chi_1(p)| \leq \varepsilon_p \rrbracket \cap \dots \cap \llbracket |\chi_k - \chi_k(p)| \leq \varepsilon_p \rrbracket$ is disjoint from $\llbracket \varphi \geq s \rrbracket$. In particular $\xi_p = |\chi_1 - \chi_1(p)| \uparrow \dots \uparrow |\chi_k - \chi_k(p)|$ is a $\bar{\Delta}$ -formula. By compactness there is a finite list p_1, \dots, p_ℓ of quantifier free types such that $\llbracket \xi_{p_1} < \varepsilon_{p_1} \rrbracket \cup \dots \cup \llbracket \xi_{p_\ell} < \varepsilon_{p_\ell} \rrbracket$ covers

$\llbracket \varphi \leq r \rrbracket$. Furthermore we still have that $\llbracket \xi_{p_1} \leq \varepsilon_{p_1} \rrbracket \cup \dots \cup \llbracket \xi_{p_\ell} \leq \varepsilon_{p_\ell} \rrbracket$ is disjoint from $\llbracket \varphi \geq s \rrbracket$. Let $\theta = (\xi_{p_1} - \varepsilon_{p_1}) \downarrow \dots \downarrow (\xi_{p_\ell} - \varepsilon_{p_\ell})$, and note that this a $\overline{\Delta}$ -formula. We have that $\llbracket \varphi \leq r \rrbracket \subseteq \llbracket \theta < 0 \rrbracket$, and by compactness there must be some $\delta > 0$ such that $\llbracket \varphi \geq s \rrbracket \subseteq \llbracket \theta > \delta \rrbracket$, as required. \square_{claim}

Let Σ be the set of formulas φ that are in prenex form with quantifier free part ψ such that ψ is a maximum of minimums of formulas of the form $\alpha - r$, $r - \alpha$, or r where α is an atomic formula and r a real number. Σ is logically complete in any signature.

Fix a type $p \in S_n(T)$, and consider a formula $\varphi \in \Sigma$ such that $p \vdash \varphi \leq 0$. Fix $\varepsilon > 0$, and let the quantifier free part of φ be $\psi = \max_{i < k} \min_{j < \ell(i)} \chi_{i,j}$. Define formulas $\chi'_{i,j}$ like so:

- If $\chi_{i,j} = r$, then $\chi'_{i,j} = r$ and $\delta_{i,j} = 1$.
- If $\chi_{i,j} = \alpha - r$, find a $\overline{\Delta}$ -formula η and real numbers $u < v$ such that $\llbracket \alpha \leq r \rrbracket \subseteq \llbracket \eta < u \rrbracket$ and $\llbracket \alpha \geq r + \varepsilon \rrbracket \subseteq \llbracket \eta > v \rrbracket$, and set $\chi'_{i,j} = \eta - u$, and let $\delta_{i,j} = \frac{v-u}{2}$.
- If $\chi_{i,j} = r - \alpha$, find a $\overline{\Delta}$ -formula η and real numbers $u < v$ such that $\llbracket \alpha \leq r - \varepsilon \rrbracket \subseteq \llbracket \eta < u \rrbracket$ and $\llbracket \alpha \geq r \rrbracket \subseteq \llbracket \eta > v \rrbracket$, and set $\chi'_{i,j} = v - \eta$, and let $\delta_{i,j} = \frac{v-u}{2}$.

Set $\delta = \min_{i,j} \delta_{i,j}$ and $\psi' = \max_{i < k} \min_{j < \ell(i)} \chi'_{i,j}$. Then let φ' be ψ' with the same quantifiers that φ has. Now we have constructed a $\overline{\Delta}$ -formula, φ' such that for any type q , if $q \vdash \varphi \leq 0$, then $q \vdash \varphi'$, and if $q \vdash \varphi' \leq \delta$, then $q \vdash \varphi \leq \varepsilon$. Since we can do this for any $\varphi \in p$ and any $\varepsilon > 0$, we have that p is entirely determined by $\{\varphi : p \vdash \varphi \in \overline{\Delta}\}$, as required. \square

Now is a convenient time to introduce the following notions:

Definition 6.1.10. Let Δ be a distortion system for T .

- (i) We say that Δ is *regular* if there is an $\varepsilon > 0$ such that for any models $\mathfrak{M}, \mathfrak{N} \models T$, any almost correlation $R \in \text{acor}(\mathfrak{M}, \mathfrak{N})$ with $\text{dis}_\Delta(R) < \varepsilon$, and any $\delta > 0$, there exists a correlation $S \in \text{cor}(\mathfrak{M}, \mathfrak{N})$ such that $S \supseteq R$ and $\text{dis}_\Delta(S) \leq \text{dis}_\Delta(R) + \delta$.
- (ii) We say that Δ is *functional* if there is an $\varepsilon > 0$ such that for any models $\mathfrak{M}, \mathfrak{N} \models T$ and any closed $R \in \text{acor}(\mathfrak{M}, \mathfrak{N})$, if $\text{dis}_\Delta(R) < \varepsilon$, then R is the graph of a uniformly continuous bijection between \mathfrak{M} and \mathfrak{N} with uniformly continuous inverse.
- (iii) We say that Δ is *uniformly uniformly continuous* or *u.u.c.* if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for any models $\mathfrak{M}, \mathfrak{N} \models T$ and any almost correlation $R \in \text{acor}(\mathfrak{M}, \mathfrak{N})$, $\text{dis}_\Delta(R^{<\delta}) \leq \text{dis}_\Delta(R) + \varepsilon$. \triangleleft

Obviously functional and u.u.c. distortion systems are regular. It's easy to construct regular distortion systems that are neither by 'gluing' together functional and u.u.c. distortion systems, such as a two-sorted theory in which both sorts are metric spaces and we simultaneously consider Gromov-Hausdorff distance on the first sort and Lipschitz distance on the second sort.

Functional distortion systems are essentially the same as Ben Yaacov's perturbations. The Gromov-Hausdorff and Kadets distances are u.u.c., and u.u.c. distortion systems are natural generalizations of the Gromov-Hausdorff and Kadets distances. Furthermore one can show that in some common cases the back-and-forth metrics of [BDNT17] must be either isomorphism itself or be equivalent to ρ_Δ for Δ , a u.u.c. distortion system.

Proposition 6.1.11. *Let Δ be a distortion system.*

- (i) *If Δ is regular, then for any (\mathfrak{M}, \bar{n}) and (\mathfrak{N}, \bar{n}) , $\rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.*

(ii) Δ is functional if and only if there is an $\varepsilon > 0$ such that for any $\delta > 0$ there is a formula $\varphi(x, y) \in \Delta$ such that for any $\mathfrak{M} \models T$ and $a, b \in \mathfrak{M}$,

- $\varphi^{\mathfrak{M}}(a, a) = 0$ and
- if $\varphi^{\mathfrak{M}}(a, b) < \varepsilon$, then $d^{\mathfrak{M}}(a, b) < \delta$.

(iii) Δ is u.u.c. if and only if it is uniformly uniformly continuous as a set of formulas, i.e. there is a single modulus $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ (continuous and with $\alpha(0) = 0$) such that for any $\bar{a}, \bar{b} \in \mathfrak{M} \models T$, $|\varphi^{\mathfrak{M}}(\bar{a}) - \varphi^{\mathfrak{M}}(\bar{b})| \leq \alpha(d^{\mathfrak{M}}(\bar{a}, \bar{b}))$.

Proof. (i) Given any almost correlations witnessing the value of a_{Δ} , regularity immediately gives us correlations witnessing the same value for ρ_{Δ} .

(ii) We will defer the proof of this until later (also labeled Proposition 6.1.11 after Proposition 6.1.20) when machinery is available to make the proof easier.

(iii) The (\Rightarrow) direction follows easily from considering the identity as a correlation on models of T . The (\Leftarrow) direction is obvious. \square

Cauchy sequences in ρ_{Δ} give us a way of constructing a limiting structure.

Lemma 6.1.12. *Let Δ be a distortion system. For every predicate symbol P and every $\varepsilon > 0$ there is a $\delta > 0$ such that if $\rho_{\Delta}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) < \delta$ then $|P^{\mathfrak{M}}(\bar{m}) - P^{\mathfrak{N}}(\bar{n})| < \varepsilon$.*

Proof. This follows from the fact that Δ is logically complete and compactness. \square

Proposition 6.1.13. *Let Δ be a distortion system for T .*

(i) *If $\{\mathfrak{M}_i, \bar{m}_i\}_{i < \omega}$ is a sequence of pre-models of T such that for each $i < \omega$,*

$$\rho_{\Delta}(\mathfrak{M}_i, \bar{m}_i; \mathfrak{M}_{i+1}, \bar{m}_i) < 2^{-i},$$

then there is a pre-structure $\mathfrak{M}_\omega \models T$ with \bar{m}_ω such that for each $i < \omega$,

$$\rho_\Delta(\mathfrak{M}_i, \bar{m}_i; \mathfrak{M}_\omega, \bar{m}_\omega) \leq 2^{-i+1}.$$

Furthermore $\#^{\text{dc}}\mathfrak{M}_\omega \leq \sup_i \#^{\text{dc}}\mathfrak{M}_i$, and if all the \mathfrak{M}_i are metrically compact, then \mathfrak{M}_ω is metrically compact.

(ii) If Δ is regular and the \mathfrak{M}_i are complete structures, then \mathfrak{M}_ω can be taken to be a complete structure.

Proof. (i) For each $i < \omega$ find closed $R_{i+\frac{1}{2}} \in \text{cor}(\mathfrak{M}_i, \bar{m}_i; \mathfrak{M}_{i+1}, \bar{m}_{i+1})$ such that $\text{dis}_\Delta(R_{i+\frac{1}{2}}) < 2^{-i}$.

Let $M_\omega^0 = \{x \in \prod_i M_i : (\forall i)(x(i), x(i+1)) \in R_{i+\frac{1}{2}}\}$. For any $\bar{a} \in M_\omega^0$ and predicate symbol P , define $P^{\mathfrak{M}_\omega^0}(\bar{a})$ to be $\lim_{i \rightarrow \infty} P^{\mathfrak{M}_i}(\bar{a}(i))$. By Lemma 6.1.12 this limit always exists. Furthermore we have that $d^{\mathfrak{M}_\omega}$ is a pseudo-metric on M_ω^0 and that all predicate symbols P obey the correct moduli of continuity for the signature \mathcal{L} . So let \mathfrak{M}_ω be M_ω^0 modded by $d^{\mathfrak{M}_\omega} = 0$, and we have that this is an \mathcal{L} -pre-structure.

Now to see that $\mathfrak{M}_\omega \models T$, we will show that for any restricted formula $\varphi(\bar{x})$, $\varphi^{\mathfrak{M}_\omega}(\bar{a}) = \lim_{i \rightarrow \infty} \varphi^{\mathfrak{M}_i}(\bar{a}(i))$ and furthermore that this convergence is uniform in \bar{a} . We already have that this is true for atomic formulas, and if F is a connective and we've shown that this holds for some tuple $\bar{\varphi}$ of formulas, then it clearly holds for $F(\bar{\varphi})$ as well. So all that we need to do is show that this is true for quantification. Let $\varphi(\bar{x}, y)$ be a formula for which $\varphi^{\mathfrak{M}_\omega}(\bar{a}, b) = \lim_{i \rightarrow \infty} \varphi^{\mathfrak{M}_i}(\bar{a}(i), b(i))$ uniformly in $\bar{a}b$. Fix $\varepsilon > 0$, and find a $j < \omega$ large enough that $|\varphi^{\mathfrak{M}_\omega}(\bar{a}, b) - \varphi^{\mathfrak{M}_i}(\bar{a}(i), b(i))| < \frac{1}{2}\varepsilon$ for all $\bar{a}b$ and $i \geq j$. Now find $c \in \mathfrak{M}_j$ such that $\varphi^{\mathfrak{M}_j}(\bar{a}(j), c) < \inf_y \varphi^{\mathfrak{M}_j}(\bar{a}(j), y) + \frac{1}{2}\varepsilon$. Extend c to a sequence $e(i) \in M_\omega^0$ such

that $e(j) = c$. Now we have that

$$\varphi^{\mathfrak{M}_\omega}(\bar{a}, e) < \varphi^{\mathfrak{M}_j}(\bar{a}(j), e(j)) + \frac{1}{2}\varepsilon,$$

so in particular

$$\inf_y \varphi^{\mathfrak{M}_\omega}(\bar{a}, y) \leq \varphi^{\mathfrak{M}_\omega}(\bar{a}, e) < \inf_y \varphi^{\mathfrak{M}_j}(\bar{a}(j), y) + \varepsilon.$$

Since we can do this for φ and $-\varphi$ and for any $\varepsilon > 0$, we have shown the required property for $\inf_y \varphi(\bar{x}, y)$. Therefore, by induction, this holds for all restricted formulas and thus, by uniform convergence, for all formulas.

Since sentences are formulas we have that for any φ such that $\mathfrak{M}_i \models \varphi \leq 0$ for all $i < \omega$, $\mathfrak{M}_\omega \models \varphi \leq 0$ as well.

To show the bound on the density character of \mathfrak{M}_ω , assume that $\#^{\text{dc}}\mathfrak{M}_i \leq \kappa$ (for some infinite κ) for each $i < \omega$, and for each such i find a dense subset $X_i \subseteq \mathfrak{M}_i$ of cardinality $\leq \kappa$. For each $a \in X_i$, choose some $b_a \in M_\omega^0$ such that $b_a(i) = a$, and let $X = \{b_a : (\exists i)a \in X_i\}$. Since d uniformly converges this clearly is a dense subset of \mathfrak{M}_ω as well.

For the statement regarding compact structures, Lemma 6.1.12 implies that the sequence of underlying metric spaces of the \mathfrak{M}_i are converging in the Gromov-Hausdorff metric to the underlying metric space of \mathfrak{M}_ω . It is well known that a sequence of compact metric spaces converging in the Gromov-Hausdorff metric converges to a compact metric space, so the result follows.

(ii) This follows easily from the fact that the correlations between the \mathfrak{M}_i and \mathfrak{M}_ω

are almost correlations between the \mathfrak{M}_i and the completion of \mathfrak{M}_ω . \square

Corollary 6.1.14. *Let Δ be a distortion system for T . Let $\text{PreMod}(T, \leq \kappa)$ be the collection of pre-models of T with density character $\leq \kappa$, $\text{Mod}(T, \leq \kappa)$ be the collection of models of T with density character $\leq \kappa$, and let $\text{Mod}(T, \leq \omega^-)$ be the collection of compact models of T . For every κ ,*

(i) $(\text{PreMod}(T, \leq \kappa), \rho_\Delta)$ is a complete pseudo-metric space.

(ii) If Δ is regular then $(\text{Mod}(T, \leq \kappa), \rho_\Delta)$ is a complete pseudo-metric space.

(iii) $(\text{Mod}(T, \leq \omega^-), \rho_\Delta)$ is a complete metric space. Furthermore, for compact models,

$$\rho_\Delta = a_\Delta.$$

Proof. (i) and (ii) are obvious from the previous proposition.

(iii) The furthermore statement follows from the fact that almost correlations between compact structures are actually correlations, by compactness.

The furthermore statement in part (i) of the previous proposition implies that $(\text{Mod}(T, \leq \omega^-), \rho_\Delta)$ is complete, so we just need to show that for compact structures, $\mathfrak{M} \approx_\Delta \mathfrak{N}$ if and only if $\mathfrak{M} \cong \mathfrak{N}$. But this is easy: Take an ultraproduct of the structures $(\mathfrak{M}, \mathfrak{N}, R_i)$ where R_i is the correlation taken as a definable subset of \mathfrak{M} and \mathfrak{N} with $\text{dis}_\Delta(R_i) \leq 2^{-i}$. Then you will get a structure of the form $(\mathfrak{M}, \mathfrak{N}, R_\omega)$ with R_ω the graph of an isomorphism. \square

There is an example of an irregular distortion system Δ for a theory T such that $(\text{Mod}(T, \leq \kappa), \rho_\Delta)$ is not complete, see Section C.4.1.

6.1.1 Induced Metrics on Type Space

Any distortion system Δ for some theory T naturally induces a family of topometrics on the type spaces of T . We will define this for one-sorted theories for readability, but the extension to many sorted theories is obvious.

Definition 6.1.15. Let Δ be a distortion system for T . For each λ and any $p, q \in S_\lambda$, let

$$\delta_\Delta^\lambda(p, q) = \inf\{\rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) : \bar{m} \models p, \bar{n} \models q\}.$$

We will typically drop the λ when it is clear from context. ◁

We're using δ instead of d to emphasize that δ is not the natural counterpart of the ordinary d -metric. Instead it is the natural counterpart of

$$\delta(p, q) = \begin{cases} 0 & p = q \\ \infty & p \neq q \end{cases},$$

i.e. a metric encoding equality of types. Later on there will be a metric, d_Δ , derived from δ_Δ that plays an analogous role to the d -metric on types. In some very special cases, such as Gromov-Hausdorff distance or Kadets distance, we will have $\delta_\Delta = d_\Delta$. This in turn will entail some nice properties of Δ -approximate isomorphism.

δ_Δ^λ enjoys the following properties.

Proposition 6.1.16. *Let Δ be a distortion system for T .*

- (i) $\delta_\Delta^\lambda(p, q) = \sup_{\varphi \in \Delta} |\varphi(p) - \varphi(q)|$, where $\varphi(r)$ means the unique value of φ entailed by the type r .

- (ii) δ_Δ^λ is a topometric on $S_\lambda(T)$, i.e. it is lower semi-continuous and refines the topology.
- (iii) (Monotonicity) For any $p, q \in S_{\lambda+\alpha}(T)$, if $p', q' \in S_\lambda(T)$ are restrictions of p and q to the first λ variables, then $\delta_\Delta^\lambda(p', q') \leq \delta_\Delta^{\lambda+\alpha}(p, q)$.
- (iv) For any $p, q \in S_\lambda(T)$ and any permutation $\sigma : \lambda \rightarrow \lambda$, $d_\Delta^n(p, q) = d_\Delta^n(\sigma p, \sigma q)$, where σr is the type $r(x_{\sigma(0)}, x_{\sigma(1)}, \dots)$.
- (v) (Extension) For any $p, q \in S_\lambda(T)$ and $p' \in S_{\lambda+\alpha}(T)$ extending p , there exists a $q' \in S_{\lambda+\alpha}(T)$ extending q such that $d_\Delta^\lambda(p, q) = d_\Delta^{\lambda+\alpha}(p', q')$.
- (vi) For any infinite λ , $\delta_\Delta^\lambda(p, q) = \sup \delta_\Delta^n(p', q')$, where p' and q' range over restrictions of p and q to finite tuples of variables.

Proof. (ii)-(iv) and (vi) all follow immediately from (i).

It will be easier to prove (i) once we have (v). To see that (v) holds, let $(\mathfrak{M}_i, \bar{m}_i, \mathfrak{N}_i, \bar{n}_i, R_i)$ be structures such that $\bar{m}_i \models p$, $\bar{n}_i \models q$, R_i is a correlation between \mathfrak{M}_i and \mathfrak{N}_i with $(\bar{m}_i, \bar{n}_i) \in R_i$ and $\text{dis}_\Delta(R_i) \leq \delta_\Delta^\lambda(p, q) + 2^{-i}$. By taking an ultraproduct of these structures we get an exact witness, i.e. a structure $(\mathfrak{M}, \bar{m}, \mathfrak{N}, \bar{n}, R)$ with $\bar{m} \models p$, $\bar{n} \models q$, and R , a correlation between \mathfrak{M} and \mathfrak{N} with $\text{dis}_\Delta(R) = \delta_\Delta^\lambda(p, q)$.

Now by compactness $(\mathfrak{M}, \bar{m}, \mathfrak{N}, \bar{n}, R)$ has a \aleph_1 -saturated elementary extension $(\mathfrak{M}', \bar{m}', \mathfrak{N}', \bar{n}', R')$ in which \mathfrak{M}' realizes $p'(\bar{x}, \bar{m}')$ with some tuple \bar{a} . By \aleph_1 -saturation, R' is a correlation, so we have that there is some tuple \bar{b} with $(\bar{a}, \bar{b}) \in R'$. So we can take $q' = \text{tp}(\bar{n}'\bar{b})$ and get the required extension.

Now for (i). It's clear that $\delta_\Delta^\lambda(p, q) \geq \sup_{\varphi \in \Delta} |\varphi(p) - \varphi(q)|$. All we need, given p, q with $\sup_{\varphi \in \Delta} |\varphi(p) - \varphi(q)| = \varepsilon$, is to build a structure $(\mathfrak{M}, \bar{m}, \mathfrak{N}, \bar{n}, R)$ witnessing that

$\delta_{\Delta}^{\lambda}(p, q) \leq \varepsilon$, but this is almost immediate from the extension property, by a ‘back-and-forth Henkin construction.’ So we have shown (i) as well. \square

As it happens, given a family of metrics $\{\delta^n\}_{n < \omega}$ satisfying some of these properties we can find a distortion system giving the same metrics. Again this is for a single-sorted theory but the extension to many-sorted theories is obvious.

Metrics satisfying these properties are very similar to the ‘perturbation (pre-)systems’ of [BY08b], but what we require here is more than a perturbation pre-system and less than a perturbation system.

Proposition 6.1.17. *Suppose that $\{\delta^n\}_{n < \omega}$ is a family of topometrics on $S_n(T)$ such that:*

- *(Monotonicity) For any $p, q \in S_{n+1}(T)$, if $p', q' \in S_n(T)$ are restrictions of p and q to the first n variables, then $\delta^n(p', q') \leq \delta^{n+1}(p, q)$.*
- *For any $p, q \in S_n(T)$ and any permutation $\sigma : n \rightarrow n$, $\delta^n(p, q) = \delta^n(\sigma p, \sigma q)$, where σr is the type $r(x_{\sigma(0)}, x_{\sigma(1)}, \dots, x_{\sigma(n-1)})$.*
- *(Extension) For any $p, q \in S_n(T)$ and $p' \in S_{n+1}(T)$ extending p there exists a $q' \in S_{n+1}(T)$ extending q such that $\delta^n(p, q) = \delta^{n+1}(p', q')$.*

Then there is a distortion system $\Delta(\delta)$ (namely the collection of δ -1-Lipschitz formulas) such that $\delta = \delta_{\Delta(\delta)}$.

Furthermore for any distortion system Δ we have that $\Delta = \Delta(\delta_{\Delta})$.

Proof. Let $\Delta(\delta)$ be the collection of formulas that are 1-Lipschitz with regards to δ (in the relevant variables). Note that by the monotonicity property $\Delta(\delta)$ is closed under

adding dummy variables (i.e. if $\varphi(\bar{x})$ is δ -1-Lipschitz in $S_n(T)$, then it is δ -1-Lipschitz in $S_{n+1}(T)$). $\Delta(\delta)$ is also clearly closed under renaming variables, 1-Lipschitz connectives, logical equivalence modulo T , and uniformly convergent limits. So all we need to do is show that $\Delta(\delta)$ is closed under quantification and that it is logically complete.

By a result of Ben Yaacov [BY10b], if (X, d) is a compact topometric space and $F, G \subseteq X$ are disjoint closed sets with $d_{\text{inf}}(F, G) := \inf\{d(x, y) : x \in F, y \in G\} > \varepsilon$, then there is a 1-Lipschitz continuous function $f : X \rightarrow [0, \varepsilon]$ such that $F \subseteq f^{-1}(0)$ and $G \subseteq f^{-1}(\varepsilon)$. This in particular implies that for any type $p \in S_n(T)$, p is determined entirely by $\{\varphi : p \vdash \varphi \in \Delta(\delta)\}$, i.e. that $\Delta(\delta)$ is logically complete. Another corollary of his result, as pointed out by him, is that $d(x, y) = \sup\{|f(x) - f(y)| : f : X \rightarrow \mathbb{R} \text{ 1-Lipschitz continuous}\}$ in any compact topometric space X , so we clearly have that $\delta^n = \delta_{\Delta(\delta)}^n$ for each $n < \omega$.

To see that $\Delta(\delta)$ is closed under quantification, let $\varphi(\bar{x}, y) \in \Delta(\delta)$. It is sufficient to show that $\inf_y \varphi(\bar{x}, y) \in \Delta(\delta)$. For any $p, q \in S_n(T)$ (where $n = |\bar{x}|$), find a realization $\bar{a} \models p$ in the monster model, and then find b such that $\models \varphi(\bar{a}, b) = \inf_y \varphi(\bar{a}, y)$. Let $p' = \text{tp}(\bar{a}b)$. Find q' extending q such that $\delta(p', q') = \delta(p, q)$. Now we have that $|\varphi(p') - \varphi(q')| \leq \delta(p', q') = \delta(p, q)$, implying that $q'(\bar{x}, y) \models \varphi(\bar{x}, y) \leq \varphi(p') + \delta(p, q)$. This implies that $q(\bar{x}) \models \inf_y \varphi(\bar{x}, y) \leq \varphi(p') + \delta(p, q)$, so by symmetry we have that $|\inf_y \varphi(p, y) - \inf_y \varphi(q, y)| \leq \delta(p, q)$. Since we can do this for any p, q we have that $\inf_y \varphi(\bar{x}, y)$ is δ -1-Lipschitz and $\inf_y \varphi(\bar{x}, y) \in \Delta(\delta)$.

Therefore $\Delta(\delta)$ is a distortion system.

For the furthermore part, we clearly have that every Δ -formula is δ_{Δ} -1-Lipschitz. We just need to show that every δ_{Δ} -1-Lipschitz formula is a Δ -formula.

Let $\varphi(\bar{x})$ be a δ_{Δ} -1-Lipschitz formula. Pick $p \in S_n(T)$ (where $n = |\bar{x}|$). We have that

$\varphi(q) \leq \varphi(p) + \delta(p, q)$ for all q , so in particular

$$\varphi(q) \leq \varphi(p) + \sup_{\psi \in \Delta} |\psi(p) - \psi(q)|,$$

for all q . For any $\varepsilon > 0$, by compactness there must be a finite set $\{\psi_1, \dots, \psi_k\} \subset \Delta$ such that

$$\varphi(q) \leq \varphi(p) + |\psi_1(p) - \psi_1(q)| \uparrow \dots \uparrow |\psi_k(p) - \psi_k(q)| + \varepsilon,$$

for all q . $|\psi_1(p) - \psi_1(\bar{x})| \uparrow \dots \uparrow |\psi_k(p) - \psi_k(\bar{x})| + \varepsilon$ is a Δ -formula, so we have shown that

$$\varphi(q) = \inf\{\psi(q) : \psi \in \Delta, \psi \geq \varphi\},$$

for all q . Now for each $i < \omega$, by compactness there must be a finite set $\{\psi_1^i, \dots, \psi_{k(i)}^i\} \subset \Delta$ such that

$$\varphi(q) \leq \psi_1(q) \downarrow \dots \downarrow \psi_k(q) < \varphi(q) + 2^{-i}$$

for all q .

So if we let $\chi_0 = \psi_1^0 \downarrow \dots \downarrow \psi_{k(0)}^0$ and $\chi_{j+1} = \chi_j \downarrow \psi_1^{j+1} \downarrow \dots \downarrow \psi_{k(j+1)}^{j+1}$, we get that $\{\chi_j\}_{j < \omega}$ is a sequence of Δ -formulas that uniformly converges to φ , so $\varphi \in \Delta$, as required. \square

6.1.2 Theories of Approximate Isomorphism

Implicit in a lot of the arguments so far has been the fact that if Δ is a distortion system for T , then for any ε there is a first-order theory whose models are precisely structures $(\mathfrak{M}, \mathfrak{N}, R)$, with R a closed almost correlation between \mathfrak{M} and \mathfrak{N} such that

$\text{dis}_\Delta(R) \leq \varepsilon$. This is how notions of approximate isomorphism are typically presented, at least implicitly. There is some kind of ambient structure relating \mathfrak{M} and \mathfrak{N} witnessing a certain degree of closeness, such as a mutual embedding into a larger structure or a certain special kind of function between them. We will give a precise characterization of these theories in our context and show that Δ can be reconstructed from them.

Definition 6.1.18. If Δ is a distortion system for T , then for any $\varepsilon \in [0, \infty]$ let $\text{Th}(\Delta, \varepsilon)$ be the common theory of all structures of the form $(\mathfrak{M}, \mathfrak{N}, R)$ with $\mathfrak{M}, \mathfrak{N} \models T$ and closed $R \in \text{cor}(\mathfrak{M}, \mathfrak{N})$ with $\text{dis}_\Delta(R) \leq \varepsilon$, where R is taken as a family of definable subsets of $s(\mathfrak{M}) \times s(\mathfrak{N})$ for sorts $s \in \mathcal{S}$. \triangleleft

Proposition 6.1.19. *Let Δ be a distortion system for T . For any $\varepsilon \in [0, \infty]$, a triple $(\mathfrak{M}, \mathfrak{N}, R) \models \text{Th}(\Delta, \varepsilon)$ if and only if R is a closed almost correlation between \mathfrak{M} and \mathfrak{N} and $\text{dis}_\Delta(R) \leq \varepsilon$.*

Proof. (\Rightarrow) Assume that $(\mathfrak{M}, \mathfrak{N}, R) \models \text{Th}(\Delta, \varepsilon)$. Clearly we have that $\mathfrak{M}, \mathfrak{N} \models T$ and that for all $(\bar{n}, \bar{m}) \in R$, $\delta_\Delta(\text{tp}(\bar{n}), \text{tp}(\bar{m})) \leq \varepsilon$, so the only thing to really check is that R is an almost correlation. This follows because it is equivalent to the first-order axiom schema consisting of

$$\sup_{x \in s(\mathfrak{M})} \inf_{y \in s(\mathfrak{N})} R_s(x, y) \uparrow \sup_{y \in s(\mathfrak{N})} \inf_{x \in s(\mathfrak{M})} R_s(x, y)$$

for each sort $s \in \mathcal{S}$, where R_s is the distance predicate for the set $R \upharpoonright s$.

(\Leftarrow) Take an \aleph_1 -saturated elementary extension of $(\mathfrak{M}', \mathfrak{N}', R') \succeq (\mathfrak{M}, \mathfrak{N}, R)$. By \aleph_1 -saturation, R' is a closed correlation between \mathfrak{M}' and \mathfrak{N}' , and we still have that $\text{dis}_\Delta(R') \leq \varepsilon$, so by definition $(\mathfrak{M}', \mathfrak{N}', R') \models \text{Th}(\Delta, \varepsilon)$, thus by elementarity, $(\mathfrak{M}, \mathfrak{N}, R) \models \text{Th}(\Delta, \varepsilon)$. \square

Unsurprisingly, Δ is recoverable from the theories $\text{Th}(\Delta, \varepsilon)$.

Proposition 6.1.20. *Fix a theory T , and suppose that $\{A_\varepsilon\}_{\varepsilon \in [0, \infty]}$ is a family of first-order theories that satisfy the following conditions:*

- For every ε , every model of A_ε is of the form $(\mathfrak{M}, \mathfrak{N}, R)$ with \mathfrak{M} and \mathfrak{N} models of T where R is a family of distance predicates R_s on $s(\mathfrak{M}) \times s(\mathfrak{N})$.
- A triple $(\mathfrak{M}, \mathfrak{N}, R)$ is a model of A_∞ if and only if R is a closed almost correlation between \mathfrak{M} and \mathfrak{N} , namely if for each sort s , it satisfies

$$\sup_{x \in s(\mathfrak{M})} \inf_{y \in s(\mathfrak{N})} R_s(x, y) \uparrow \sup_{y \in s(\mathfrak{N})} \inf_{x \in s(\mathfrak{M})} R_s(x, y).$$

- A triple $(\mathfrak{M}, \mathfrak{N}, R)$ is a model of A_0 if and only if R is the graph of an isomorphism between \mathfrak{M} and \mathfrak{N} .
- For each $\varepsilon < \delta$, A_ε logically entails A_δ and $\bigcup_{\delta > \varepsilon} A_\delta$ is logically equivalent to A_ε .
- (Symmetry) If $(\mathfrak{M}, \mathfrak{N}, R) \models A_\varepsilon$, then $(\mathfrak{N}, \mathfrak{M}, R^{-1}) \models A_\varepsilon$, where $R^{-1} := \{(y, x) : (x, y) \in R\}$.
- (Composition) For every $\varepsilon, \delta > 0$ if $(\mathfrak{M}, \mathfrak{N}, R) \models A_\varepsilon$ and $(\mathfrak{N}, \mathfrak{D}, S) \models A_\delta$ and $(\mathfrak{M}, \mathfrak{N}, R)$ and $(\mathfrak{N}, \mathfrak{D}, S)$ are \aleph_1 -saturated, then $(\mathfrak{M}, \mathfrak{D}, \overline{S \circ R}) \models A_{\varepsilon + \delta}$, where $\overline{S \circ R}$ is understood to mean the family of distance predicates of the metric closure of the relation $S \circ R$.
- (Sub-structure) If $(\mathfrak{M}, \mathfrak{N}, R) \models A_\varepsilon$ and $\mathfrak{M}' \preceq \mathfrak{M}$ and $\mathfrak{N}' \preceq \mathfrak{N}$ are elementary sub-structures such that $(\mathfrak{M}, \mathfrak{N}, R)$, \mathfrak{M}' , and \mathfrak{N}' are all \aleph_1 -saturated, and $R' = R \upharpoonright \mathfrak{M}' \times_{\mathcal{S}} \mathfrak{N}'$ is a correlation, then $(\mathfrak{M}', \mathfrak{N}', R') \models A_\varepsilon$.

Then there is a distortion system Δ such that $A_\varepsilon \equiv \text{Th}(\Delta, \varepsilon)$ for every $\varepsilon \in [0, \infty]$.

Proof. First we will show that $\{A_\varepsilon\}_{\varepsilon \in [0, \infty]}$ induces a family of topometrics $\{\delta_A^n\}_{n < \omega}$ satisfying the conditions of Proposition 6.1.17. Then we will show that $A_\varepsilon = \text{Th}(\Delta(\delta), \varepsilon)$ for every $\varepsilon \in [0, \infty]$.

Let $\delta_A(p, q) = \inf\{\varepsilon : (\mathfrak{M}, \mathfrak{N}, R) \models A_\varepsilon, R \in \text{cor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}), \bar{m} \models p, \bar{n} \models q\}$.

It's clear that $\delta_A(p, q) \geq 0$ and that $\delta_A(p, p) = 0$. By symmetry we have that $\delta_A(p, q) = \delta_A(q, p)$.

Pick $p, q, r \in S_n(T)$. Pick $(\mathfrak{M}, \mathfrak{N}, R)$ witnessing $\delta_A(p, q) \leq \varepsilon$ and $(\mathfrak{N}', \mathfrak{D}, S)$ witnessing that $\delta_A(q, r) \leq \delta$. By passing to elementary extensions we can find triples $(\mathfrak{M}', \mathfrak{N}'', R')$ and $(\mathfrak{N}'', \mathfrak{D}', S')$ with tuples \bar{m} and \bar{n} such that $\mathfrak{M}' \ni \bar{m} \models p$ and $\mathfrak{N}'' \ni \bar{n} \models q$, and $\mathfrak{D}' \ni \bar{o} \models r$ such that all structures involved are \aleph_1 -saturated. By composition we have that $(\mathfrak{M}', \mathfrak{D}', \overline{S \circ R}) \models A_{\varepsilon + \delta}$, witnessing that $\delta(p, r) \leq \varepsilon + \delta$. Since we can do this for any ε and δ , we have that $\delta(p, r) \leq \delta(p, q) + \delta(q, r)$.

Finally by taking ultraproducts of witnesses it's clear that if $\delta(p, q) = 0$ then $p = q$.

So we have that δ^n are metrics. They are clearly lower semi-continuous, again by taking ultraproducts of relevant witnesses, so they are a family of topometrics. Now we just need to verify the other conditions of Proposition 6.1.17. Monotonicity and permutation invariance are both clear. For extension, suppose that $(\mathfrak{M}, \mathfrak{N}, R)$ is an exact witness for the value of $\delta^n(p, q)$, i.e. there are $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$ such that $\bar{m} \models p$ and $\bar{n} \models q$ and $(\bar{m}, \bar{n}) \in R$, a correlation. Then by passing to a saturated enough elementary extension we can find a such that $\models p'(\bar{m}, a)$. By picking some b correlated to a and picking $q' = \text{tp}(\bar{n}b)$, we get the required extension.

So we have that Proposition 6.1.17 applies and $\Delta(\delta_A)$ is a distortion system with $\delta_A = \delta_{\Delta(\delta_A)}$.

So now clearly by construction we have that for any $\varepsilon \in [0, \infty]$, $A_\varepsilon \vdash \text{Th}(\Delta(\delta_A), \varepsilon)$. So all we need to do is show that $\text{Th}(\Delta(\delta_A), \varepsilon) \vdash A_\varepsilon$. Let $(\mathfrak{M}, \mathfrak{N}, R) \models \text{Th}(\Delta(\delta_A), \varepsilon)$. Assume that $(\mathfrak{M}, \mathfrak{N}, R)$ is \aleph_1 -saturated, by passing to an elementary extension if necessary. By construction for every pair of finite tuples, $(\bar{m}, \bar{n}) \in R$ there exists $(\mathfrak{A}_{\bar{m}, \bar{n}}, \mathfrak{B}_{\bar{m}, \bar{n}}, S_{\bar{m}, \bar{n}}) \models A_\varepsilon$ such that $(\bar{m}, \bar{n}) \in S_{\bar{m}, \bar{n}}$. Let \mathcal{F} be the filter on $R^{<\omega}$ ordered by extensions of tuples, and let \mathcal{U} be an ultrafilter extending \mathcal{F} . Take the ultraproduct $(\mathfrak{A}', \mathfrak{B}', S') = \prod_{(\bar{m}, \bar{n}) \in R^{<\omega}} (\mathfrak{A}_{\bar{m}, \bar{n}}, \mathfrak{B}_{\bar{m}, \bar{n}}, S_{\bar{m}, \bar{n}}) / \mathcal{U}$, and assume that this is \aleph_1 -saturated (taking an elementary extension if necessary). By construction we have that $\mathfrak{M} \preceq \mathfrak{A}'$, $\mathfrak{N} \preceq \mathfrak{B}'$, and $R = S' \upharpoonright \mathfrak{M} \times_S \mathfrak{N}$ is a correlation, so by the sub-structure property we have that $(\mathfrak{M}, \mathfrak{N}, R) \models A_\varepsilon$. Since we can do this for any theory completing $\text{Th}(\Delta, \varepsilon)$, we have that $\text{Th}(\Delta, \varepsilon) \vdash A_\varepsilon$, so $\text{Th}(\Delta, \varepsilon) = A_\varepsilon$ as required. \square

Now we can finally tie up a loose end.

Proposition 6.1.11. *Let Δ be a distortion system.*

(ii) *Δ is functional if and only if there is an $\varepsilon > 0$ such that for any $\delta > 0$ there is a formula $\varphi(x, y) \in \Delta$ such that for any $\mathfrak{M} \models T$ and $a, b \in \mathfrak{M}$,*

- $\varphi^{\mathfrak{M}}(a, a) = 0$ and
- if $\varphi^{\mathfrak{M}}(a, b) < \varepsilon$, then $d^{\mathfrak{M}}(a, b) < \delta$.

Proof. (\Rightarrow) For any $\varphi(x, y) \in \Delta$, let $\chi_\varphi(x, y) = \frac{1}{2}|\varphi(x, y) - \varphi(x, x)|$, and note that χ_φ is always a Δ -formula.

Assume that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for any $\varphi \in \Delta$ there exists $a, b \in \mathfrak{M} \models T$, either $\varphi^{\mathfrak{M}}(a, a) \neq 0$ or $(\varphi^{\mathfrak{M}}(a, b) < \varepsilon$ and $d^{\mathfrak{M}}(a, b) \geq \delta)$.

In particular this implies that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for any $\varphi_1, \dots, \varphi_k \in \Delta$ there exists $a, b \in \mathfrak{M} \models T$, $\chi_{\varphi_1}^{\mathfrak{M}}(a, b) \uparrow \dots \uparrow \chi_{\varphi_k}^{\mathfrak{M}}(a, b) < \varepsilon$ and $d^{\mathfrak{M}}(a, b) \geq \delta$.

By compactness this implies that for any $\varepsilon > 0$ there is a $\delta > 0$ and $c, e \in \mathfrak{N} \models T$ such that for every $\varphi \in \Delta$, $\chi_{\varphi}^{\mathfrak{N}}(c, e) \leq \varepsilon$ and $d^{\mathfrak{N}}(c, e) \geq \delta$. In particular this implies that $\delta_{\Delta}(\text{tp}(ce), \text{tp}(cc)) \leq 2\varepsilon$. So we can build a structure witnessing this and we have that Δ cannot be functional at $\varepsilon > 0$. Since we can do this at any $\varepsilon > 0$, Δ is not functional.

(\Leftarrow) Assume that there is an $\varepsilon > 0$ such that for any $\delta > 0$ there is a formula $\varphi_{\delta}(x, y) \in \Delta$ such that for any $a, b \in \mathfrak{M} \models T$, if $\varphi_{\delta}^{\mathfrak{M}}(a, b) < \varepsilon$, then $d^{\mathfrak{M}}(a, b) < \delta$. Pick $0 < \gamma < \varepsilon$, and let $(\mathfrak{M}, \mathfrak{N}, R)$ be an \aleph_1 -saturated model of $\text{Th}(\Delta, \gamma)$. For each $a, b \in \mathfrak{M}$ and $c \in \mathfrak{N}$ with $(a, c), (b, c) \in R$, we have that $|\varphi_{\delta}^{\mathfrak{M}}(a, b) - \varphi_{\delta}^{\mathfrak{N}}(c, c)| \leq \gamma$, so in particular $\varphi_{\delta}^{\mathfrak{M}}(a, b) \leq \gamma < \varepsilon$. So we have that $d^{\mathfrak{M}}(a, b) < \delta$. Since we can do this for any $\delta > 0$, we have that $d^{\mathfrak{M}}(a, b) = 0$ and $a = b$.

Therefore R is the graph of a bijection in every \aleph_1 -saturated model of $\text{Th}(\Delta, \gamma)$. By Proposition 2.3.80, this implies that it is actually the graph of a definable bijection, so this fact must be true in every model of $\text{Th}(\Delta, \gamma)$. By compactness this implies that there is a modulus α_{γ} such that in every model of $\text{Th}(\Delta, \gamma)$, R and R^{-1} are α_{γ} -uniformly continuous. So we have that every closed $R \in \text{acor}(\mathfrak{M}, \mathfrak{N})$ with $\text{dis}_{\Delta}(R) < \varepsilon$ is the graph of a uniformly continuous bijection with uniformly continuous inverse, therefore Δ is functional. \square

A corollary of this is that when checking functionality of Δ it is enough to check closed correlations, rather than closed almost correlations.

6.2 Special Cases

Here we will examine a few specific cases of notions of approximate isomorphism arising from distortion systems.

6.2.1 Elementary and Finitary Gromov-Hausdorff-Kadets Distances

Definition 6.2.1. Let Δ_0 and Δ_1 be distortion systems.

We say that Δ_1 *uniformly dominates* Δ_0 if for every $\varepsilon > 0$ there is a $\delta > 0$ such that if $\delta_{\Delta_1}(p, q) < \delta$ then $\delta_{\Delta_0}(p, q) < \varepsilon$. We may also say that Δ_0 is *coarser* than Δ_1 or that Δ_1 is *finer* than Δ_0 .

If Δ_0 and Δ_1 uniformly dominate each other we say that they are *uniformly equivalent*. ◁

Note that Δ_1 uniformly dominates Δ_0 if and only if the collection of Δ_0 -formulas are u.u.c. with regards to δ_{Δ_1} .

Proposition 6.2.2. *Fix a signature \mathcal{L} .*

- (i) *There is a collection of formulas, $e\text{GHK}_0$, such that for any \mathcal{L} -theory T , $e\text{GHK}_0$ generates the finest u.u.c. distortion system for T , up to uniform equivalence. Furthermore $\delta_{e\text{GHK}_0} = d$, the d -metric on types.*
- (ii) *If \mathcal{L} is countable then there is a collection of formulas, $f\text{GHK}_0$, such that for any \mathcal{L} -theory T , $f\text{GHK}_0$ generates the coarsest distortion system for T , up to uniform equivalence.*

Proof. (i) For any \mathcal{L} -formula φ , let $\chi_\varphi(\bar{x}) = \inf_{\bar{y}} \varphi(\bar{y}) + d(\bar{x}, \bar{y})$. χ_φ has the property that it is 1-Lipschitz in any \mathcal{L} -structure and furthermore that if φ is 1-Lipschitz in every model of T , then $T \models \chi_\varphi = \varphi$. Let $\text{eGHK}_0 = \{\chi_\varphi : \varphi \in \mathcal{L}\}$. Note that $\varphi \in \text{eGHK}_0$ for any sentence φ .

By a previously mentioned result of Ben Yaacov [BY10b], for any types p, q in the same complete theory, $d(p, q) = \delta_{\text{eGHK}_0}(p, q)$ (and for types in different complete theories $\delta_{\text{eGHK}_0}(p, q) = \infty$). This implies that eGHK_0 uniformly dominates any u.u.c. distortion system.

(ii) Let $\{P_i\}_{i < \omega}$ be an enumeration of all atomic \mathcal{L} -predicates. For each i , if P_i 's codomain interval is $[a, b]$, let $r_i = 1 + |a| \uparrow |b|$.

Let $\text{fGHK}_0 = \left\{ \frac{1}{2^{i r_i}} P_i \right\}_{i < \omega}$. This is clearly atomically complete. By Lemma 6.1.12, in any theory T , any distortion system for T uniformly dominates $\overline{\text{fGHK}_0}$. \square

Definition 6.2.3. Fix a theory T and an enumeration of atomic \mathcal{L} -formulas.

- $\text{eGHK} = \overline{\text{eGHK}_0}$, as defined in the proof of Proposition 6.2.2. $\rho_{\text{eGHK}}(\mathfrak{M}, \mathfrak{N})$ is the *elementary Gromov-Hausdorff-Kadets distance between \mathfrak{M} and \mathfrak{N}* .
- If T is countable, let $\text{fGHK} = \overline{\text{fGHK}_0}$, as defined in the proof of Proposition 6.2.2. $\rho_{\text{fGHK}}(\mathfrak{M}, \mathfrak{N})$ is ‘the’ *finitary Gromov-Hausdorff-Kadets distance between \mathfrak{M} and \mathfrak{N}* . \triangleleft

Clearly ρ_{fGHK} depends on the choice of enumeration, but \approx_{fGHK} does not.

Proposition 6.2.4. $\mathfrak{M} \approx_{\text{fGHK}} \mathfrak{N}$ if and only if for every $\varepsilon > 0$, finite collection $\mathcal{S}_0 \subseteq \mathcal{S}$ of sorts, and finite collection Σ of atomic \mathcal{L} -formulas whose variables are from sorts in \mathcal{S}_0 , there exists a correlation $R \subseteq \mathfrak{M} \times_{\mathcal{S}_0} \mathfrak{N}$ such that $\text{dis}_\Sigma(R) < \varepsilon$.

You may have noticed that the condition in Proposition 6.2.4 makes sense in uncountable languages. Indeed there is a canonical uniform structure analog of ρ_{fGHK} , given by a family of pseudo-metrics. It is possible to develop the whole theory of distortion systems with this more general context in mind, similar to [Iov99]. Rather than a single collection of formulas we would need a directed family of collections of formulas. In the absence of motivating examples we opted to develop this simpler framework.

Clearly, for the empty signature, ρ_{fGHK} is uniformly equivalent to ρ_{GH} . This is true of ρ_K in the theory of (unit balls of) Banach spaces as well, justifying the name.

Proposition 6.2.5. *Let T be the theory of unit balls of Banach spaces. ρ_{fGHK} is uniformly equivalent to d_K , the Kadets distance.*

Proof. Fix two unit ball Banach space structures \mathfrak{M} and \mathfrak{N} .

Claim: For every $\varepsilon > 0$, there is a $\delta > 0$ such that if $\rho_{fGHK}(\mathfrak{M}, \mathfrak{N}) < \delta$, then there exists a correlation $R \in \text{cor}(\mathfrak{M}, \mathfrak{N})$ such that:

- $(\mathbf{0}^{\mathfrak{M}}, \mathbf{0}^{\mathfrak{N}}) \in R$
- If $(a, c), (b, d) \in R$, then $(\frac{1}{2}(a + b), \frac{1}{2}(c + d)) \in R$.
- If $(a, b) \in R$ then $(-a, -b) \in R$.
- (For complex Banach spaces) If $(a, b) \in R$ then $(e^i a, e^i b) \in R$.
- If $(a, c), (b, d) \in R$, then $|\|a - b\|_{\mathfrak{M}} - \|c - d\|_{\mathfrak{N}}| \leq \varepsilon$.

Proof of claim: Fix $\varepsilon > 0$. There is a $\delta > 0$ such that for any $S \in \text{cor}(\mathfrak{M}, \mathfrak{N})$ with $\text{dis}_{fGHK}(S) \leq \delta$, then:

- For every $(a, b) \in S$, $|\|a\|_{\mathfrak{M}} - \|b\|_{\mathfrak{N}}| \leq \frac{1}{5}\varepsilon$.

- For every $(a, c), (b, d) \in S$, $|\|a - b\|_{\mathfrak{M}} - \|c - d\|_{\mathfrak{N}}| \leq \frac{1}{5}\varepsilon$, $|\|a + b\|_{\mathfrak{M}} - \|c + d\|_{\mathfrak{N}}| \leq \frac{1}{5}\varepsilon$,
and (if the Banach spaces are complex) $|\|e^i a - b\|_{\mathfrak{M}} - \|e^i c - d\|_{\mathfrak{N}}| \leq \frac{1}{5}\varepsilon$.
- For every $(a, d), (b, e), (c, f) \in S$, $|\|\frac{1}{2}(a + b) - c\|_{\mathfrak{M}} - \|\frac{1}{2}(d + e) - f\|_{\mathfrak{N}}| \leq \frac{1}{5}\varepsilon$.

Such a δ exists because this is a finite list of atomic formulas. Consider the correlation

$$R = \left\{ (a, b) : (\exists (c, d) \in S) \|a - c\|_{\mathfrak{M}} \leq \frac{1}{5}\varepsilon \wedge \|b - d\|_{\mathfrak{N}} \leq \frac{1}{5}\varepsilon \right\}.$$

Now we have what we want:

- $(\mathbf{0}^{\mathfrak{M}}, \mathbf{0}^{\mathfrak{N}})$ is in R because if $(a, \mathbf{0}^{\mathfrak{N}}) \in S$, then the distance between a and $\mathbf{0}^{\mathfrak{M}}$ is $\leq \frac{1}{5}\varepsilon$.
- If $(a, c), (b, d), (\frac{1}{2}(a + b), e) \in S$, then the distance between e and $\frac{1}{2}(c + d)$ is $\leq \frac{1}{5}\varepsilon$.
- If $(a, b), (-a, c) \in S$, then the distance between c and $-b$ is $\leq \frac{1}{5}$.
- (For a complex Banach space) If $(a, b), (e^i a, c) \in S$, then the distance between c and $e^i b$ is $\leq \frac{1}{5}\varepsilon$.
- If $(a, c), (b, d) \in R$, then there are $(a', c'), (b', d') \in S$ each distance $\frac{1}{5}$ to the corresponding element. By several applications of the triangle inequality this implies that $|\|a - b\|_{\mathfrak{M}} - \|c - d\|_{\mathfrak{N}}| \leq \frac{4}{5}\varepsilon < \varepsilon$. □_{claim}

By iterating the second bullet point in the claim we get the following: For any $n < \omega$, and any $(a_1, b_1), \dots, (a_{2^n}, b_{2^n}) \in R$, $(2^{-n} \sum_i a_i, 2^{-n} \sum_i b_i) \in R$. By using duplicates we get that if $\lambda_1, \dots, \lambda_n$ are a sequence of positive dyadic rationals with $\sum_i \lambda_i = 1$, then for any $(a_1, b_1), \dots, (a_m, b_m) \in R$, $(\sum_i \lambda_i a_i, \sum_i \lambda_i b_i) \in R$.

Using the third and fourth bullet points we get that if $\lambda_1, \dots, \lambda_n$ are a sequence of numbers of the form ad with $a = \pm e^{ik}$ with $k < \omega$ and d a positive dyadic rational ≤ 1 , if $\sum_i |\lambda_i| = 1$, then for any $(a_1, b_1), \dots, (a_m, b_m) \in R$, $(\sum_i \lambda_i a_i, \sum_i \lambda_i b_i) \in R$, so in particular

$$\left| \left\| \sum_i \lambda_i a_i \right\|_{\mathfrak{M}} - \left\| \sum_i \lambda_i b_i \right\|_{\mathfrak{N}} \right| \leq \varepsilon.$$

Since λ_i of this form are dense in the set of all coefficients γ_i satisfying $\sum_i |\gamma_i| = 1$ (for complex Banach spaces this relies on the fact that e^i is an irrational rotation), and by Fact 3.4 in [BY14], this implies that $d_K(\mathfrak{M}, \mathfrak{N}) \leq \varepsilon$.

The other direction follows from the minimality of ρ_{fGHK} under uniform domination. □

6.2.2 Banach-Mazur Distance

Difficulty arises with the Banach-Mazur distance in that the witnessing correlations are bijections between the entire Banach spaces in question. To deal with this we will use Ben Yaacov's emboudment concept [BY08a] to encode the entire Banach space as a bounded structure.

We could in principle do this more cleanly using the full logic for unbounded structures in [BY08a], but then we would have to re-develop the machinery of distortion systems in that broader context. We should note that Ben Yaacov does develop a theory of perturbations for unbounded metric structures in [BY08a].

Definition 6.2.6. An *embounded Banach space structure* is a metric structure $\{\mathfrak{M}, d, \mathbf{0}, \infty, P, S_r\}_{r \in K}$, where \mathfrak{M} is a Banach space over the field $K \in \{\mathbb{R}, \mathbb{C}\}$ together with an additional point ∞ .

Let $\theta(x) = \frac{x}{1+x}$. The metric is

$$d^{\mathfrak{M}}(x, y) = \frac{\theta(\|x - y\|)}{1 + \|x\| \downarrow \|y\|}$$

for $x, y \neq \infty$ and $d(x, \infty) = \frac{1}{1+\|x\|}$. We also have

$$P(x, y, z) = \frac{\theta(\|x + y - z\|)}{1 + \|x\| \uparrow \|y\| \uparrow \|z\|}$$

if $x, y, z \neq \infty$, and $P(x, y, z) = 0$ if any of x, y, z are ∞ . Finally,

$$S_r(x, y) = \frac{\theta(\|rx - y\|)}{1 + \|x\| \uparrow \|y\|}$$

if $x, y \neq \infty$, and $S_r(x, y) = 0$ if either x or y are ∞ . ◁

Note that even though the language as stated is uncountable it is actually interdefinable with a finite sub-language² in unit ball Banach space structures.

In order to describe the formulas that capture the Banach-Mazur distance we will freely use the following facts:

Fact 6.2.7. *There is a theory whose models are precisely embeddings of Banach spaces. Let T be that theory.*

1. *For any $r > 0$ there is a formula that is the distance predicate of the ball of (norm) radius r , B_r , in any model of T .*
2. *For any $r > 0$ there is a formula that defines $\|x\|$ in B_r in any model of T .*

²For $K = \mathbb{R}$, $S_{\frac{1}{2}}$ is sufficient, and for $K = \mathbb{C}$, we also need S_i .

3. For any $r > 0$ there is a formula that defines the function $+ : B_r^2 \rightarrow B_{2r}$ in any model of T .

4. For any $s \in K$ and any $r > 0$ there is a formula that defines the function $(x \mapsto sx) : B_r \rightarrow B_{|s|r}$ in any model of T .

Also note that inclusion maps between definable sets are always uniformly definable.

This lemma follows immediately from the previous set of facts, although a careful proof would be slightly involved.

Lemma 6.2.8. *Let T be the theory of emboundments of Banach space structures. There are formulas that define the quantities*

- $\varphi_r(x, y, z) = [r - r^{-2} \log(\|x\| \uparrow \|y\| \uparrow \|z\|)]_0^1 \cdot [(2 - r^{-1}) \log\|x + y - z\|]_{-r}^r$ and
- $\psi_{r,s}(x, y) = [r - r^{-2} \log(\|x\| \uparrow \|y\|)]_0^1 \cdot [(2 - r^{-1}) \log\|sx - y\|]_{-r}^r$

in any model of T for any real $r > 0$ and $t \in K$, where these quantities are understood to be 0 if any of their inputs are ∞ .

To clarify what we're doing, intuitively we're after expressions of the form $2 \log\|\dots\|$, with \dots replaced with various linear combinations, to capture the Banach-Mazur distance. These are unbounded so we need to use bounded approximations. Unlike with the Lipschitz metric, we can't just use $[2 \log\|\dots\|]_{-r}^r$ as they aren't by themselves continuous on the emboundment (specifically the problem is at ∞). Given this, we need the more complicated expressions of the form $[\dots]_0^1$ as cutoff functions which are 1 whenever the maximum norm of the inputs is less than $e^{r^3 - r^2}$ and which are 0 whenever it is greater than e^{r^3} . The specific form of these cutoffs and the coefficient $2 - r^{-1}$ are chosen so that the (\Leftarrow) direction of the next result, Proposition 6.2.10, will work.

Definition 6.2.9. Let BM_0 be the formulas in Lemma 6.2.8 allowing substitution of the constant $\mathbf{0}$. Let $\text{BM} = \overline{\text{BM}_0}$. \triangleleft

To see that BM_0 is atomically complete, note that by choosing large enough values for r (so that the cutoff function is 1) and appropriate values of s and t , the formulas in BM_0 clearly fix the values of $d(x, y)$, $P(x, y, z)$, $S_r(x, y)$ for any $x, y, z \in \{a, b, c, \mathbf{0}\}$ with a, b, c any triple of elements of a structure. The only unclear thing is determining the value of $d(a, \infty)$, but this is $1 - d(a, \mathbf{0})$, so BM_0 is atomically complete. Therefore BM is a distortion system.

Proposition 6.2.10. *Let \mathfrak{M} and \mathfrak{N} be embeddings of the Banach spaces X and Y , respectively. For $R \in \text{cor}(\mathfrak{M}, \mathfrak{N})$ a closed correlation, $\text{dis}_{\text{BM}}(R) \leq \varepsilon < \infty$ if and only if R is the graph of a linear bijection between X and Y (together with the tuple $(\infty^{\mathfrak{M}}, \infty^{\mathfrak{N}})$) such that $\|R\| \leq \sqrt{e^\varepsilon}$ and $\|R^{-1}\| \leq \sqrt{e^\varepsilon}$.*

Proof. (\Rightarrow) Assume that R is a closed correlation between \mathfrak{M} and \mathfrak{N} with $\text{dis}_{\text{BM}}(R) \leq \varepsilon < \infty$.

Pick $m \in \mathfrak{M}$, and assume that $(m, \infty^{\mathfrak{N}}) \in R$. Consider the formula $\psi_{r,s}(x, \mathbf{0})$. Since $\psi_{r,s}^{\mathfrak{N}}(\infty^{\mathfrak{N}}, \mathbf{0}^{\mathfrak{N}}) = 0$ for any r, s , we have that $|\psi_{r,s}^{\mathfrak{M}}(m, \mathbf{0}^{\mathfrak{M}})| \leq \varepsilon$ for any r, s . Assume that $a \neq \infty^{\mathfrak{M}}$. When r is large enough (relative to the choice of s) we have that $\psi_{r,s}^{\mathfrak{M}}(m, \mathbf{0}^{\mathfrak{M}}) = [(2 - r^{-1}) \log \|sm\|]_{-r}$, but this quantity is unbounded in r and s (even if $a = \mathbf{0}^{\mathfrak{M}}$), so we must have that $a = \infty^{\mathfrak{M}}$.

By symmetry $(\infty^{\mathfrak{M}}, \infty^{\mathfrak{N}})$ is the only instance of a pair containing either copy of ∞ .

Pick $a, b \in X$, and consider $d, e, f \in Y$ such that $(a, d), (b, e), (a + b, f) \in R$. For any sufficiently large r , we have that $\varphi_r^{\mathfrak{M}}(a, b, a + b) = -r$. Assume that $f \neq d + e$. Then for

any sufficiently large r , we have that

$$\varphi_r^{\mathfrak{m}}(d, e, f) = (2 - r^{-1}) \log \|b + e - f\| > \log \|b + e - f\| > -\infty.$$

Since we can choose r arbitrarily large, this contradicts that $\text{dis}_{\text{BM}}(R) \leq \varepsilon$. Therefore $f = d + e$.

The same argument shows that if $(a, b), (ua, c) \in R$, then $b = ua$. In particular this implies that $(\mathbf{0}^{\mathfrak{m}}, \mathbf{0}^{\mathfrak{n}})$ is the only correlation involving a copy of $\mathbf{0}$.

Therefore, by symmetry, $R \upharpoonright X \times Y$ is the graph of a linear bijection.

Now consider $a \in X \setminus \{\mathbf{0}\}$ and $b \in Y \setminus \{\mathbf{0}\}$ such that $(a, b) \in R$. Considering the formula $\varphi_r(x, \mathbf{0}, \mathbf{0})$ for sufficiently large r , we have that

$$(2 - r^{-1}) |\log \|a\| - \log \|b\|| = (2 - r^{-1}) \left| \log \frac{\|a\|}{\|b\|} \right| \leq \varepsilon.$$

Since we can do this for arbitrarily large r , this yields

$$2 \left| \log \frac{\|a\|}{\|b\|} \right| \leq \varepsilon.$$

So we have that $\|R\| \leq e^{\varepsilon/2} = \sqrt{e^\varepsilon}$ and by symmetry $\|R^{-1}\| \leq \sqrt{e^\varepsilon}$.

(\Leftarrow) Let A be a linear bijection between X and Y such that $\|R\|, \|R^{-1}\| \leq \sqrt{e^\varepsilon}$. Let $R = A \cup \{\infty^{\mathfrak{m}}, \infty^{\mathfrak{n}}\}$. We need to compute $\text{dis}_{\text{BM}}(R) = \text{dis}_{\text{BM}_0}(R)$. Since we know that $(\mathbf{0}^{\mathfrak{m}}, \mathbf{0}^{\mathfrak{n}}) \in R$, we only need to check the formulas in Lemma 6.2.8.

Let $(a, e), (b, f), (c, g) \in R$, and consider the quantity $|\varphi_r^{\mathfrak{m}}(a, b, c) - \varphi_r^{\mathfrak{n}}(e, f, g)|$. If any of a, b, c, e, f, g are ∞ then this is 0, so assume that none of them are. To estimate

this we will need the following facts:

$$|x_1 y_1 - x_0 y_0| \leq |x_1 - x_0|(|y_0| \uparrow |y_1|) + |y_1 - y_0|(|x_0| \uparrow |x_1|) \text{ and}$$

$$|[x]_a^b - [y]_a^b| \leq |x - y|.$$

Applying these to this case gives

$$|\varphi_r^{\mathfrak{M}}(a, b, c) - \varphi_r^{\mathfrak{N}}(e, f, g)| \leq r^{-2} \left| \log \frac{\|a\| \uparrow \|b\| \uparrow \|c\|}{\|e\| \uparrow \|f\| \uparrow \|g\|} \right| r$$

$$+ (2 - r^{-1}) \left| \log \frac{\|a + b - c\|}{\|e + f - g\|} \right|,$$

since the first term in φ_r can have magnitude at most 1 and the second term can have magnitude at most r . Now finally note that we must have

$$\left| \log \frac{\|a\| \uparrow \|b\| \uparrow \|c\|}{\|e\| \uparrow \|f\| \uparrow \|g\|} \right| \leq \frac{\varepsilon}{2} \text{ and}$$

$$\left| \log \frac{\|a + b - c\|}{\|e + f - g\|} \right| \leq \frac{\varepsilon}{2}.$$

Putting this all together gives

$$|\varphi_r^{\mathfrak{M}}(a, b, c) - \varphi_r^{\mathfrak{N}}(e, f, g)| \leq r^{-1} \frac{\varepsilon}{2} + (2 - r^{-1}) \frac{\varepsilon}{2} = \varepsilon.$$

The same proof works for $\psi_{r,s}$, so we have that $\text{dis}_{\text{BM}}(R) \leq \varepsilon$. □

Corollary 6.2.11. *If X and Y are Banach spaces and \mathfrak{M} and \mathfrak{N} are their corresponding embeddings, then $\rho_{\text{BM}}(\mathfrak{M}, \mathfrak{N}) = d_{\text{BM}}(X, Y)$.*

Proof. Clearly we have $d_{\text{BM}}(X, Y) \leq \rho_{\text{BM}}(\mathfrak{M}, \mathfrak{N})$. To get the other direction, let $A :$

$X \rightarrow Y$ be a linear bijection with $\|A\| \cdot \|A^{-1}\| \leq e^\varepsilon$. If we set $r = \sqrt{\frac{\|A^{-1}\|}{\|A\|}}$, then we have that rA is a linear bijection between X and Y with $\|rA\| \leq \sqrt{e^\varepsilon}$ and $\|(rA)^{-1}\| \leq \sqrt{e^\varepsilon}$, so we get $\rho_{\text{BM}}(\mathfrak{M}, \mathfrak{N}) \leq d_{\text{BM}}(X, Y)$. \square

6.2.3 Approximate Isomorphism in Discrete Logic

Perhaps surprisingly, the concept of a distortion system is non-trivial in discrete logic.

Definition 6.2.12. A *stratified language* is a language \mathcal{L} together with a designated sequence of sub-languages $\{\mathcal{L}_i\}_{i < \omega}$ whose union is \mathcal{L} . (Note that the sub-languages may have fewer sorts than the full language.)

In the context of a stratified language \mathcal{L} , two \mathcal{L} -structures $\mathfrak{M}, \mathfrak{N}$ are said to be *approximately isomorphic*, written $\mathfrak{M} \approx_{\mathcal{L}} \mathfrak{N}$, if $\mathfrak{M} \upharpoonright \mathcal{L}_i \cong \mathfrak{N} \upharpoonright \mathcal{L}_i$ for every $i < \omega$. In general let $\rho_{\mathcal{L}}(\mathfrak{M}, \mathfrak{N}) = 2^{-i}$ where i is the largest such that $\mathfrak{M} \upharpoonright \mathcal{L}_i \cong \mathfrak{N} \upharpoonright \mathcal{L}_i$ but $\mathfrak{M} \upharpoonright \mathcal{L}_{i+1} \not\cong \mathfrak{N} \upharpoonright \mathcal{L}_{i+1}$, or 0 if no such i exists.

We may drop the subscript \mathcal{L} if the relevant stratified language is clear by context. \triangleleft

Clearly $\rho_{\mathcal{L}}$ is a pseudo-metric on \mathcal{L} -structures.

Proposition 6.2.13. *Let T be a discrete first-order theory (i.e. every predicate is $\{0, 1\}$ -valued in every model of T), and let Δ be a distortion system for T .*

- (i) *For every finite set $\mathcal{S}_0 \subseteq \mathcal{S}$ there is an $\varepsilon > 0$ such that if $\text{dis}_{\Delta}(R) < \varepsilon$, then R restricted to the sorts in \mathcal{S}_0 is the graph of a bijection. For every predicate symbol P there is an $\varepsilon_P > 0$ such that whenever $\text{dis}_{\Delta}(R) < \varepsilon_P$, then R is the graph of a bijection that respects R .*

(ii) There is a stratification of \mathcal{L} such that ρ_Δ and $\rho_{\mathcal{L}}$ are uniformly equivalent. In particular $\mathfrak{M} \approx_\Delta \mathfrak{N}$ if and only if $\mathfrak{M} \approx_{\mathcal{L}} \mathfrak{N}$.

Proof. (i) This follows immediately from Lemma 6.1.12.

(ii) Choose ε_P as in part (i) for all predicate symbols. For each $i < \omega$, let \mathcal{S}_i be the set of sorts such that $\varepsilon_{=s} \geq 2^{-i}$. Set \mathcal{L}_i to be the set of all predicate symbols P such that $\varepsilon_P \geq 2^{-i}$ and for every sort s of a variable in P , $\varepsilon_{=s} \geq 2^{-i}$.

Then, for sufficiently small distances, $\rho_{\mathcal{L}}$ and ρ_Δ never differ by more than a factor of 4, so they are uniformly equivalent. \square

Note that Δ for a discrete theory will still contain continuous formulas (since we are implicitly considering it as a continuous theory) and these will be what gives it its structure.

6.3 Scott Sentences for Functional Approximation Fragments

Here we will develop back-and-forth pseudo-metrics, r_α^Δ , for arbitrary distortion systems, an extension of [BDNT17]. In the case of separable structures with functional or u.u.c. distortion systems, r_∞^Δ will be equal to the corresponding ρ_Δ , but for some irregular distortion systems we will show that $r_\infty^\Delta \neq \rho_\Delta$ (in particular because $r_\infty^\Delta \leq a_\Delta < \rho_\Delta$).

As a corollary of this we will explicitly exhibit Scott sentences for Δ -equivalence with functional Δ (which are precisely the same as Ben Yaacov's perturbations [BY08b]). This covers Banach-Mazur equivalence for Banach spaces and Lipschitz equivalence for metric spaces, which were not expressible in the framework of [BDNT17], although the

existence of these was shown indirectly by the continuous Lopez-Escobar theorem in [BDNT17] and results in [CDK18].

Many of the proofs in this section are nearly identical to the corresponding proofs in [BDNT17], so we will only sketch the important parts. We should pause to emphasize that **for bookkeeping purposes in this section we are not treating all variables as interchangeable**. For $n < m$ we are thinking of x_m as being (potentially) ‘more sensitive’ than x_n , so more formulas are allowed to have x_m as a variable than x_n . See section 2 of [BDNT17].

Definition 6.3.1. Let Δ be a collection of formulas closed under renaming variables (typically a distortion system).

For any \mathcal{L} -structures $\mathfrak{M}, \mathfrak{N}$, $\bar{m} \in \mathfrak{M}$, $\bar{n} \in \mathfrak{N}$, and weak modulus Ω , we define the (Δ, Ω) -back-and-forth pseudo-metrics, $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$, as follows:

- $r_0^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ is

$$\sup\{|\psi^{\mathfrak{M}}(\bar{m}) - \psi^{\mathfrak{N}}(\bar{n})| : \psi \in \Delta, \psi \text{ respects } \Omega \text{ in every } \mathcal{L}\text{-structure}\}.$$

- $r_{\alpha+1}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ is

$$\sup_{a \in \mathfrak{M}} \inf_{b \in \mathfrak{N}} r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a; \mathfrak{N}, \bar{n}b) \uparrow \sup_{b \in \mathfrak{N}} \inf_{a \in \mathfrak{M}} r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a; \mathfrak{N}, \bar{n}b).$$

- $r_\lambda^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ is

$$\sup_{\alpha < \lambda} r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}),$$

for λ a limit or ∞ .

◁

This is the analog of Lemma 3.2 in [BDNT17]; the proofs are essentially identical.

Lemma 6.3.2. (i) For fixed $\alpha \in \text{Ord} \cup \{\infty\}$ and k , $r_\alpha^{\Delta, \Omega}$ is a pseudo-metric on the class of all pairs (\mathfrak{M}, \bar{m}) , with $|\bar{m}| = k$.

(ii) For every α , \mathfrak{M} , and $\bar{a}, \bar{b} \in \mathfrak{M}$, $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{a}; \mathfrak{M}, \bar{b}) \leq d^\Omega(\bar{a}, \bar{b})$.

(iii) For every α , \mathfrak{M} , and \mathfrak{N} , and k , the function $(\bar{m}, \bar{n}) \mapsto r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ on pairs of k -tuples is uniformly continuous on $\mathfrak{M}^k \times \mathfrak{N}^k$ with regards to the max metric.

This is the analog of Lemma 3.3 in [BDNT17]; again the proofs are essentially identical.

Lemma 6.3.3. (i) For every $\alpha \leq \beta$, $r_\alpha^{\Delta, \Omega} \leq r_\beta^{\Delta, \Omega}$.

(ii) For every pair of structures $\mathfrak{M}, \mathfrak{N}$ with $\#\text{dc}\mathfrak{M}, \#\text{dc}\mathfrak{N} \leq \kappa$, there is an $\alpha < \kappa^+$ such that $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = r_{\alpha+1}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ for all pairs of tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$, which implies that in fact $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ for all such pairs of tuples.

This is the analog of Proposition 3.4 in [BDNT17]. See [BDNT17] for the definition of *shift increasing*.

Proposition 6.3.4. Let $\mathfrak{M}, \mathfrak{N} \models T$ be separable. For any $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$, $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) < \varepsilon$ if and only if there exists tail-dense sequences $\{a_i\}_{i < \omega} \subseteq \mathfrak{M}$ and $\{b_i\}_{i < \omega} \subseteq \mathfrak{N}$ starting with \bar{m} and \bar{n} , respectively, such that

$$\sup_{n < \omega} r_0^{\Delta, \Omega}(\mathfrak{M}, a_{< n}; \mathfrak{N}, b_{< n}) < \varepsilon,$$

where a sequence is tail-dense if every final segment of it is metrically dense.

Corollary 6.3.5. *Let Δ be a distortion system for T , a theory in a countable language.*

Let Ω be a weak modulus.

- (i) *For any models \mathfrak{M} and \mathfrak{N} and tuples \bar{m} and \bar{n} , we have that $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$. (In particular since $r_\infty^{\Delta, \Omega}$ is a pseudo-metric, this implies that the function $(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \mapsto r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ is 2-Lipschitz in ρ_Δ .)*
- (ii) *If Δ is u.u.c. and Ω is shift increasing with the property that for any $\varphi \in \Delta$, $\varphi(x_0, x_1, x_2, \dots)$ is an Ω -formula, then for any separable models \mathfrak{M} and $\mathfrak{N} \models T$ and tuples \bar{m} and \bar{n} , we have that $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = \rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.*
- (iii) *If Δ is functional then for any sequence $\{\varphi_i\}_{i < \omega} \subset \Delta$ dense in Δ in the uniform norm and if Ω is shift increasing such that for any $i < \omega$ there exists an $n < \omega$ such that $\varphi_i(x_n, x_{n+1}, \dots, x_{n+k})$ is an Ω -formula, then there is an $\varepsilon > 0$ such that for any separable models \mathfrak{M} and $\mathfrak{N} \models T$, if $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) < \varepsilon$, then $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) = \rho_\Delta(\mathfrak{M}, \mathfrak{N})$.*

Proof. (i) If \mathfrak{M} and \mathfrak{N} are separable we can just use the previous proposition.

The idea of the proof is that we can use an almost correlation between \mathfrak{M} and \mathfrak{N} as a back-and-forth strategy. We will proceed by induction, showing that for every α , $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.

$r_0^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ clearly holds, as does the limit case.

Assume that $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ holds for all tuples \bar{m} and \bar{n} , and consider $r_{\alpha+1}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$. Fix an $\varepsilon > 0$, and find $R \in \text{acor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ such that $\text{dis}_\Delta(R) < a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) + \frac{1}{2}\varepsilon$.

Now for any $a \in \mathfrak{M}$ find an a' such that there is some $b \in \mathfrak{N}$ with $(a', b) \in R$ and

such that a and a' are close enough that

$$|r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a; \mathfrak{N}, \bar{n}b) - r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a'; \mathfrak{N}, \bar{n}b)| \leq \frac{1}{2}\varepsilon.$$

(This exists since $r_\alpha^{\Delta, \Omega}$ is uniformly continuous in the tuple arguments.) By the induction hypothesis, $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a'; \mathfrak{N}, \bar{n}b) \leq a_\Delta(\mathfrak{M}, \bar{m}a'; \mathfrak{N}, \bar{n}b) \leq \text{dis}_\Delta(R)$, so we have that $r_\alpha^{\Delta, \Omega}(\mathfrak{M}, \bar{m}a; \mathfrak{N}, \bar{n}b) < a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) + \varepsilon$. Since we can do this for any $\varepsilon > 0$ and we can do the same thing for any $b \in \mathfrak{N}$, we have that $r_{\alpha+1}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq a_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.

(ii) From part (i) we already have that $r_\infty^{\Delta, \Omega} \leq \rho_\Delta$, so we just need to show that $\rho_\Delta \leq r_\infty^{\Delta, \Omega}$.

Fix $\varepsilon > 0$, and assume that $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) < \varepsilon$. As guaranteed by Proposition 6.3.4, let \bar{a}, \bar{b} be tail-dense sequences in \mathfrak{M} and \mathfrak{N} which begin with \bar{m} and \bar{n} , respectively, such that for every $n < \omega$, $r_0^{\Delta, \Omega}(\mathfrak{M}, a_{<n}; \mathfrak{N}, b_{<n}) < \varepsilon$. By the condition on Ω this implies that $R = \{(a_i, b_i) : i < \omega\}$ is an almost correlation between (\mathfrak{M}, \bar{m}) and (\mathfrak{N}, \bar{n}) such that $\text{dis}_\Delta(R) \leq \varepsilon$. Since Δ is regular this implies that $\rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq \varepsilon$. Since we can do this for any $\varepsilon > r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$, we have that $\rho_\Delta(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \leq r_\infty^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$.

(iii) From part (i) we already have that $r_\infty^{\Delta, \Omega} \leq \rho_\Delta$, so we just need to show that $\rho_\Delta \leq r_\infty^{\Delta, \Omega}$.

By Proposition 6.1.11 part (ii), there is a $\delta > 0$ such that for any $\gamma > 0$ there is a formula $\psi(x, y) \in \Delta$ such that for any $\mathfrak{M} \models T$ and $a, b \in \mathfrak{M}$, $\psi^{\mathfrak{M}}(a, a) = 0$, and if $\psi^{\mathfrak{M}}(a, b) < \delta$, then $d^{\mathfrak{M}}(a, b) < \gamma$. By the density of the sequence $\{\varphi_i\}_{i < \omega}$, there is an $m(\gamma)$ such that $\|\varphi_{m(\gamma)} - \psi\|_\infty < \frac{1}{3}\delta$ (in any \mathcal{L} -structure). This implies that for any $\mathfrak{M} \models T$ and $a, b \in \mathfrak{M}$, if $\varphi_{m(\gamma)}^{\mathfrak{M}}(a, a) < \frac{1}{3}\delta$ and $\varphi_{m(\gamma)}^{\mathfrak{M}}(a, b) < \frac{2}{3}\delta$, then $d^{\mathfrak{M}}(a, b) < \gamma$.

Now assume that $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) < \frac{1}{6}\delta$ and pick η such that $r_\infty^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) < \eta < \frac{1}{6}\delta$. As

guaranteed by Proposition 6.3.4, let \bar{a}, \bar{b} be tail-dense sequences in \mathfrak{M} and \mathfrak{N} , respectively, such that for every $n < \omega$, $r_0^{\Delta, \Omega}(\mathfrak{M}, a_{<n}; \mathfrak{N}, b_{<n}) < \eta$.

Let R be the set of all pairs $(c, e) \in \mathfrak{M} \times \mathfrak{N}$ such that there exists a sequence $\{i(j)\}_{j < \omega}$ of natural numbers such that $\lim_{j \rightarrow \infty} i(j) = \infty$ and $\{a_{i(j)}\}_{j < \omega}$ is a Cauchy sequence limiting to c and $\{b_{i(j)}\}_{j < \omega}$ is a Cauchy sequence limiting to e . By the uniform continuity of formulas and the fact that Ω is shift-increasing, it's clear that $\text{dis}_{\{\varphi_i\}}(R) \leq \eta$. Since $\{\varphi_i\}$ is dense in Δ this implies that $\text{dis}_{\Delta}(R) \leq \eta$.

So now we just need to show that R is a correlation. Pick $c \in \mathfrak{M}$, and let $\{i(j)\}_{j < \omega}$ be a sequence of natural numbers such that $\lim_{j \rightarrow \infty} i(j) = \infty$ and $\{a_{i(j)}\}_{j < \omega}$ is a Cauchy sequence limiting to c , which must exist by the tail-denseness of $\{a_i\}_{i < \omega}$. Consider the sequence $\{b_{i(j)}\}_{j < \omega}$.

Pick $\gamma > 0$, and consider the formula $\varphi_{m(\gamma)}$, as specified above. Find a $\sigma > 0$ such that if $d(xy, zw) < \sigma$, then $|\varphi_{m(\gamma)}(x, y) - \varphi_{m(\gamma)}(z, w)| < \frac{1}{6}\delta$ (in any \mathcal{L} -structure).

Find an $N(\gamma)$ such that $\varphi_{m(\gamma)}(x_{N(\gamma)}, x_{N(\gamma)+1})$ is an Ω -formula and such that for all $j, k \geq N(\gamma)$, $d^{\mathfrak{M}}(a_{i(j)}, a_{i(k)}) < \sigma$. This implies that for any $j, k \geq N(\gamma)$,

$$|\varphi_{m(\gamma)}^{\mathfrak{M}}(a_{i(j)}, a_{i(k)}) - \varphi_{m(\gamma)}^{\mathfrak{M}}(a_{i(j)}, a_{i(j)})| < \frac{1}{6}\delta,$$

so in particular

$$\varphi_{m(\gamma)}^{\mathfrak{M}}(a_{i(j)}, a_{i(k)}) < \frac{1}{3}\delta + \frac{1}{6}\delta = \frac{1}{2}\delta < \frac{2}{3}\delta.$$

This implies that for any $j, k \geq N(\gamma)$,

$$\varphi_{m(\gamma)}^{\mathfrak{N}}(b_{i(j)}, b_{i(k)}) < \frac{1}{2}\delta + \eta < \frac{1}{2}\delta + \frac{1}{6}\delta = \frac{2}{3}\delta.$$

By construction this implies that $d^{\mathfrak{N}}(b_{i(j)}, b_{i(k)}) < \gamma$. Since we can do this for any $\gamma > 0$, we have that $\{b_{i(j)}\}_{j < \omega}$ is a Cauchy sequence in \mathfrak{N} , converging to some point e , so we have that $(c, e) \in R$.

By symmetry we can do the same for Cauchy sequences in \mathfrak{N} , showing that R is a correlation, so we have that $\rho_{\Delta}(\mathfrak{M}, \mathfrak{N}) \leq r_{\infty}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N})$. Therefore $r_{\infty}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) = \rho_{\Delta}(\mathfrak{M}, \mathfrak{N})$ whenever $r_{\infty}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) < \varepsilon = \frac{1}{6}\delta$. \square

We should note that the case of u.u.c. distortion systems is very close to something that can be captured by the original formalism in [BDNT17]. In particular if Δ is a u.u.c. distortion system for a first-order theory T , then T is interdefinable with a theory T' in a uniformly Lipschitz language [Han] and the back-and-forth pseudo-metric coming from the 1-Lipschitz weak modulus will be uniformly equivalent to the original ρ_{Δ} for separable structures.

In cases where we know that a_{Δ} is not a pseudo-metric, we automatically know from part (i) that $r_{\infty}^{\Delta, \Omega} < a_{\Delta} < \rho_{\Delta}$, since $r_{\infty}^{\Delta, \Omega}$ and ρ_{Δ} are pseudo-metrics.

So now we can continue on to construct Scott sentences. This is the analog of Definition 3.6 in [BDNT17].

Definition 6.3.6. Let Δ be a collection of formulas closed under renaming variables and Ω a weak modulus. For a pair of models $\mathfrak{M}, \mathfrak{N} \models T$, $\alpha_{\mathfrak{M}, \mathfrak{N}}$ is the least ordinal α such that for all tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n} \in \mathfrak{N}$, $r_{\alpha}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) = r_{\alpha+1}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$. This is called the (Δ, Ω) -Scott rank of the pair \mathfrak{M} and \mathfrak{N} . If $\mathfrak{M} = \mathfrak{N}$ we just write $\alpha_{\mathfrak{M}}$, which is the (Δ, Ω) -Scott rank of \mathfrak{M} . \triangleleft

Just like Lemma 3.7 in [BDNT17] we have that if $r_{\infty}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) = 0$, then $\alpha_{\mathfrak{M}} = \alpha_{\mathfrak{M}, \mathfrak{N}} = \alpha_{\mathfrak{N}}$.

To define Scott sentences we need to specify what we mean by $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$.

Definition 6.3.7. Given a collection of first-order formulas Δ , closed under renaming variables, and a weak modulus Ω , n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formulas are defined inductively. We also need to inductively define the *codomain interval*, written $I(\varphi)$, of such formulas. For first-order formulas this is the convex closure of the set of possible values of the formula in \mathcal{L} -structures, which is always a compact interval.

- If $\varphi \in \Delta$ and $\varphi(x_0, \dots, x_{n-1})$ obeys Ω , then $\varphi(x_0, \dots, x_{n-1})$ is an n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula. $I(\varphi)$ is the codomain interval of φ as a first-order formula.
- For any compact interval I , if $\{\varphi_i\}_{i < \omega}$ is a sequence of n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formulas such that I is the closure of $\bigcup_{i < \omega} I(\varphi_i)$, then $\psi = \sup_i \varphi_i$ and $\chi = \inf_i \varphi_i$ are n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formulas and $I(\psi) = I(\chi) = I$.
- If φ is an $(n + 1)$ -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula then $\psi = \sup_{x_n} \varphi$ and $\chi = \inf_{x_n} \varphi$ are n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formulas and $I(\psi) = I(\chi) = I$.
- If $\varphi_1, \dots, \varphi_k$ is a finite list of n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formulas and $F : \mathbb{R}^k \rightarrow \mathbb{R}$ is a 1-Lipschitz connective, then $\psi = F(\varphi_1, \dots, \varphi_k)$ is an n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula and $I(\psi)$ is the image of $I(\varphi_1) \times \dots \times I(\varphi_k)$ under F (which is always a compact interval).

An $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula is an n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula for some n , and an $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -sentence is a 0-ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula. ◁

The interpretation of an $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula in a \mathcal{L} -structure is obvious. It's also not hard to show that an n -ary $\mathcal{L}_{\omega_1\omega}^{\Delta,\Omega}$ -formula φ is always $\Omega \upharpoonright n$ -uniformly continuous and can only take on values in the interval $I(\varphi)$ in \mathcal{L} -structures. Next, just like in [BDNT17] we get

that for every countable ordinal α , $n < \omega$, separable model $\mathfrak{M} \models T$, and tuple $\bar{m} \in \mathfrak{M}$, there is an $\mathcal{L}_{\omega_1\omega}^{\Delta, \Omega}$ -formula $\varphi_{\alpha, \mathfrak{M}, \bar{m}}$ such that for all models $\mathfrak{N} \models T$ and tuples $\bar{n} \in \mathfrak{N}$,

$$\varphi_{\alpha, \mathfrak{M}, \bar{m}}^{\mathfrak{M}}(\bar{n}) = r_{\alpha}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n}) \downarrow 1.$$

Fix a countable dense sub-pre-structure $\mathfrak{M}_0 = \{a_i\}_{i < \omega}$. We define these formulas inductively.

$$\varphi_{0, n, \mathfrak{M}, \bar{m}}(x_0, \dots, x_{n-1}) = \sup_i |\psi_i^{\mathfrak{M}}(\bar{a}) - \psi_i(\bar{x})|,$$

where $\{\psi_i\}_{i < \omega}$ is a countable sequence of n -ary Δ -formulas respecting Ω that are dense in the uniform norm in the collection of n -ary Δ -formulas respecting Ω .

For a successor stage we define $\varphi_{\alpha+1, n, \mathfrak{M}, \bar{m}}(x_0, \dots, x_{n-1})$ to be

$$\sup_i \inf_{x_n} \varphi_{\alpha, n+1, \mathfrak{M}, \bar{m}a_i}(x_0, \dots, x_n) \uparrow \sup_{x_n} \inf_i \varphi_{\alpha, n+1, \mathfrak{M}, \bar{m}a_i}(x_0, \dots, x_n).$$

And then for limit λ , obviously we define

$$\varphi_{\lambda, n, \mathfrak{M}, \bar{m}}(x_0, \dots, x_{n-1}) = \sup_{\alpha < \lambda} \varphi_{\alpha, n, \mathfrak{M}, \bar{m}}(x_0, \dots, x_{n-1}).$$

Now finally if $\alpha = \alpha_{\mathfrak{M}}$ is the (Δ, Ω) -Scott rank of \mathfrak{M} , then we define the (Δ, Ω) -Scott sentence, $\sigma_{\mathfrak{M}}$, by

$$\sigma_{\mathfrak{M}} = \varphi_{\alpha, 0, \mathfrak{M}} \uparrow \sup_{n < \omega, \bar{a} \in \mathfrak{M}_0^n} \sup_{x_0, \dots, x_{n-1}} \frac{1}{2} |\varphi_{\alpha, n, \mathfrak{M}, \bar{a}} - \varphi_{\alpha+1, n, \mathfrak{M}, \bar{a}}|,$$

i.e. $\mathfrak{N} \models \sigma_{\mathfrak{M}} \leq 0$ if and only if $\alpha_{\mathfrak{N}} = \alpha_{\mathfrak{M}}$ and $r_{\alpha}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) = 0$.

Now we get the following analog of Theorem 3.8 in [BDNT17].

Theorem 6.3.8. *If Δ is a u.u.c. or functional distortion system for T , then for any separable $\mathfrak{M}, \mathfrak{N} \models T$, $\mathfrak{N} \models \sigma_{\mathfrak{M}} \leq 0$ if and only if $\rho_{\Delta}(\mathfrak{M}, \mathfrak{N}) = 0$.*

Given the comment after Corollary 6.3.5, we know that Theorem 6.3.8 is simply false for irregular distortion systems (and we will examine an example in Section C.4.1), so two natural questions arise:

Question 6.3.9. *Do Corollary 6.3.5 and Theorem 6.3.8 hold for any regular distortion system with the appropriate choice of Ω ?*

Question 6.3.10. *For a fixed separable model $\mathfrak{M} \models T$, are either of the collections of separable models $\{\mathfrak{N} : \rho_{\Delta}(\mathfrak{M}, \mathfrak{N}) = 0\}$ or $\{\mathfrak{N} : a_{\Delta}(\mathfrak{M}, \mathfrak{N}) = 0\}$ Borel in the sense of [BDNT17] when Δ is an irregular distortion system?*

6.4 Parameters in Distortion Systems and d_{Δ}

Eventually we will need the correct analog of the d -metric for counting types in stability considerations and other things. This concept was introduced by Ben Yaacov in the slightly less general context of perturbations [BY08b]. Our δ_{Δ} is analogous to his \mathfrak{p} and our d_{Δ} is analogous to his $\mathfrak{p}_{\bar{a}}^0$.

Given a complete theory T and a collection of parameters A in some $\mathfrak{M} \models T$, T_A is the theory in the language \mathcal{L}_A with constants added for the parameters A . Given a distortion system Δ for T , there is a natural way to extend it to a distortion system $\Delta(A)$ for T_A .

Definition 6.4.1. Let Δ be a distortion system for a complete theory T . Let $A \subseteq \mathfrak{M} \models T$ be some set of parameters. $\Delta(A) = \{\varphi(\bar{x}, \bar{a}) : \varphi \in \Delta, \bar{a} \in A\}$. ◁

Clearly $\Delta(A)$ is still a distortion system and it's easy to see that for models $\mathfrak{M}, \mathfrak{N}$ containing A , we have $\rho_{\Delta(A)}(\mathfrak{M}, \mathfrak{N}) = \rho_{\Delta}(\mathfrak{M}, A; \mathfrak{N}, A)$.

Definition 6.4.2. For any distortion system Δ and set of fresh constant symbols C , let $D(\Delta, C)$ be $\overline{D_0(\Delta, C)}$, where $D_0(\Delta, C) = \Delta \cup \{d(x, c)\}_{c \in C}$. \triangleleft

Proposition 6.4.3. *If Δ is a distortion system, then for any set of fresh constant symbols C , $D(\Delta, C)$ is a distortion system.*

Proof. This follows immediately from Proposition 6.1.9. \square

The fact that $D(\Delta, C)$ is a distortion system isn't what is important about it, although it is convenient. What is important is that $\delta_{D(\Delta, \bar{c})}^0$ plays an analogous role to that of the d -metric in type spaces. In particular if Δ is the collection of all formulas and $p, q \in S_n(T)$ are two types, then $\delta_{D(\Delta, \bar{c})}^0(p(\bar{c}), q(\bar{c})) = d(p, q)$ (note that in this expression $p(\bar{c})$ and $q(\bar{c})$ are 0-types, i.e. complete $\mathcal{L}_{\bar{c}}$ -theories). To this end we will introduce notation to make the analogy more prominent.

Definition 6.4.4. If T is a complete theory, Δ is a distortion system for T , and A is some set of parameters in some model of T , then for any λ and $p, q \in S_{\lambda}(A)$, we let

$$d_{\Delta, A}(p, q) = \delta_{D(\Delta(A), \bar{c})}^0(p(\bar{c}, A), q(\bar{c}, A)).$$

We will drop A when it is empty. \triangleleft

Given how many layers there are to the definition of $d_{\Delta, A}$, the following will be useful for computing estimates of $d_{\Delta, A}$ and is really the best way to think about it.

Proposition 6.4.5. *Let T be a complete theory, let Δ be a distortion system for T .*

(i) For every $\varepsilon > 0$ there is a $\delta > 0$ such that if there are models $\mathfrak{M}, \mathfrak{N} \models T$ both containing some set of parameters A , tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n}, \bar{b} \in \mathfrak{N}$ such that $\bar{m} \models p$ and $\bar{n} \models q$, and an $R \in \text{cor}(\mathfrak{M}, A\bar{m}; \mathfrak{N}, A\bar{b})$, and $\text{dis}_\Delta(R), d^{\mathfrak{M}}(\bar{n}, \bar{b}) \leq \delta$, then $d_{\Delta, A}(p, q) \leq \varepsilon$.

(ii) For any set of parameters A , if $d_{\Delta, A}(p, q) \leq \varepsilon$, then there exists models $\mathfrak{M}, \mathfrak{N} \models T$ containing A , tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n}, \bar{b} \in \mathfrak{N}$ such that $\mathfrak{M} \models p(\bar{m})$ and $\mathfrak{N} \models q(\bar{n})$, and an $R \in \text{cor}(\mathfrak{M}, A\bar{m}; \mathfrak{N}, A\bar{b})$ such that $\text{dis}_\Delta(R) \leq \varepsilon$ and $d^{\mathfrak{N}}(\bar{n}, \bar{b}) \leq \varepsilon$.

Proof. (i) Fix $\varepsilon > 0$, By Lemma 6.1.12, there is a $\delta > 0$ such that if $\delta_\Delta(\text{tp}(ab), \text{tp}(ce)) \leq \delta$, then $|d(a, b) - d(c, e)| \leq \frac{1}{2}\varepsilon$. Without loss assume that $\delta < \frac{1}{2}\varepsilon$. Assume that there are models $\mathfrak{M}, \mathfrak{N} \models T_A$ with tuples $\bar{m} \in \mathfrak{M}$ and $\bar{n}, \bar{b} \in \mathfrak{N}$ such that $\bar{m} \models p$ and $\bar{n} \models q$, and an $R \in \text{cor}(\mathfrak{M}, A\bar{m}; \mathfrak{N}, A\bar{b})$, such that $\text{dis}_\Delta(R) \leq \delta$ and $d^{\mathfrak{N}}(\bar{n}, \bar{b}) \leq \delta$.

We need to compute $\text{dis}_{D(\Delta(A), \bar{c})}(R) = \text{dis}_{D_0(\Delta(A), \bar{c})}(R)$. Clearly we already have that $\text{dis}_{D_0(\Delta(A), \bar{c})}(R) \geq \text{dis}_{\Delta(A)}(R)$. We just need to compute $\sup_{c \in \bar{c}, (u, v) \in R} |d^{\mathfrak{M}}(u, c) - d^{\mathfrak{N}}(v, c)|$. For any $i < |\bar{m}|$, we have that $(m_i, b_i) \in R$. So for any $(u, v) \in R$, we have that $|d^{\mathfrak{M}}(m_i, u) - d^{\mathfrak{N}}(b_i, v)| \leq \frac{1}{2}\varepsilon$, so we also have $|d^{\mathfrak{M}}(m_i, u) - d^{\mathfrak{N}}(n_i, v)| \leq \frac{1}{2}\varepsilon + \delta < \varepsilon$. Therefore all together we have that $\text{dis}_{D(\Delta(A), \bar{c})}(R) = \text{dis}_{D_0(\Delta(A), \bar{c})}(R) \leq \varepsilon$, as required.

(ii) We have that since $d_\Delta(p, q) = \delta_{D(\Delta(A), \bar{c})}^0 \leq \varepsilon$, we can construct models $(\mathfrak{M}, \bar{m}) \models p(\bar{m})$ and $(\mathfrak{N}, \bar{n}) \models q(\bar{n})$ and an $R \in \text{cor}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$ such that $\text{dis}_{D(\Delta(A), \bar{c})}(R) \leq \varepsilon$. Find $\bar{b} \in \mathfrak{N}$ such that $(\bar{m}, \bar{b}) \in R$. Now we have that $|d^{\mathfrak{M}}(\bar{m}, \bar{m}) - d^{\mathfrak{N}}(\bar{b}, \bar{n})| \leq \varepsilon$, so in particular $d^{\mathfrak{N}}(\bar{n}, \bar{b}) \leq \varepsilon$, as required. \square

In particular $d_{\Delta, A}$ is always uniformly dominated by d on $S_n(A)$ and δ_Δ restricted to $S_n(A)$. Moreover as witnessed by the identity correlation on a sufficiently saturated model of T_A , $d_{\Delta, A} \leq d$ always holds.

Corollary 6.4.6. *If T is a complete theory and Δ is a u.u.c. distortion system for T , then for every $\varepsilon > 0$ there is a $\delta > 0$ such that for any λ and any types $p, q \in S_\lambda(T)$, if $d_\Delta(p, q) < \delta$ then $\delta_\Delta(p, q) < \varepsilon$, i.e. δ_Δ and d_Δ are uniformly equivalent.*

In fact it's not hard to see that δ_Δ and d_Δ being uniformly equivalent like this (in a way that is uniform across all parameter free type spaces) characterizes u.u.c. distortion systems. Furthermore this means that with u.u.c. Δ we don't need to be careful about the distinction between $(S_n(A), d_\Delta)$ and $(S_n(A), d_{\Delta, A})$, as these two metrics are always uniformly equivalent.

6.5 Approximately Atomic Types

The material in this section will only be important in Subsection 6.6.2 in the context of inseparable approximate categoricity.

Corollary 6.5.1. *Fix a complete first-order theory T , distortion system Δ for T , parameter set A , and a type $p \in S_n(A)$.*

- (i) *p is $\delta_{\Delta(A)}$ -atomic if and only if there is an A -formula φ such that $\varphi(p) = 0$ and for any formula $\psi(\bar{x}, \bar{y}) \in \Delta$ and $\bar{a} \in A$, $|\psi(q, \bar{a}) - \psi(p, \bar{a})| \downarrow 1 \leq \varphi(q)$ for all $q \in S_n(A)$.*
- (ii) *p is $d_{\Delta, A}$ -atomic if and only if there is an A -formula φ such that $\varphi(p) = 0$ and for any formula $\psi(\bar{c}, \bar{y}) \in D(\Delta(A), \bar{c})$ and $\bar{a} \in A$, $|\psi(q, \bar{a}) - \psi(p, \bar{a})| \downarrow 1 \leq \varphi(q)$ for all $q \in S_n(A)$.*

We will make use of the following fact [BY08c].

Fact 6.5.2. *If (X, d) is a topometric space then any continuous function $f : X \rightarrow \mathbb{R}$ is uniformly continuous with regards to d .*

Proposition 6.5.3. *Fix a complete first-order theory T , distortion system Δ for T , parameter set C , and tuples \bar{a} and \bar{b} with $|\bar{a}| = n$ and $|\bar{b}| = m$.*

(i) *$\text{tp}(\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic and $\text{tp}(\bar{a}/C\bar{b})$ is $\delta_{\Delta(C\bar{b})}$ -atomic if and only if $\text{tp}(\bar{a}\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic.*

(ii) *If $\text{tp}(\bar{b}/C)$ is $d_{\Delta, C}$ -atomic and $\text{tp}(\bar{a}/C\bar{b})$ is $d_{\Delta, C\bar{b}}$ -atomic, then $\text{tp}(\bar{a}\bar{b}/C)$ is $d_{\Delta, C}$ -atomic.*

(iii) *If $\text{tp}(\bar{a}\bar{b}/C)$ is $d_{\Delta, C}$ -atomic, then $\text{tp}(\bar{b}/C)$ is $d_{\Delta, C}$ -atomic.*

(iv) *If Δ is u.u.c. and $\text{tp}(\bar{a}\bar{b}/C)$ is $d_{\Delta, C}$ -atomic, then $\text{tp}(\bar{a}/C\bar{b})$ is $d_{\Delta, C\bar{b}}$ -atomic.*

(v) *If $\text{tp}(\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic or d -atomic, then it is $d_{\Delta, C}$ -atomic.*

Proof. (i, \Rightarrow) Let $\varphi(\bar{y})$ be a C -formula witnessing that $\text{tp}(\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic, and let $\psi(\bar{x}, \bar{y})$ be a C -formula such that $\psi(\bar{x}, \bar{b})$ witnesses that $\text{tp}(\bar{a}/C\bar{b})$ is $\delta_{\Delta(C\bar{b})}$ -atomic. By the fact above we have that $\psi(\bar{x}, \bar{y})$, a function on $S_{n+m}(C)$, is uniformly continuous in the metric $\delta_{\Delta(C)}$. Let $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ be a modulus of uniform continuity for $\psi(\bar{x}, \bar{y})$ with regards to $\delta_{\Delta(C)}$, i.e. α is a continuous function with $\alpha(0) = 0$ such that for any tuples $\bar{c}, \bar{e}, \bar{u}, \bar{v}$, $|\psi(\bar{c}, \bar{e}) - \psi(\bar{u}, \bar{v})| \leq \alpha(\delta_{\Delta(C)}(\text{tp}(\bar{c}\bar{e}/C), \text{tp}(\bar{u}\bar{v}/C)))$. We may assume that α is strictly increasing and in particular invertible.

Consider the formula $\chi(\bar{x}, \bar{y}) = \psi(\bar{x}, \bar{y}) + \alpha(\varphi(\bar{y}))$. Pick $\varepsilon > 0$ with $\varepsilon < 1$, and let \bar{c}, \bar{e} be such that $\chi(\bar{c}, \bar{e}) < \varepsilon$. This implies that $\varphi(\bar{e}) < \alpha^{-1}(\varepsilon)$, so $\delta_{\Delta(C)}(\text{tp}(\bar{e}/C), \text{tp}(\bar{b}/C)) < \alpha^{-1}(\varepsilon)$. Let $(\mathfrak{M}, \mathfrak{N}, R)$ be a structure witnessing this, i.e. $R \in \text{cor}(\mathfrak{M}, C\bar{e}; \mathfrak{N}, C\bar{b})$ and $\text{dis}_{\Delta}(R) < \alpha^{-1}(\varepsilon)$. Furthermore assume that $\bar{c} \in \mathfrak{M}$ as well, and let $\bar{a}' \in \mathfrak{N}$ be such that

$(\bar{c}, \bar{a}') \in R$. Now by construction we have that $\delta_{\Delta(C)}(\text{tp}(\bar{c}\bar{c}/C), \text{tp}(\bar{a}'\bar{b}/C)) < \alpha^{-1}(\varepsilon)$. This implies that $|\psi(\bar{c}, \bar{c}) - \psi(\bar{a}', \bar{b})| < \alpha(\alpha^{-1}(\varepsilon)) = \varepsilon$, so in particular $\psi(\bar{a}', \bar{b}) < 2\varepsilon$, and by construction $\delta_{\Delta(C)}(\text{tp}(\bar{a}', \bar{b}/C), \text{tp}(\bar{a}\bar{b}/C)) \downarrow 1 < 2\varepsilon$. By the triangle inequality this implies that $\delta_{\Delta(C)}(\text{tp}(\bar{c}\bar{c}/C), \text{tp}(\bar{a}\bar{b}/C)) < 2\varepsilon + \alpha^{-1}(\varepsilon)$. Since α is strictly increasing, α^{-1} is strictly increasing and the function $\varepsilon \mapsto 2\varepsilon + \alpha^{-1}(\varepsilon)$ is invertible. Let f be its inverse. We now have that $f(\psi(\bar{x}, \bar{y}) \downarrow 1)$ is a formula witnessing that $\text{tp}(\bar{a}\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic.

(i, \Leftarrow) Let $\varphi(\bar{x}, \bar{y})$ be a C -formula witnessing that $\text{tp}(\bar{a}\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic. Consider the formula $\psi(\bar{y}) = \inf_{\bar{x}} \varphi(\bar{x}, \bar{y})$. To see that this witnesses that $\text{tp}(\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic, let $\chi(\bar{y})$ be a $\Delta(C)$ -formula such that $\models \chi(\bar{b}) \leq 0$. Let \bar{b}' be a tuple such that $\models \psi(\bar{b}') \leq 0$. In a sufficiently saturated model there exists a tuple \bar{a}' such that $\models \varphi(\bar{a}'\bar{b}') \leq 0$, so we have that $\bar{b}' \equiv_C \bar{b}$. Now let \bar{c} be any tuple, and assume that $\psi(\bar{c}) < \varepsilon \leq 1$. That implies that there exists a tuple \bar{e} such that $\varphi(\bar{e}, \bar{c}) < \varepsilon$, so $\delta_{\Delta(C)}(\text{tp}(\bar{a}\bar{b}/C), \text{tp}(\bar{e}\bar{c}/C)) < \varepsilon$. This implies that $\delta_{\Delta(C)}(\text{tp}(\bar{b}/C), \text{tp}(\bar{c}/C)) < \varepsilon$ as well, so $\text{tp}(\bar{b}/C)$ is $\delta_{\Delta(C)}$ -atomic.

Finally, consider the formula $\varphi(\bar{x}, \bar{b})$. This witnesses that $\text{tp}(\bar{a}/C\bar{b})$ is $\delta_{\Delta(C\bar{b})}$ -atomic by the previous corollary and since $\Delta(C\bar{b})$ is the same set of formulas as $\{\varphi(\bar{x}, \bar{b}) : \varphi(\bar{x}, \bar{y}) \in \Delta(C)\}$.

(ii) We will prove a stronger version of this in Lemma 6.5.4.

(iii) The proof of this is the same as the proof of the corresponding statement in part (i).

(iv) If Δ is u.u.c. then $d_{\Delta, C}$ and $\delta_{\Delta(C)}$ are uniformly equivalent and $d_{\Delta, C\bar{b}}$ and $\delta_{\Delta(C\bar{b})}$ are uniformly equivalent, so the statement follows from part (i).

(v) This follows from the facts that $d_{\Delta, C} \leq \delta_{\Delta(C)}$ and $d_{\Delta, C} \leq d$ on $S_n(C)$. \square

Compare this proposition to these analogous non-approximate statements:

- (i) $\text{tp}(\bar{a}\bar{b}/C)$ is topologically isolated in $S_{n+m}(C)$ if and only if $\text{tp}(\bar{b}/C)$ is topologically isolated in $S_m(C)$ and $\text{tp}(\bar{a}/C\bar{b})$ is topologically isolated in $S_n(C\bar{b})$.
- (ii) If $\text{tp}(\bar{b}/C)$ and $\text{tp}(\bar{a}/C\bar{b})$ are d -atomic, then $\text{tp}(\bar{a}\bar{b}/C)$ is d -atomic.
- (iii) If $\text{tp}(\bar{a}\bar{b}/C)$ is d -atomic, then $\text{tp}(\bar{b}/C)$ is d -atomic.
- (v) If $\text{tp}(\bar{b}/C)$ is topologically isolated in $S_m(C)$, then it is d -atomic.

Topological isolation in $S_n(C)$ is of course the same thing as $\delta_{\Delta(C)}$ -atomicity when Δ is the collection of all formulas.

What fails in the statement ‘ $\text{tp}(\bar{a}\bar{b}/C)$ $d_{\Delta,C}$ -atomic implies $\text{tp}(\bar{b}/C\bar{a})$ is $d_{\Delta,C\bar{a}}$ -atomic’ is just the fact that $d_{\Delta,C\bar{a}}$ will in general be larger than $d_{\Delta,C}$ on $S_m(C\bar{a})$. This is the exact same phenomenon that happens with the ordinary d -metric.

A slightly beefier version of Proposition 6.5.3 part (ii) will be useful later for ‘approximately constructible’ models. We will prove this for a sequence of singletons for the sake of notational simplicity, but the proof is the same for a sequence of finite tuples.

Lemma 6.5.4. *Let T be a complete theory, Δ a distortion system for T , C a parameter set, \bar{a} a tuple of elements, and $\{b_i\}_{i < \omega}$ a sequence of elements. If $\text{tp}(b_{<i}/C)$ is $d_{\Delta,C}$ -atomic for every $i < \omega$ and $\text{tp}(\bar{a}/Cb_{<\omega})$ is $d_{\Delta,Cb_{<\omega}}$ -atomic, then $\text{tp}(\bar{a}b_{<i}/C)$ is $d_{\Delta,C}$ -atomic for every $i < \omega$.*

Proof. For each $i < \omega$, let $\varphi_i(y_0, \dots, y_{i-1})$ be a $[0, 1]$ -valued C -formula witnessing that $\text{tp}(b_{<i}/C)$ is $d_{\Delta,C}$ -atomic, and let $\psi(\bar{x}, \bar{y})$ be a C -formula such that $\psi(\bar{x}, b_{<\omega})$ witnesses that $\text{tp}(\bar{a}/Cb_{<\omega})$ is $d_{\Delta,Cb_{<\omega}}$ -atomic. Since $\psi(\bar{x}, \bar{y})$ can be realized as a uniformly convergent limit of formulas in finitely many variables, we have that $\psi(\bar{x}, \bar{y})$ is uniformly

continuous with regards to the metric

$$d^\dagger(\bar{x}\bar{y}, \bar{x}'\bar{y}') = \sup_{i < \omega} 2^{-i} d_{\Delta, C}(\text{tp}(\bar{x}y_{<i}/C), \text{tp}(\bar{x}'y'_{<i}/C)).$$

Let α be a modulus of uniform continuity for ψ with regards to d^\dagger . We may assume that α is strictly increasing and so in particular invertible. Consider the formulas

$$\chi_j(\bar{x}, y_0, \dots, y_{j-1}) = \inf_{y_j, y_{j+1}, \dots} \psi(\bar{x}, \bar{y}) + \sup_{i < \omega} 2^{-i} \alpha(\varphi_i(y_{<i})).$$

(Despite appearances these are first-order formulas.) For any $\bar{a}'b'_{<i}$ such that $\models \chi_i(\bar{a}'b'_{<i}) \leq 0$ there is an elementary extension with an ω -tuple $b'_i b'_{i+1} \dots$ such that $\models \psi(\bar{a}', b'_{<\omega}) \leq 0$ and $\models \varphi_i(b_{<i}) \leq 0$ for each $i < \omega$, so $\bar{a}'b'_{<i} \equiv_C \bar{a}b_{<i}$.

Let $\beta : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous strictly increasing function with $\beta(0) = 0$ that witnesses Proposition 6.4.5 part (i) in this situation, i.e. for every $\varepsilon > 0$ if there are models $\mathfrak{M}, \mathfrak{N} \models T$ both containing C , tuple $\bar{m} \in \mathfrak{M}$ and $\bar{n}, \bar{b} \in \mathfrak{N}$, and an $R \in \text{cor}(\mathfrak{M}, C\bar{m}; \mathfrak{N}, C\bar{n})$ with $\text{dis}_\Delta(R), d^{\mathfrak{M}}(\bar{n}, \bar{b}) \leq \beta(\varepsilon)$, then $d_{\Delta, C}(\text{tp}(\bar{m}/C), \text{tp}(\bar{n}/C)) \leq \varepsilon$.

Pick $\varepsilon > 0$ with $\varepsilon \leq 1$ and $j < \omega$, and let $\bar{c}e_{<j}$ be a tuple such that

$$\models \chi_j(\bar{c}, e_{<j}) < \frac{1}{2}\beta^{-1}\left(\frac{1}{2}\varepsilon\right) \downarrow 2^{-j}\beta^{-1}\left(\alpha^{-1}\left(\frac{1}{2}\beta^{-1}\left(\frac{1}{2}\varepsilon\right)\right)\right) \downarrow 2^{-j-1}\varepsilon.$$

Let $e_j e_{j+1} \dots$ be a tuple witnessing the infimum in $\chi_j(\bar{c}, e_{<j})$, so in particular

$$\models \psi(\bar{c}, e_{<\omega}) < \frac{1}{2}\beta^{-1}\left(\frac{1}{2}\varepsilon\right) \downarrow 2^{-j}\beta^{-1}\left(\alpha^{-1}\left(\frac{1}{2}\beta^{-1}\left(\frac{1}{2}\varepsilon\right)\right)\right) \downarrow 2^{-j-1}\varepsilon.$$

We also have that $\models 2^{-j}\varphi_j(e_{<j}) < 2^{-j}\beta^{-1}(\alpha^{-1}(\frac{1}{2}\beta^{-1}(\frac{1}{2}\varepsilon)) \downarrow 2^{-j-1}\varepsilon)$, so

$$d_{\Delta,C}(\text{tp}(e_{<k}/C), \text{tp}(b_{<j}/C)) < \beta^{-1} \left(\alpha^{-1} \left(\frac{1}{2}\beta^{-1} \left(\frac{1}{2}\varepsilon \right) \right) \downarrow 2^{-j-1}\varepsilon \right).$$

By Proposition 6.4.5 part (ii) we can find a triple $(\mathfrak{M}, \mathfrak{N}, R)$ such that $e_{<j} \in \mathfrak{M}$, $b_{<j} \in \mathfrak{N}$, $R \in \text{cor}(\mathfrak{M}, Ce_{<j}; \mathfrak{N}, C\bar{n})$, and there is a tuple $\bar{n} \in \mathfrak{N}$ such that $d^{\mathfrak{M}}(\bar{n}, b_{<j})$ and $\text{dis}_{\Delta}(R)$ are both less than $\beta^{-1}(\alpha^{-1}(\frac{1}{2}\beta^{-1}(\frac{1}{2}\varepsilon)) \downarrow 2^{-j-1}\varepsilon)$. We may assume that all of $b_{<\omega}$ is in \mathfrak{N} .

If we let $\bar{a}' \in \mathfrak{N}$ be a tuple such that $(\bar{e}, \bar{a}') \in R$, then we have that $d_{\Delta,C}(\text{tp}(\bar{c}e_{<k}/C), \text{tp}(\bar{a}'b_{<k}/C)) < \alpha^{-1}(\frac{1}{2}\beta^{-1}(\frac{1}{2}\varepsilon)) \downarrow 2^{-j-1}\varepsilon$, for each $k \leq j$, by Proposition 6.4.5 part (i). By construction this implies that

$$d^{\dagger}(\bar{c}e_{<\omega}, \bar{a}'b_{<\omega}) < \alpha^{-1} \left(\frac{1}{2}\beta^{-1} \left(\frac{1}{2}\varepsilon \right) \right) \downarrow 2^{-j-1}\varepsilon.$$

Now by the choice of α we get that $|\psi(\bar{c}, e_{<\omega}) - \psi(\bar{a}', b_{<\omega})| < \frac{1}{2}\beta^{-1}(\frac{1}{2}\varepsilon)$ and in particular $\psi(\bar{a}', b_{<\omega}) < \beta^{-1}(\frac{1}{2}\varepsilon)$, since $\psi(\bar{c}, e_{<\omega}) < \frac{1}{2}\beta^{-1}(\frac{1}{2}\varepsilon)$.

This implies that $d_{\Delta,Cb_{<\omega}}(\text{tp}(\bar{a}'/Cb_{<\omega}), \text{tp}(\bar{a}/Cb_{<\omega})) < \beta^{-1}(\frac{1}{2}\varepsilon)$. Proposition 6.4.5 now implies that $d_{\Delta,C}(\text{tp}(\bar{a}'b_{<j}/C), \text{tp}(\bar{a}b_{<j}/C)) < \frac{1}{2}\varepsilon$.

Since $d^{\dagger}(\bar{c}e_{<\omega}, \bar{a}'b_{<\omega}) < 2^{-j-1}\varepsilon$, we have by definition that

$$2^{-j}d_{\Delta,C}(\text{tp}(\bar{c}e_{<j}/C), \text{tp}(\bar{a}'b_{<j})) < 2^{-j-1}\varepsilon$$

and so $d_{\Delta,C}(\text{tp}(\bar{c}e_{<j}/C), \text{tp}(\bar{a}'b_{<j}/C)) < \frac{1}{2}\varepsilon$.

By the triangle inequality we have that $d_{\Delta,C}(\text{tp}(\bar{c}e_{<j}/C), \text{tp}(\bar{a}b_{<j}/C)) < \frac{1}{2}\varepsilon + \frac{1}{2}\varepsilon = \varepsilon$. Since we can do this for any $\varepsilon > 0$ and $j < \omega$, we have that $\text{tp}(\bar{a}b_{<j}/C)$ is $d_{\Delta,C}$ -atomic

for every $j < \omega$. □

6.6 Approximate Categoricity

Given an approximate notion of isomorphism it is natural to consider a corresponding notion of approximate categoricity.

Definition 6.6.1. Fix a theory T and a distortion system Δ for T .

- T is Δ - κ -categorical if for any two models $\mathfrak{M}, \mathfrak{N} \models T$ of density character κ , $a_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$.
- T is *strongly* Δ - κ -categorical if for any two models $\mathfrak{M}, \mathfrak{N} \models T$ of density character κ , $\mathfrak{M} \approx_\Delta \mathfrak{N}$. ◁

Of course Δ - κ -categoricity and strong Δ - κ -categoricity are equivalent when Δ is regular. The motivating examples of distortion systems are regular and certainly strong Δ - κ -categoricity is the more compelling notion, but all of the results in this section easily generalize to what we are calling Δ - κ -categoricity, with no apparent gain from assuming strong Δ - κ -categoricity and no clear way to characterize strong Δ - κ -categoricity. This, together with the necessity of introducing a weaker version of Δ - κ -categoricity in Section 6.6.2, lead us to this naming convention.

Of course a natural question is whether or not this distinction even matters.

Question 6.6.2. *Does there exist a theory T , a distortion system Δ for T , and an infinite cardinal κ such that T is Δ - κ -categorical but not strongly Δ - κ -categorical?*

6.6.1 Separable Categoricity

We can restate Ben Yaacov's generalization of the Engeler–Ryll–Nardzewski–Svenonius theorem to the context of perturbations in our language. In this section we will extend this result to distortion systems in general.

Theorem 6.6.3 (Ben Yaacov [BY08b]). *Let T be a complete theory with non-compact models in a countable language and Δ a functional distortion system for T .*

- *The following are equivalent.*
 - *T is Δ - ω -categorical.*
 - *For every finite tuple of parameters \bar{a} and $n < \omega$, every type in $S_n(\bar{a})$ is weakly d_Δ -atomic-in- $S_n(\bar{a})$.*
 - *The same but with n restricted to 1.*
- *If $(S_n(T), d_\Delta)$ is metrically compact for every $n < \omega$ (equivalently if every \emptyset -type is d_Δ -isolated), then T is Δ - ω -categorical.*

d_Δ on $S_1(\bar{a})$ is the restriction of d_Δ on $S_{1+|\bar{a}|}(T)$ to the subspace corresponding to $S_1(\bar{a})$. This is \tilde{d}_p in Ben Yaacov's notation. Note that d_Δ on $S_1(\bar{a})$ is not the same thing as $d_{\Delta, \bar{a}}$ on $S_1(\bar{a})$. For example, if Δ is the collection of all formulas, so that ρ_Δ corresponds to isomorphism, then $d_\Delta = d$ for the type spaces $S_n(T)$, but over parameters the topometric space $(S_n(\bar{a}), d)$, where d is the ordinary d -metric on $S_n(\bar{a})$, is not the same topometric space as $(S_n(\bar{a}), d_\Delta) = (S_n(\bar{a}), d^{S_{n+|\bar{a}|}(T)})$ in general, where $d^{S_{n+|\bar{a}|}(T)}$ is the d -metric on $S_{n+|\bar{a}|}(T)$.

Lemma 6.6.4. *Fix a countable first-order theory T and a distortion system Δ for T . For any type $p(\bar{x}) \in S_n(T)$ for some n , and any extension $q(\bar{x}, \bar{y}) \in S_{n+m}(T)$, there is*

a separable model $\mathfrak{M} \models T$ such that \mathfrak{M} realizes p and for any n -tuple $\bar{a} \in \mathfrak{M}$ and any $\varepsilon > 0$ there is an m -tuple $\bar{b} \in \mathfrak{M}$ such that $d_\Delta(q, \text{tp}(\bar{a}\bar{b})) < d_\Delta(p, \text{tp}(\bar{a})) + \varepsilon$.

Proof. Let \mathfrak{M}_0 be a countable pre-model of T realizing p . Proceed inductively. At stage i , given \mathfrak{M}_i , a countable pre-model of T , find $\mathfrak{M}_{i+1} \succeq \mathfrak{M}_i$ such that for every n -tuple $\bar{a} \in \mathfrak{M}_i$ there is an m -tuple $\bar{b} \in \mathfrak{M}_{i+1}$ such that $d_\Delta(q, \text{tp}(\bar{a}\bar{b})) < d_\Delta(p, \text{tp}(\bar{a})) + 2^{-i}$. Note that this is always possible since d_Δ has the extension property.

Now let \mathfrak{M} be the completion of the union $\bigcup_{i < \omega} \mathfrak{M}_i$. Let \bar{a} be an n -tuple in \mathfrak{M} . For any $\varepsilon > 0$, find $\bar{a}_i \in \mathfrak{M}_i$ such that $d(\bar{a}, \bar{a}_i) < \frac{1}{3}\varepsilon$ and such that $2^{-i} < \frac{1}{3}\varepsilon$. Then by construction there is $\bar{b} \in \mathfrak{M}$ such that $d_\Delta(q, \text{tp}(\bar{a}_i\bar{b})) < d_\Delta(p, \text{tp}(\bar{a}_i)) + \frac{1}{3}\varepsilon$. This implies that $d_\Delta(q, \text{tp}(\bar{a}_i\bar{b})) < d_\Delta(p, \text{tp}(\bar{a})) + \frac{2}{3}\varepsilon$, but we also have that $d_\Delta(q, \text{tp}(\bar{a}\bar{b})) < d_\Delta(q, \text{tp}(\bar{a}_i\bar{b})) + d(\bar{a}_i\bar{b}, \bar{a}\bar{b})$. Finally $d(\bar{a}_i\bar{b}, \bar{a}\bar{b}) = d(\bar{a}_i, \bar{a}) < \frac{1}{3}\varepsilon$, so putting this together gets $d_\Delta(q, \text{tp}(\bar{a}\bar{b})) < d_\Delta(q, \text{tp}(\bar{a}_i\bar{b})) + \varepsilon$, as required. \square

Definition 6.6.5. A structure \mathfrak{M} is *approximately Δ - ω -saturated* if for every $\bar{a} \in \mathfrak{M}$, every $p \in S_n(\bar{a})$, and every $\varepsilon > 0$, there is $\bar{b} \in \mathfrak{M}$ such that $d_\Delta(p, \text{tp}(\bar{a}\bar{b})) < \varepsilon$. \triangleleft

When Δ is the collection of all formulas, a structure is approximately Δ - ω -saturated if and only if it is approximately ω -saturated, hence the redundant sounding name.

Proposition 6.6.6. *A structure \mathfrak{M} is approximately Δ - ω -saturated if and only if it is for 1-types, i.e. for every $\bar{a} \in \mathfrak{M}$, every $p \in S_1(\bar{a})$, and every $\varepsilon > 0$, there is $b \in \mathfrak{M}$ such that $d_\Delta(p, \text{tp}(\bar{a}b)) < \varepsilon$.*

Proof. The \Rightarrow direction is obvious, so we only need to show that if \mathfrak{M} is Δ - ω -saturated for 1-types, then it is Δ - ω -saturated.

Let $\bar{a} \in \mathfrak{M}$ be a tuple, and let $p \in S_n(\bar{a})$ be some type. Pick $\varepsilon > 0$. For each i with $0 < i \leq n$, let p_i be the restriction of p to the first i variables.

First find $b_1 \in \mathfrak{M}$ such that $d_\Delta(p_1, \text{tp}(\bar{a}b_1)) < \frac{\varepsilon}{n}$.

Now at any stage $i \geq 1$, given $b_1 \dots b_i$, find $q_{i+1} \in S_1(\bar{a}b_1 \dots b_i)$ such that $d_\Delta(p_{i+1}, q_{i+1}) = d_\Delta(p_i, \text{tp}(\bar{a}b_1 \dots b_i))$. \square

Lemma 6.6.7. *If \mathfrak{M} is approximately Δ - ω -saturated, then for every tuple \bar{a} and every type $p(\bar{a}, \bar{y}) \in S_n(\bar{a})$ and $\varepsilon > 0$ there is $\bar{b}\bar{c} \in \mathfrak{M}$ such that $d(\bar{a}, \bar{b}) < \varepsilon$ and $\delta_\Delta(p, \text{tp}(\bar{b}\bar{c})) < \varepsilon$.*

Proof. Pick $\varepsilon > 0$. Let $\bar{b}_0 = \bar{a}$, and find \bar{c}_0 such that $d_\Delta(p, \text{tp}(\bar{b}_0\bar{c}_0)) < \frac{1}{2}\varepsilon$. By Proposition 6.4.5 there exists a type q_0 such that $\delta_\Delta(p, q_0) < \frac{1}{2}\varepsilon$ and $d(q_0, \text{tp}(\bar{b}_0\bar{c}_0)) < \frac{1}{2}\varepsilon$. Let r_0 be a completion of the type $\{d(\bar{b}_0\bar{c}_0, \bar{x}\bar{y}) < \frac{1}{2}\varepsilon, q_0(\bar{x}, \bar{y})\}$ (i.e. a type in twice as many variables).

Now at stage i , given $\bar{b}_i\bar{c}_i$ and r_i , find $\bar{b}_{i+1}\bar{c}_{i+1}$ such that $d_\Delta(r_i, \text{tp}(\bar{b}_i\bar{c}_i\bar{b}_{i+1}\bar{c}_{i+1})) < 2^{-i-1}\varepsilon$. So in particular $d_\Delta(q_i, \text{tp}(\bar{b}_{i+1}\bar{c}_{i+1})) < 2^{-i-1}\varepsilon$. By Proposition 6.4.5 there exists a type q_{i+1} such that $\delta_\Delta(q_i, q_{i+1}) < 2^{-i-1}\varepsilon$ and $d(q_{i+1}, \text{tp}(\bar{b}_{i+1}\bar{c}_{i+1})) < 2^{-i-1}\varepsilon$. Let r_{i+1} be a completion of the type $\{d(\bar{b}_{i+1}\bar{c}_{i+1}, \bar{x}\bar{y}) < 2^{-i-1}\varepsilon, q_{i+1}(\bar{x}, \bar{y})\}$.

Now by construction $\{\bar{b}_i\bar{c}_i\}_{i < \omega}$ is a Cauchy sequence in d . Let it limit to $\bar{b}\bar{c}$. By construction we have that $d(\bar{a}, \bar{b}) < \varepsilon$. $\{q_i\}_{i < \omega}$ is also a Cauchy sequence in δ_Δ . Let it limit to q . By construction we have that $\delta_\Delta(p, q) < \varepsilon$. Furthermore since $d(q_{i+1}, \text{tp}(\bar{b}_{i+1}\bar{c}_{i+1})) < 2^{-i-1}\varepsilon$, we have that $\bar{b}\bar{c} \models q$, so in particular $\delta_\Delta(p, \text{tp}(\bar{b}\bar{c})) < \varepsilon$, as required. \square

Proposition 6.6.8. *For any countable theory T (not necessarily complete) and distortion system Δ for T , if $\mathfrak{M}, \mathfrak{N} \models T$ are separable models that are both approximately Δ - ω -saturated, then $a_\Delta(\mathfrak{M}, \mathfrak{N})$ is as small as possible, i.e. $a_\Delta(\mathfrak{M}, \mathfrak{N}) = \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$. In particular if $\mathfrak{M} \equiv \mathfrak{N}$ then $a_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$.*

Proof. Pick $\varepsilon > \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$. Let $\{m_{2i}\}_{i < \omega}$ be an enumeration of a tail-dense (i.e. every final segment is dense) sequence in \mathfrak{M} , and let $\{n_{2i+1}\}_{i < \omega}$ be a tail-dense sequence

in \mathfrak{N} . We will construct an almost correlation between \mathfrak{M} and \mathfrak{N} with a back-and-forth argument with the typical continuous modification that the sequence built will need to ‘slide around’ a little at each stage to make things line up.

$\{a_i^j\}_{j \leq i < \omega}$ will be an array of elements of \mathfrak{M} and $\{b_i^j\}_{j \leq i < \omega}$ will be an array of elements of \mathfrak{N} chosen so that for each fixed i , a_i^j and b_i^j are Cauchy sequences in j . Their limits, a_i^ω and b_i^ω , will be the desired correlation with distortion $\leq \varepsilon$.

On stage 0, let $a_0^0 = m_0$, and find b_0^0 such that

$$\delta_\Delta(\text{tp}(a_0^0), \text{tp}(b_0^0)) = \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N})) < \varepsilon.$$

On odd stage $2k+1$, let $b_i^{2k+1} = b_i^{2k}$ for $i < 2k+1$, and let $b_{2k+1}^{2k+1} = n_{2k+1}$. By the induction hypothesis, $\delta_\Delta(\text{tp}(a_{\leq 2k}^{2k}), \text{tp}(b_{\leq 2k}^{2k})) < \varepsilon$. Let p_{2k+1} be an extension of the type $\text{tp}(a_{\leq 2k}^{2k})$ such that $\delta_\Delta(\text{tp}(a_{\leq 2k}^{2k}), \text{tp}(b_{\leq 2k}^{2k})) = \delta_\Delta(p_{2k+1}, \text{tp}(b_{\leq 2k+1}^{2k+1}))$. Now by Lemma 6.6.7 we can find $a_{\leq 2k}^{2k+1}$ and a_{2k+1}^{2k+1} such that $d^{\mathfrak{M}}(a_{\leq 2k}^{2k}, a_{\leq 2k}^{2k+1}) < 2^{-k}$ and with $\delta_\Delta(p_{2k+1}, \text{tp}(a_{\leq 2k+1}^{2k+1}))$ small enough that $\delta_\Delta(\text{tp}(a_{\leq 2k+1}^{2k+1}), \text{tp}(b_{\leq 2k+1}^{2k+1})) < \varepsilon$.

On even stage $2k+2$ we do the same with the roles reversed.

Now clearly for each fixed i , a_i^j and b_i^j are Cauchy sequences in j , so let a_i^ω and b_i^ω be their limits. Note that $d^{\mathfrak{M}}(a_{2k}^\omega, m_{2k}) \leq 2^{-2k+1}$ and similarly for b_{2k+1}^ω and n_{2k+1} , so we have that $\{a_i^\omega\}_{i < \omega}$ is dense in \mathfrak{M} and $\{b_i^\omega\}_{i < \omega}$ is dense in \mathfrak{N} by the tail-density of the sequences $\{m_{2i}\}$ and $\{n_{2i+1}\}$. So $R = \{(a_i^\omega, b_i^\omega) : i < \omega\}$ is an almost correlation between \mathfrak{M} and \mathfrak{N} .

By induction we have for each $j < \omega$ that $\delta_\Delta(\text{tp}(a_{\leq j}^k), \text{tp}(b_{\leq j}^k)) < \varepsilon$ for all k such that this quantity is defined. By lower semi-continuity of δ_Δ this implies that $\delta_\Delta(\text{tp}(a_{\leq j}^\omega), \text{tp}(b_{\leq j}^\omega)) \leq \varepsilon$ for all $j < \omega$. Therefore we have that $\text{dis}_\Delta(R) \leq \varepsilon$ as well.

Since we can do this for any $\varepsilon > \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$, we have that $a_\Delta(\mathfrak{M}, \mathfrak{N}) \leq \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$. Since $a_\Delta(\mathfrak{M}, \mathfrak{N}) \geq \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$ always holds, we have $a_\Delta(\mathfrak{M}, \mathfrak{N}) = \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$, as required. \square

Corollary 6.6.9. *If Δ is a regular distortion system and $\mathfrak{M}, \mathfrak{N} \models T$ are approximately Δ - ω -saturated separable models then $\rho_\Delta(\mathfrak{M}, \mathfrak{N}) = \delta_\Delta(\text{Th}(\mathfrak{M}), \text{Th}(\mathfrak{N}))$. In particular if $\mathfrak{M} \equiv \mathfrak{N}$, then $\mathfrak{M} \approx_\Delta \mathfrak{N}$.*

Proposition 6.6.10. *For any countable complete theory T and distortion system Δ for T , if T is Δ - ω -categorical, then every separable model of T is approximately Δ - ω -saturated.*

Proof. Fix a separable $\mathfrak{M} \models T$. Fix $\bar{a} \in \mathfrak{M}$ and $p(\bar{x}, \bar{a}) \in S_m(\bar{a})$. Let \mathfrak{N} be the model for $\text{tp}(\bar{a})$ and p guaranteed by the previous lemma. Pick $\gamma > 0$. Find $\delta > 0$ according to Proposition 6.4.5 (ii) with $\varepsilon < \frac{1}{4}\gamma$. Without loss assume that $\delta < \frac{1}{4}\gamma$. Let $R \in \text{acor}(\mathfrak{M}, \mathfrak{N})$ be closed and have $\text{dis}_\Delta(R) < \delta$. Find \bar{a}' in the domain of R such that $d^{\mathfrak{M}}(\bar{a}, \bar{a}') < \delta$, and let $\bar{b} \in \mathfrak{N}$ be such that $(\bar{a}', \bar{b}) \in R$. By passing to an \aleph_1 -saturated elementary extension of $(\mathfrak{M}, \mathfrak{N}, R)$ and by using Proposition 6.4.5 we get that $d_\Delta(\text{tp}(\bar{a}), \text{tp}(\bar{b})) \leq \varepsilon < \frac{1}{4}\gamma$.

This implies that there exists $\bar{c} \in \mathfrak{N}$ such that $d_\Delta(p, \text{tp}(\bar{b}\bar{c})) < d_\Delta(\text{tp}(\bar{a}), \text{tp}(\bar{b})) + \frac{1}{4}\gamma$. Now find $\bar{c}' \in \mathfrak{N}$ in the range of R such that $d^{\mathfrak{N}}(\bar{c}, \bar{c}') < \frac{1}{4}\gamma$. Now we have that $d_\Delta(p, \text{tp}(\bar{b}\bar{c}')) < d_\Delta(p, \text{tp}(\bar{b}\bar{c})) + d(\bar{c}, \bar{c}')$ and thus $d_\Delta(p, \text{tp}(\bar{b}\bar{c}')) < d_\Delta(p, \text{tp}(\bar{b}\bar{c})) + \frac{1}{4}\gamma$. Putting this all together gives that $d_\Delta(p, \text{tp}(\bar{b}\bar{c}')) < \frac{3}{4}\gamma$.

Now find $\bar{e} \in \mathfrak{M}$ such that $(\bar{e}, \bar{c}') \in R$, so in particular we have that $(\bar{a}'\bar{e}, \bar{b}\bar{c}') \in R$. By passing to an \aleph_1 -saturated elementary extension of $(\mathfrak{M}, \mathfrak{N}, R)$ and using Proposition 6.4.5 again, we have that $d_\Delta(\text{tp}(\bar{a}'\bar{e}), \text{tp}(\bar{b}\bar{c}')) \leq \varepsilon < \frac{1}{4}\gamma$. By the triangle inequality this

gives that $d_\Delta(p, \text{tp}(\bar{a}\bar{e})) < \frac{1}{4}\gamma + \frac{3}{4}\gamma = \gamma$.

Since we can do this for any separable $\mathfrak{M} \models T$, any \bar{a} , any $p \in S_m(\bar{a})$, and any $\gamma > 0$, we have that every separable models of T is approximately Δ - ω -saturated. \square

Corollary 6.6.11. *A countable theory T with distortion system Δ is Δ - ω -categorical if and only if every separable model is approximately Δ - ω -saturated.*

Proposition 6.6.12. *Fix a countable complete theory T and a distortion system Δ for T . For any parameters \bar{a} and a type $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$ the following are equivalent:*

- *For every $\mathfrak{M} \models T$ containing $\bar{b} \equiv \bar{a}$ and for every $\varepsilon > 0$, there is $\bar{c} \in \mathfrak{M}$ such that $d_\Delta(p(\bar{x}, \bar{b}), \text{tp}(\bar{c}\bar{b})) < \varepsilon$.*
- *p is weakly d_Δ -atomic-in- $S_n(\bar{a})$.*

Proof. The second bullet point clearly implies the first bullet, since if p is weakly d_Δ -atomic-in- $S_n(\bar{a})$, then for every $\varepsilon > 0$, $\text{int}_{S_n(\bar{a})} B_{\leq \varepsilon}^{d_\Delta}(p)$ is non-empty and non-empty open subsets of type space are always realized.

So assume that the second bullet point fails. This implies that there is an $\varepsilon > 0$ such that $\text{int}_{S_n(\bar{a})} B_{\leq \varepsilon}^{d_\Delta}(p) = \emptyset$. $\text{int}_{S_n(\bar{a})} B_{\leq \varepsilon}^{d_\Delta}(p)$ is a closed set, so we can build a pre-model $\mathfrak{M}_0 \ni \bar{a}$ omitting it. Passing to the completion $\mathfrak{M} = \overline{\mathfrak{M}_0}$, note that any type realized in \mathfrak{M} must be in the metric closure under the ordinary d -metric of the set of types realized in \mathfrak{M}_0 , but since $d \geq d_\Delta$ this implies that they're in the metric closure under d_Δ of the set of types realized in \mathfrak{M}_0 , so we have that any type $q \in S_n(\bar{a})$ realized in \mathfrak{M} must have $d_\Delta(p, q) \geq \varepsilon$, contradicting the first bullet point. \square

Theorem 6.6.13. *For any countable complete theory T and distortion system Δ for T ,*

(i) T is Δ - ω -categorical if and only if for every tuple of parameters \bar{a} and every $n < \omega$, every $p \in S_n(\bar{a})$ is weakly d_Δ -atomic-in- $S_n(\bar{a})$.

(ii) Same as the previous statement but only considering types in $S_1(\bar{a})$.

(iii) If every $S_n(T)$ is metrically compact relative to d_Δ , then T is Δ - ω -categorical.

Proof. (i) If T is Δ - ω -categorical then for any finite tuple \bar{a} , any type $p(\bar{x}, \bar{a})$, and any separable model $\mathfrak{M} \ni \bar{a}$, the condition in the first bullet point of Proposition 6.6.12 holds (since \mathfrak{M} is approximately Δ - ω -saturated), i.e. the type must be ‘approximately realized.’ Therefore p must be weakly d_Δ -atomic-in- $S_n(\bar{a})$.

Conversely, if for every \bar{a} and $p \in S_n(\bar{a})$, p is weakly d_Δ -atomic-in- $S_n(\bar{a})$, then by Proposition 6.6.12 no type over a tuple of parameters can be ‘approximately omitted,’ i.e. the first bullet point in Proposition 6.6.12 always holds. This implies that every separable model of T is approximately Δ - ω -saturated, therefore T is Δ - ω -categorical.

(ii) Clearly if every $p \in S_n(\bar{a})$ is weakly d_Δ -atomic-in- $S_n(\bar{a})$, then the same is true restricting n to 1.

So assume that for every finite tuple of parameters \bar{a} , every $p \in S_1(\bar{a})$ is weakly d_Δ -atomic-in- $S_1(\bar{a})$. It’s clear that the proof of Proposition 6.6.8 only requires approximate Δ - ω -saturation for 1-types, so we get that for any two separable models $\mathfrak{M}, \mathfrak{N} \models t$, $a_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$, i.e. T is Δ - ω -categorical.

(iii) This is enough to imply that every $S_n(\bar{a})$ is metrically compact with regards to d_Δ as well (since d_Δ is just the restriction to $S_n(\bar{a})$ as a subspace of $S_{n+|\bar{a}|}(T)$). Therefore every $p \in S_n(\bar{a})$ is d_Δ -atomic, and therefore in particular weakly d_Δ -atomic. \square

Corollary 6.6.14. *If T is a countable theory, Δ is a distortion system for T , \bar{a} is any tuple of parameters, and T is Δ - ω -categorical, then $T_{\bar{a}}$ is $D(\Delta, \bar{a})$ - ω -categorical, i.e. for*

any two separable models $\mathfrak{M}, \mathfrak{N} \models T_{\bar{a}}$ and $\varepsilon > 0$ there is an almost correlation R between \mathfrak{M} and \mathfrak{N} such that $\text{dis}_{\Delta}(R) < \varepsilon$ and for all $(\bar{m}, \bar{n}) \in R$, $|d^{\mathfrak{M}}(\bar{m}, \bar{a}^{\mathfrak{M}}) - d^{\mathfrak{N}}(\bar{n}, \bar{a}^{\mathfrak{N}})| < \varepsilon$.

Proof. Unpacking definitions gives that $d_{D(\Delta, \bar{a})}$ in $S_n(\bar{a}\bar{b})$ is the same thing as d_{Δ} in $S_n(\bar{a}\bar{b})$. So we clearly have that $T_{\bar{a}}$ is $D(\Delta, \bar{a})$ - ω -categorical. All we need to do is verify that the alternative statement of $D(\Delta, \bar{a})$ - ω -categoricity is equivalent, but this follows from the definition of $D_0(\Delta, \bar{a})$. \square

Note that in particular if Δ is the collection of all formulas, so that Δ - ω -categoricity corresponds to ω -categoricity, then $D(\Delta, \bar{a})$ - ω -categoricity is not the same thing as ω -categoricity for $T_{\bar{a}}$. Rather, it says that given any two separable models $\mathfrak{M}, \mathfrak{N} \models T_{\bar{a}}$, for any $\varepsilon > 0$ there is an isomorphism $f : \mathfrak{M} \cong \mathfrak{N}$ such that $d(f(\bar{a}^{\mathfrak{M}}), \bar{a}^{\mathfrak{N}}) < \varepsilon$. Similar weakenings occur when Δ is a distortion system with an ‘obvious’ extension to $T_{\bar{a}}$, although in some cases, such as with the Gromov-Hausdorff distance, the obvious extension is equivalent to $D(\Delta, \bar{a})$ (in particular because δ_{Δ} and d_{Δ} are uniformly equivalent).

Recall the definition of the ‘elementary Gromov-Hausdorff distance’ (Definition 6.2.3).

Corollary 6.6.15. *If a countable theory T is eGHK- ω -categorical, then it is ω -categorical.*

Proof. $\delta_{\text{eGHK}} = d$ and for any type $p \in S_n(T)$ (it is important that there are no parameters), p is weakly d -atomic if and only if it is d -atomic. Every type in $S_n(T)$ being d -atomic is equivalent to $S_n(T)$ being metrically compact, so it follows that T is ω -categorical. \square

It may seem that there is a contradiction here, given the knowledge that ordinary ω -categoricity isn’t always preserved under adding constants. This stems from the fact that given a theory T and the corresponding $\text{eGHK}(T)$, $\text{eGHK}(T)(\bar{a})$ is not

the same thing as $\text{eGHK}(T_{\bar{a}})$ and the difference is related to what happens with ω -categorical theories that fail to be ω -categorical after adding constants. Witnesses for $\rho_{\text{eGHK}(T)(\bar{a})}(\mathfrak{M}, \mathfrak{N}) < \varepsilon$ for $\mathfrak{M}, \mathfrak{N} \models T_{\bar{a}}$ are elementary embeddings $f : \mathfrak{M} \preceq \mathfrak{C}$ and $g : \mathfrak{N} \preceq \mathfrak{C}$ such that $d_H^{\mathfrak{C}}(f(\mathfrak{M}), g(\mathfrak{N})) < \varepsilon$ and $d^{\mathfrak{C}}(f(\bar{a}), g(\bar{a})) < \varepsilon$. On the other hand witnesses for $\rho_{\text{eGHK}(T_{\bar{a}})}(\mathfrak{M}_{\bar{a}}, \mathfrak{N}_{\bar{a}}) < \varepsilon$ require that $f(\bar{a}) = g(\bar{a})$ because in this case we're thinking of the big model \mathfrak{C} as a model of $T_{\bar{a}}$ and we need f and g to be elementary embeddings for $T_{\bar{a}}$, not just T .

We were unable to show the analogous result for uncountable cardinalities but we were also unable to construct a counterexample, so a natural question arises.

Question 6.6.16. *Does there exist a countable theory T and an uncountable cardinality κ such that T is $\text{eGHK-}\kappa$ -categorical but not κ -categorical?*

This question is obviously trivial in single-sorted discrete theories, but there is a specific case of it with a many-sorted discrete theory that is non-trivial and can be resolved negatively (i.e. $\text{eGHK-}\kappa$ -categoricity implies κ -categoricity). If we have a many-sorted discrete theory with sorts $\{\mathcal{S}_i\}_{i < \omega}$ and we take the metric on sort \mathcal{S}_i to be $\{0, 2^{-i}\}$ -valued, then the question is non-trivial. This is equivalent to taking a Morleyized many-sorted language \mathcal{L} and letting \mathcal{L}_i be the set of all formulas with free variables in the first i sorts (but, crucially, we're implicitly allowing quantification over arbitrarily high sorts). Assume that such a theory is $\text{eGHK-}\kappa$ -categorical for some uncountable κ . By Corollary 6.6.34 in the next section, this implies that it is $\text{GHK-}\lambda$ -categorical for every uncountable λ (since this is a discrete theory). Therefore it cannot have any Vaughtian pairs. Now to show that it is actually uncountably categorical we just need to show that it is ω -stable. What we have is that it is $\text{eGHK-}\omega$ -stable (as defined in the next section).

Unpacking the definition in this case, this implies that it cannot have an infinite binary tree whose parameters all come from a finite collection of sorts (whereas in principle a binary tree in general could have parameters from all of the sorts). However, by stability and in particular the fact that the sorts are stably embedded, all of the formulas in the binary tree are definable with parameters from the sort that the tree is in, so we have that in this case eGHK- ω -stability implies ω -stability. Therefore the theory is actually uncountably categorical. By the results of [Han20], this is enough to resolve Question 6.6.16 in the context of continuous theories with totally disconnected type spaces (which include ω -stable ultrametric theories) since these are bi-interpretable with many-sorted discrete theories.

6.6.2 Inseparable Approximate Categoricity

Definition 6.6.17. For any topological space X and metric $d : X^2 \rightarrow \mathbb{R}$, a (d, ε) -perfect tree in X is a family of non-empty closed sets $\{F_\sigma\}_{\sigma \in 2^{<\omega}}$ such that for any $\sigma \in 2^{<\omega}$ and $i < 2$, $F_{\sigma \frown i} \subset \text{int}_X F_\sigma$ and $d_{\text{inf}}(F_{\sigma \frown 0}, F_{\sigma \frown 1}) > \varepsilon$. \triangleleft

Definition 6.6.18. Fix a complete theory T and a distortion system Δ for T .

- T is Δ - κ -stable if for any parameter set A of size $\leq \kappa$, $\#^{\text{dc}}(S_1(A), d_{\Delta, A}) \leq \kappa$.
- T is Δ -totally transcendental or Δ -t.t. if for any parameter set A there does not exist a $(d_{\Delta, A}, \varepsilon)$ -perfect tree in $S_1(A)$. \triangleleft

Proposition 6.6.19. Fix a theory T and a distortion system Δ for T .

- (i) If T is Δ - ω -stable, then it is Δ -t.t.
- (ii) If T is Δ -t.t., then it is Δ - κ -stable for any $\kappa \geq |\mathcal{L}|$.

(iii) If T is countable and Δ - κ -stable for some $\kappa = \kappa^\omega$, then it is stable. In particular a Δ - ω -stable countable theory is stable.

Proof. (i) Assume that for some set of parameters A and some $\varepsilon > 0$ there is a $(d_{\Delta,A}, \varepsilon)$ -perfect tree $\{F_\sigma\}_{\sigma \in 2^{<\omega}}$ in $S_1(A)$.

Fix $\sigma \in 2^{<\omega}$. Fix $p \in F_{\sigma \smallfrown 0}$ and $q \in F_{\sigma \smallfrown 1}$. We have by assumption that $d_{\Delta,A}(p, q) > \varepsilon$. This implies that there exists a $D(\Delta(A), \bar{c})$ -sentence, $\varphi_{\sigma,p,q}$, such that $|\varphi_{\sigma,p,q}(p) - \varphi_{\sigma,p,q}(q)| > \varepsilon$. There are open sets $U \ni p$ and $V \ni q$ such that for all $r \in U$ and $s \in V$, $|\varphi_{\sigma,p,q}(r) - \varphi_{\sigma,p,q}(s)| > \varepsilon$. Therefore by compactness there is a finite collection Σ_σ of $D(\Delta(A), \bar{c})$ -sentences such that for any $p \in F_{\sigma \smallfrown 0}$ and $q \in F_{\sigma \smallfrown 1}$ there is a $\varphi \in \Sigma_\sigma$ such that $|\varphi(p) - \varphi(q)| > \varepsilon$.

For each $\sigma \in 2^{<\omega}$, let ψ_σ be a restricted A -formula such that $\llbracket \psi_\sigma \leq \frac{2}{3} \rrbracket \subseteq \text{int}_{S_1(A)} F_\sigma$ and such that $F_{\sigma \smallfrown 0}, F_{\sigma \smallfrown 1} \subseteq \llbracket \psi_\sigma < \frac{1}{3} \rrbracket$. Now let A_0 be the collection of all parameters used in some $\varphi \in \Sigma_\sigma$ or ψ_σ for some $\sigma \in 2^{<\omega}$. Note that A_0 is a countable set of parameters. For each $\sigma \in 2^{<\omega}$, let $G_\sigma = \llbracket \psi_\sigma \leq \frac{2}{3} \rrbracket$, and note that by construction for any $\sigma \in 2^{<\omega}$ and $i < 2$, $G_{\sigma \smallfrown i} \subseteq \text{int}_{S_1(A_0)} G_\sigma$.

To verify that $\{G_\sigma\}_{\sigma \in 2^{<\omega}}$ is a $(d_{\Delta,A_0}, \varepsilon)$ -perfect tree in $S_1(A_0)$ we now just need to verify that $d_{\Delta,A_0}(p, q) > \varepsilon$ for any $p \in G_{\sigma \smallfrown 0}$ and $q \in G_{\sigma \smallfrown 1}$, but this follows easily from the inclusion of the parameters from Σ_σ . For any such p, q there is some sentence $\varphi \in \Sigma_\sigma$ such that $|\varphi(p) - \varphi(q)| > \varepsilon$. So we have that $\{G_\sigma\}$ is a $(d_{\Delta,A_0}, \varepsilon)$ -perfect tree. Therefore $\#^{\text{dc}}(S_1(A_0), d_{\Delta,A_0}) \geq 2^{\aleph_0}$ and T is not Δ - ω -stable.

(ii) Suppose that T fails to be Δ - κ -stable for some $\kappa \geq |\mathcal{L}|$. Let A be a collection of parameters of cardinality $\leq \kappa$ such that $\#^{\text{dc}}(S_1(A), d_{\Delta,A}) > \kappa$. This implies that there is some $\varepsilon > 0$ such that $\text{ent}_{>\varepsilon}(S_1(A), d_{\Delta,A}) \geq \kappa^+$.

Let \mathcal{B} be a base for the topology of $S_1(A)$ of cardinality κ (this exists because $\kappa \geq |\mathcal{L}|$).

Define a transfinite sequence of closed subsets:

- $X_0 = S_1(A)$,
- $X_{\alpha+1} = X_\alpha \setminus \bigcup\{U \in \mathcal{B} : \text{ent}_{>\varepsilon}(U \cap X_\alpha) \leq \kappa\}$, and
- $X_\lambda = \bigcap_{\alpha < \lambda} X_\alpha$ for λ a limit or ∞ .

For each $U \in \mathcal{B}$, let $\alpha(U)$ be the smallest ordinal such that $\text{ent}_{>\varepsilon}(U \cap X_{\alpha(U)}) \leq \kappa$ if it exists, and ∞ otherwise. Let $\beta = \sup\{\alpha(U) : \alpha(U) < \infty\}$. Since κ^+ is a regular cardinal, we must have that $\beta < \kappa^+$. In particular this implies that $X_\beta = X_{\beta+1} = X_\infty$.

Now assume that $X_\beta = \emptyset$. Let Y be a ($> \varepsilon$)-separated subset of $S_1(A)$ of cardinality κ^+ (this must exist because κ^+ is a regular cardinal). Since κ^+ is a regular cardinal, there must be some α such that $|Y \cap (X_\alpha \setminus X_{\alpha+1})| = \kappa^+$. Furthermore since \mathcal{B} has cardinality κ this implies that there must be a $U \in \mathcal{B}$ such that $|Y \cap (X_\alpha \setminus X_{\alpha+1}) \cap U| = \kappa^+$, but this implies that $\text{ent}_{>\varepsilon}(U \cap X_\alpha) \geq \kappa^+$, which is a contradiction. Therefore X_β is non-empty.

Now X_β must have the property that for any $U \in \mathcal{B}$, $\text{ent}_{>\varepsilon}(U \cap X_\beta) \geq \kappa^+$, so in particular in any non-empty open subset of X_β there exists p, q with $d_{\Delta, A}(p, q) > \varepsilon$.

Let $F_\emptyset = S_1(A)$. For each $\sigma \in 2^{<\omega}$, given F_σ , a closed set whose interior has non-empty intersection with X_β , find $p, q \in X_\beta \cap \text{int}_{S_1(A)} F_\sigma$ such that $d_{\Delta, A}(p, q) > \varepsilon$. Find $F_{\sigma \frown 0}$, a closed set such that $p \in \text{int}_{S_1(A)} F_{\sigma \frown 0}$, $F_{\sigma \frown 0} \subset \text{int}_{S_1(A)} F_\sigma$, and such that $F_{\sigma \frown 0} \cap B_{\leq \varepsilon}^{d_{\Delta, A}}(q) = \emptyset$. Then find a $F_{\sigma \frown 1}$, a closed set such that $q \in \text{int}_{S_1(A)} F_{\sigma \frown 1}$, $F_{\sigma \frown 1} \subset \text{int}_{S_1(A)} F_\sigma$, and such that $F_{\sigma \frown 1} \cap F_{\sigma \frown 0}^{d_{\Delta, A} \leq \varepsilon} = \emptyset$. Then by construction $\{F_\sigma\}_{\sigma \in 2^{<\omega}}$ is a $(d_{\Delta, A}, \varepsilon)$ -perfect tree, so T is not Δ -t.t.

(iii) The cardinality of a complete metric space of density character λ is always either λ or λ^ω . This implies that the cardinality of $S_1(A)$ for any set of parameters A with

$|A| \leq \kappa$ must be $\leq \kappa$, since it's $d_{\Delta, A}$ -density character is $\leq \kappa = \kappa^\omega$. Therefore T is stable with regards to the discrete metric and in particular stable. \square

Proposition 6.6.20. *For any theory T and distortion system Δ for T , if T is Δ -t.t., then for any set of parameters A and any closed set $F \subseteq S_n(A)$, $d_{\Delta(A)}$ -atomic-in- F types are dense in F .*

Proof. This follows from Corollary 3.8 of [BY08c]. \square

Proposition 6.6.21. *For any countable first-order theory T , there are models of any density character that only realize separably many types over countable parameter sets.*

Proof. This follows from Proposition 3.37 of [BY05]. \square

Corollary 6.6.22. *If a countable theory T with distortion system Δ is Δ - κ -categorical for uncountable κ , then T is Δ - ω -stable, so in particular it is Δ -t.t. and Δ - κ -stable for every $\kappa \geq \aleph_0$.*

Proof. Assume that T is not Δ - ω -stable. Then there is some countable parameter set A and an $\varepsilon > 0$ such that there is an uncountable $(d_{\Delta, A} > \varepsilon)$ -separated set of types P in $S_1(A)$. Let \mathfrak{M} be a model of density character κ that realizes A and every element of P . Let \mathfrak{N} be a model of T of density character κ such that for any countable $B \subset \mathfrak{N}$, the set of types in $S_1(B)$ realized in \mathfrak{N} is separable with respect to the d -metric.

Find a $\delta > 0$ small enough that $\delta_\Delta(\text{tp}(ab), \text{tp}(ce)) < \delta$ then $|d(a, b) - d(c, e)| < \frac{\varepsilon}{9}$. Find an almost correlation R between \mathfrak{M} and \mathfrak{N} with $\text{dis}_\Delta(R) < \frac{\varepsilon}{9} \downarrow \delta$. By replacing each element of A with a Cauchy sequence if necessary we may assume that A is in the domain of R (extending each element of P so that they are still realized in \mathfrak{M} , note that

they are still $(d_{\Delta,A} > \varepsilon)$ -separated). Let B be some ω -tuple of elements of \mathfrak{N} correlated with A by R .

Now for each $p \in P$, find a type p' with $d(p, p') < \frac{\varepsilon}{3}$ such that p' is still realized in \mathfrak{M} but by something in the domain of R . P' is now $(d_{\Delta,A} > \frac{\varepsilon}{3})$ -separated. Now for each $p' \in P'$, let $q \in S_1(B)$ be a type such that q is realized in \mathfrak{N} by something correlated to a realization of p' in \mathfrak{M} . Let Q be the collection of these types.

By unpacking definitions we have that for any countable parameter set E , the metric $d_{\Delta,E}$ on $S_1(E)$ is equivalent to the metric $\delta_{D(\Delta,c)}^\omega$ restricted to $S_1(E)$ as a subspace of $S_\omega(T_c)$, where the ω variables correspond to an enumeration of E , c is a fresh constant symbol corresponding to the free variables in the types in $S_1(E)$, and T_c is the (incomplete) theory of T with an extra constant (and no new axioms).

Pick $q_0, q_1 \in Q$, corresponding to $p'_0, p'_1 \in P'$. The choice of δ implies that $\delta_{D(\Delta,c)}^\omega(p'_0, q_0) < \frac{\varepsilon}{9}$ and $\delta_{D(\Delta,c)}^\omega(q_1, p'_1) < \frac{\varepsilon}{9}$. Since

$$d_{\Delta,A}(p'_0, p'_1) \leq \delta_{D(\Delta,c)}^\omega(p'_0, q_0) + \delta_{D(\Delta,c)}^\omega(q_0, q_1) + \delta_{D(\Delta,c)}^\omega(q_1, p'_1) \text{ and}$$

$$d_{\Delta,A}(p'_0, p'_1) \leq \delta_{D(\Delta,c)}^\omega(p'_0, q_0) + d(q_0, q_1) + \delta_{D(\Delta,c)}^\omega(q_1, p'_1),$$

this implies that $d(q_0, q_1) > \frac{\varepsilon}{3} - \frac{2\varepsilon}{9} = \frac{\varepsilon}{9}$. So Q is $(d > \frac{\varepsilon}{9})$ -separated which is a contradiction. Therefore T is Δ - ω -stable. \square

Definition 6.6.23. For a countable theory T with distortion system Δ , a model $\mathfrak{M} \models T$ is Δ - κ -saturated if for any $A \subseteq \mathfrak{M}$ with $|A| < \kappa$, \mathfrak{M} realizes a $d_{\Delta,A}$ -dense subset of $S_1(A)$.

If $\#\text{dc}\mathfrak{M} = \kappa$, we say that \mathfrak{M} is Δ -saturated. \triangleleft

When Δ is the collection of all formulas, a structure is Δ - κ -saturated if and only if it is κ -saturated.

At this point there is a notable omission. Given $\mathfrak{M}, \mathfrak{N} \models T$, both Δ -saturated with the same uncountable density character, it's unclear whether or not we can conclude $\rho_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$. If we could then we would be able to prove the full analog of Morley's theorem for distortion systems.

Proposition 6.6.24. *If T is a complete, countable, Δ - κ -stable theory with non-compact models for Δ , a distortion system for T , then for every uncountable regular $\lambda \leq \kappa$, T has a Δ - λ -saturated model of density character κ .*

Proof. Let \mathfrak{M}_0 be any pre-model of density character and cardinality κ . Form a continuous elementary chain $\{\mathfrak{M}_i\}_{i < \lambda}$ of length κ of pre-models of density character and cardinality κ such that for each i , \mathfrak{M}_{i+1} realizes a $d_{\Delta, \mathfrak{M}_i}$ -dense set of types in $S_1(\mathfrak{M}_i)$.

Finally let \mathfrak{M} be the completion of the union. Clearly $\#^{\text{dc}}\mathfrak{M} = \kappa$. Any subset A of \mathfrak{M} of cardinality $< \lambda$ is in the closure of some \mathfrak{M}_i . Then $S_1(\mathfrak{M}_i A) = S_1(\mathfrak{M}_i)$, so we have that a $d_{\Delta, \mathfrak{M}_i A}$ -dense subset of $S_1(\mathfrak{M}_i A)$ is realized in \mathfrak{M} . This implies that a $d_{\Delta, A}$ -dense subset of $S_1(A)$ is realized in \mathfrak{M} , and so \mathfrak{M} is Δ - λ -saturated. \square

Corollary 6.6.25. *If a countable theory T with distortion system Δ is Δ - κ -categorical for some $\kappa \geq \aleph_1$, then every model of T of density character κ is Δ -saturated.*

Proof. Let \mathfrak{M} be a model of T of density character $\kappa \geq \aleph_1$. For any regular uncountable $\lambda \leq \kappa$, let \mathfrak{N} be a Δ - λ -saturated model of T of density character κ . Let $A \subset \mathfrak{M}$ be any subset of cardinality $< \lambda$. Pick $p \in S_1(A)$ and $\varepsilon > 0$. Find a $\delta > 0$ small enough that $\delta_\Delta(\text{tp}(ab), \text{tp}(ce)) < \delta$ then $|d(a, b) - d(c, e)| < \frac{\varepsilon}{3}$. Find an almost correlation R between \mathfrak{M} and \mathfrak{N} with distortion $< \frac{\varepsilon}{3} \downarrow \delta$. Every element of A is a metric limit of points in the domain of R . Let A' be a set containing a sequence limiting to each element of A (note that we still have $|A'| < \lambda$), and let p' be some extension of p to $S_1(A')$. Go to a

large enough elementary extension of $(\mathfrak{M}, \mathfrak{N}, R)$ that p' is realized over A' by some m and such that R is a correlation. Let $B \subseteq \mathfrak{N}$ be correlated to A' by R so that $|B| < \lambda$, and let e be correlated to m in the elementary extension. Let q be $\text{tp}(e/B)$.

By Δ - λ -saturation, there is some type $r \in S_1(B)$ such that \mathfrak{N} realizes r with some e' and $d_{\Delta, B}(q, r) < \frac{\varepsilon}{5}$. Find e'' such that $d(e', e'') < \frac{\varepsilon}{3}$ and such that e'' is correlated to some $m' \in \mathfrak{M}$ by R .

Just like in the proof of Corollary 6.6.22, we have that

$$\begin{aligned} d_{\Delta, A}(p', \text{tp}(m'/A')) &\leq \delta_{D(\Delta(A), c)}^\omega(p', q) \\ &\quad + d(e', e'') \\ &\quad + \delta_{D(\Delta(A), c)}^\omega(\text{tp}(e''/B), \text{tp}(m'/A')), \end{aligned}$$

where we're thinking of A and B as ω -tuples when computing $\delta_{D(\Delta(A), c)}^\omega$. So we have that $d_{\Delta, A}(p', \text{tp}(m'/A')) < \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon$.

Since we can do this for any A with $|A| < \lambda$, any $p \in S_1(A)$, and any $\varepsilon > 0$, we have that \mathfrak{M} is Δ - λ -saturated. Since we can do this for any regular $\lambda \leq \kappa$, we have that \mathfrak{M} is Δ -saturated. \square

Definition 6.6.26. We say that T is *weakly Δ - κ -categorical* if every model of T of density character κ is Δ -saturated. \triangleleft

It's immediate from Proposition 6.6.21 that a weakly Δ - κ -categorical theory for some uncountable κ is Δ - ω -stable.

Definition 6.6.27. Let T be a complete theory and Δ a distortion system for T .

- For any parameter set A , a Δ -*construction over A* is a sequence $\{b_i\}_{i < \lambda}$ such that

for every $i < \lambda$, $\text{tp}(b_i/Ab_{<i})$ is $d_{\Delta, Ab_{<i}}$ -atomic.

- A model $\mathfrak{M} \models T$ is Δ -constructible over A if $A \subseteq \mathfrak{M}$ and there is an enumeration of a dense subset of \mathfrak{M} which is a Δ -construction over A . \triangleleft

Proposition 6.6.28. *For any complete theory T , Δ , a distortion system for T , and parameter set A , if $\{b_i\}_{i < \lambda}$ is a Δ -construction over A , then for any finite tuple $\bar{b} \in \{b_i\}$, $\text{tp}(\bar{b}/A)$ is $d_{\Delta, A}$ -atomic.*

Proof. We have already done most of the work in Proposition 6.5.3 and Lemma 6.5.4. We just need to apply it inductively. Assume that we're shown that for any $i_0 < i_1 < \dots < i_k < \alpha$ that $\text{tp}(b_{i_0}b_{i_1} \dots b_{i_k}/A)$ is $d_{\Delta, A}$ -atomic. Let $j_0 < j_1 < \dots < j_\ell \leq \alpha$ be any tuple of indices, and consider $b_{j_0}b_{j_1} \dots b_{j_\ell}$. By Proposition 6.5.3 we know there exists a countable set of elements of the form b_i with $i < \alpha$ such that $\text{tp}(b_{j_0}b_{j_1} \dots b_{j_\ell}/AB)$ is $d_{\Delta, A}$ -atomic. By Lemma 6.5.4 this implies that $\text{tp}(b_{j_0}b_{j_1} \dots b_{j_\ell}/A)$ is $d_{\Delta, A}$ -atomic, as required. \square

Proposition 6.6.29. *Fix T , a countable complete Δ - ω -stable theory with non-compact models, for Δ , a distortion system for T . If $\mathfrak{M} \models T$ with $\#\text{dc}\mathfrak{M} = \kappa \geq \aleph_1$, then for any regular $\lambda \leq \kappa$, \mathfrak{M} has arbitrarily large elementary extensions \mathfrak{N} such that for any $A \subset \mathfrak{M}$ with $|A| < \lambda$ the set of types realized by \mathfrak{N} in $S_1(A)$ is contained in the $d_{\Delta, A}$ -closure of the set of types realized by \mathfrak{M} in $S_1(A)$.*

Proof. Fix regular $\lambda \leq \kappa$. Find an $\varepsilon > 0$ such that $\text{ent}_{>\varepsilon}(\mathfrak{M}, d) \geq \lambda$, which must exist because λ is regular. Let Q be a maximal ($>\varepsilon$)-separated subset of \mathfrak{M} of cardinality $\geq \lambda$. Let $X \subset S_1(\mathfrak{M})$ be the set of all types p such that for every open neighborhood $U \ni p$, $|U \cap Q| \geq \lambda$. Note that X is a closed set and by compactness is non-empty. Also

note that since Q was chosen to be maximal, $\mathfrak{M} \cap X = \emptyset$, since for any realized type p there is some $a \in \mathfrak{M}$ such that $d(p, a) < \varepsilon$, which is an open neighborhood of p whose intersection with Q has cardinality 1, which is in particular less than λ .

Now since T is Δ - ω -stable, $d_{\Delta, \mathfrak{M}}$ -atomic-in- X types are dense in X . Let p be some $d_{\Delta, \mathfrak{M}}$ -atomic-in- X type. Let a be a realization of that type. By Proposition 6.6.20, there exists a Δ -constructible model \mathfrak{N} over $\mathfrak{M}a$ (specifically, $d_{\Delta, A}$ -atomic types are dense in every $S_1(A)$). Note that \mathfrak{N} is a proper elementary extension of \mathfrak{M} .

Let b be some element of \mathfrak{N} , and let $E \subset \mathfrak{M}$ be some set of parameters of cardinality $\mu < \lambda$. Since \mathfrak{N} is Δ -constructible over $\mathfrak{M}a$, $\text{tp}(b/\mathfrak{M}a)$ is $d_{\Delta, \mathfrak{M}a}$ -atomic. Let $\varphi(x, \bar{m}, a)$ be a formula (with \bar{m} possibly an ω -tuple) such that $\mathfrak{N} \models \varphi(b, \bar{m}, a) \leq 0$ and for every type $q \in S_1(\mathfrak{M}a)$, $d_{\Delta, \mathfrak{M}a}(q, \text{tp}(b/\mathfrak{M}a)) \leq \varphi(q, \bar{m}, a)$. In particular what this means is that if we think of c as a fresh constant symbol, then for any $D(\Delta(\mathfrak{M}a), c)$ -sentence $\psi(c, \bar{m}, a)$, $|\psi^{\mathfrak{N}}(b, \bar{m}, a) - \psi(q, \bar{m}, a)| \leq \varphi(q, \bar{m}, a)$. We may assume that $\bar{m} \in E$. Let $\{\psi_i^0(c, \bar{e}_i)\}_{i < \mu}$ be a collection of $D(\Delta(\mathfrak{E}), c)$ -sentences which are dense in the uniform norm (treating c as a variable). Then for each $i < \mu$, let $\psi_i(c, \bar{e}_i) = |\psi_i^0(c, \bar{e}_i) - \psi_i^{0, \mathfrak{N}}(b, \bar{e}_i)|$, and note that each $\psi_i(c, \bar{e}_i)$ is also a $D(\Delta(E), c)$ -sentence. Now we have that for each $i < \mu$ and for any $q \in S_1(\mathfrak{M}a)$, $\mathfrak{N} \models \sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, a)$ and of course $\mathfrak{N} \models \inf_x \varphi(x, \bar{m}, a)$, so these two statements are parts of $\text{tp}(a/\mathfrak{M}) = p$.

Let χ be a \mathfrak{M} -formula such that $\chi(p) = 0$ and for every $q \in X$, $d_{\Delta, \mathfrak{M}}(p, q) \leq \chi(q)$, so in particular for any $D(\Delta(\mathfrak{M}), c)$ -sentence $\theta(c, \bar{m})$, we have $|\theta(p, \bar{m}) - \theta(q, \bar{m})| \leq \chi(q)$. Again by extending E we may assume that χ is an E -formula.

Now pick $\delta > 0$, and find a $\gamma > 0$ such that if $d_{\Delta, \mathfrak{M}}(\text{tp}(a'b'/\mathfrak{M}), \text{tp}(a''b''/\mathfrak{M})) \leq \gamma$ then $|\varphi(b', \bar{m}, a') - \varphi(b'', \bar{m}, a'')| < \frac{\delta}{3}$, and note that this implies that if $d_{\Delta, \mathfrak{M}}(\text{tp}(a'/\mathfrak{M}), \text{tp}(a''/\mathfrak{M})) \leq \gamma$ then for any b' , $|\varphi(b', \bar{m}, a') - \varphi(b', \bar{m}, a'')| < \frac{\delta}{3}$ and also the same thing for formulas of

the form $\sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, y)$. Consider the sets

$$H_{-1} = Q \cap [\chi < \gamma] \setminus \left[\inf_x \varphi(x, \bar{m}, y) < \frac{\delta}{2} \right] \quad \text{and}$$

$$H_i = Q \cap [\chi < \gamma] \setminus \left[\sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, y) < \frac{\delta}{2} \right].$$

for $i < \mu$ (as subsets of $S_1(\mathfrak{M})$). Note that since $[\chi < \gamma]$ is an open neighborhood of p , it must be the case that $|Q \cap [\chi < \gamma]| \geq \lambda$. Now assume that $|H_i|$ is $\geq \lambda$. This implies that $X \cap \text{cl}_{S_1(\mathfrak{M})} H_i \neq \emptyset$, so in particular there is some $q \in X$ such that $\chi(q) \leq \gamma$ and yet either $\inf_x \varphi(x, \bar{m}, q) \geq \frac{\delta}{2}$ or $\sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, q) \geq \frac{\delta}{2}$, which contradicts the choice of γ , so it must be the case that $|H_i| < \lambda$ for every i with $-1 \leq i < \mu$. Therefore since λ is a regular cardinal we must have that

$$\left| Q \cap [\chi < \gamma] \cap \left[\inf_x \varphi(x, \bar{m}, y) < \frac{\delta}{2} \right] \cap \bigcap_{i < \mu} \left[\sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, y) < \frac{\delta}{2} \right] \right| \geq \lambda,$$

so there exists some $a' \in Q$ such that $\mathfrak{M} \models \inf_x \varphi(x, \bar{m}, a') < \frac{\delta}{2}$ and $\sup_x \psi_i(x, \bar{e}_i) \div \varphi(x, \bar{m}, a') < \frac{\delta}{2}$ for each $i < \mu$. Let b' be an element of \mathfrak{M} such that $\mathfrak{M} \models \varphi(b', \bar{m}, a') < \frac{\delta}{2}$. Then we have that for each $i < \mu$, $\mathfrak{M} \models \psi_i(b', \bar{e}_i) \div \varphi(b', \bar{m}, a') < \frac{\delta}{2}$, so together this implies that $\mathfrak{M} \models \psi_i(b', \bar{e}_i) < \delta$ for every $i < \mu$. By the choice of the ψ_i 's, this implies that $d_{\Delta, E}(\text{tp}(b'/E), \text{tp}(b/E)) \leq \delta$. Since we can do this for any $\delta > 0$, we have that the set of types in $S_1(E)$ realized in \mathfrak{N} is in the $d_{\Delta, E}$ -metric closure of the set of types in $S_1(E)$ realized in \mathfrak{M} .

Now we are free to iterate this process to form arbitrarily large elementary extensions $\mathfrak{N}' \succ \mathfrak{M}$ such that for any set of parameters $E \subset \mathfrak{M}$ with $|E| < \lambda$, if \mathfrak{N}' realizes some type in $S_1(E)$, then it is in the $d_{\Delta, E}$ -metric closure of the types in $S_1(E)$ realized in

\mathfrak{M} . □

Corollary 6.6.30. *Fix T a countable complete theory with non-compact models and Δ , a distortion system for T . For any $\kappa \geq \aleph_1$, if T is weakly Δ - κ -categorical, then it is weakly Δ - λ -categorical for every $\lambda \geq \aleph_1$ with $\lambda \leq \kappa$.*

Proof. Assume that T fails to be weakly Δ - κ -categorical. Let \mathfrak{M} be a model of density character κ that fails to be Δ - κ -saturated. Let $A \subset \mathfrak{M}$ be a set of parameters with $|A| < \kappa$ such that for some type $p \in S_1(A)$ and some $\varepsilon > 0$, every type $q \in S_1(A)$ realized in \mathfrak{M} has $d(p, q) \geq \varepsilon$. Then by the previous proposition \mathfrak{M} has arbitrarily large elementary extensions with the same property, so in particular for any $\lambda > \kappa$, T has a model of density character λ that fails to be Δ -saturated, so T is not weakly Δ - λ -categorical. □

Theorem 6.6.31. *Fix T a countable complete theory with non-compact models and Δ , a distortion system for T . For any $\kappa \geq \aleph_1$, if T is weakly Δ - κ -categorical, then it is weakly Δ - λ -categorical for every $\lambda \geq \aleph_1$.*

Proof. By the previous Corollary 6.6.30, we know that the collection of uncountable cardinalities for which a theory T is weakly Δ - κ -categorical is always an initial segment of the uncountable cardinals. Assume that it is not all of them. Find κ large enough that T is not weakly Δ - κ -categorical and such that for any $\lambda < \kappa$, $\lambda^\omega < \kappa$ and such that $\text{cf}(\kappa) \geq \omega_1$ (such a cardinal can always be found).

Let \mathfrak{M} be a model of T of cardinality κ which is not Δ -saturated. Let $A \subset \mathfrak{M}$ be a set of parameters with $|A| < \kappa$ such that for some type $p \in S_1(A)$ and some $\varepsilon > 0$, every type $q \in S_1(A)$ realized in \mathfrak{M} has $d_{\Delta, A}(p, q) \geq \varepsilon$ (i.e. A witnesses that \mathfrak{M} is not Δ -saturated). Since for any superset $A' \supseteq A$, the natural restriction map $(S_1(A'), d_{\Delta, A'}) \rightarrow$

$(S_1(A), d_{\Delta, A})$ is 1-Lipschitz, we can freely pass to a larger set of parameters and preserve this condition. So we may assume that $|A| = |A|^\omega$ by passing to a larger parameter set if necessary (we can do this since we have ensured that $\lambda^\omega < \kappa$ for any $\lambda < \kappa$). Let $\lambda = |A|$. Since T is Δ - ω -stable, it is stable and so in particular λ -stable.

Find sufficiently small $\varepsilon > 0$ such that we can find $\{b_i\}_{i < \lambda^+}$, a $(>\varepsilon)$ -separated sequence of elements of \mathfrak{M} . By Proposition 4.17 of [BY05] there exists (in the monster model) an array $\{c_i^j\}_{i < \lambda^+, j < \omega}$ such that $d(c_i^j, c_i^{j+1}) < 2^{-j}$, such that for each $j < \omega$ there is a sub-sequence $I \subseteq \lambda^+$ such that $\{c_i^j\}_{i < \lambda^+} \equiv_A \{b_i\}_{i \in I}$, and such that the sequence of limits $\{c_i^\omega\}_{i < \lambda^+}$ is an A -indiscernible sequence. Let $C = \{c_i^\omega\}_{i < \omega}$.

Let Σ be a countable dense subset of the collection of finitary $D(\Delta, x)$ formulas (where x is being treated as the fresh constant symbol). Formulas in Σ are of the form $\varphi(x, \bar{y})$. The important thing is that Σ has the property that for any set of parameters E , $d_{\Delta, E}(\text{tp}(f_0/E), \text{tp}(f_1/E)) = \sup_{\varphi \in \Sigma, \bar{e} \in E} |\varphi(f_0, \bar{e}) - \varphi(f_1, \bar{e})|$. We may assume that Σ is closed under $\varphi \mapsto (\varphi \div r)$ for each rational r .

(*) Note that for each restricted AC -formula $\varphi(x, \bar{a}, \bar{c})$ such that $\models \inf_x \varphi(x, \bar{a}, \bar{c}) \leq 0$, there is an A -formula $\psi_{\varphi(\cdot, \bar{a})}(x, \bar{a}_{\varphi(\cdot, \bar{a}, \bar{c})})$ with $\psi \in \Sigma$ such that $p(x) \vdash \psi_{\varphi(\cdot, \bar{a})}(x, \bar{a}_{\varphi(\cdot, \bar{a})}) \leq 0$ and such that

$$\left\{ \varphi(x, \bar{a}, \bar{c}) < \frac{1}{2}, \psi_{\varphi(\cdot, \bar{a})}(x, \bar{a}_{\varphi(\cdot, \bar{a}, \bar{c})}) > \frac{\varepsilon}{2} \right\}$$

is consistent. This holds because every type realized in \mathfrak{M} has $d_{\Delta, A}$ -distance $\geq \varepsilon$ from p and because we can approximate C with something of the form $\{b_i\}_{i \in I}$ for some $I \subseteq \lambda^+$, which is a set of elements in \mathfrak{M} .

Now let $A_0 = \emptyset$, and for each $n < \omega$, given A_n , let A_{n+1} be the collection of all a 's occurring in some tuple of the form $\bar{a}_{\varphi(\cdot, \bar{a}, \bar{c})}$ where $\varphi(x, \bar{a}, \bar{c})$ is a restricted $A_n C$ -formula.

Finally let $A_\omega = \bigcup_{n < \omega} A_n$. Clearly A_ω is a countable set.

Now by construction we have that condition (*) holds with A_ω replacing A . Also note that C is still an indiscernible sequence over A_ω (because $A_\omega \subseteq A$). Let C' be an extension of C to an indiscernible sequence of length ω_1 . Now let \mathfrak{N} be a Δ -constructible model over $A_\omega C'$. Note that $\#\text{dc}\mathfrak{N} = \aleph_1$. If q is the restriction of p to A_ω , we would like to argue that for every type $r \in S_1(A_\omega)$ realized in \mathfrak{N} , $d_{\Delta, A_\omega}(q, r) \geq \frac{\varepsilon}{2}$. Assume that this is false and that there is some $e \in \mathfrak{N}$ such that $d_{\Delta, A}(q, \text{tp}(e/A_\omega)) \leq \delta < \frac{\varepsilon}{2}$. By construction, $\text{tp}(e/A_\omega C')$ is $d_{\Delta, A_\omega C'}$ -atomic. This implies that there is a restricted formula $\chi(x, \bar{a}, \bar{c})$ such that $\models \chi(e, \bar{a}, \bar{c}) \leq 0$ and for any f if $\models \chi(f, \bar{a}, \bar{c}) < \frac{1}{2}$, then $d_{\Delta, A}(\text{tp}(e/A_\omega C'), \text{tp}(f/A_\omega C')) \leq \frac{\varepsilon}{2} - \delta$. In particular this implies that $d_{\Delta, A}(\text{tp}(f/A_\omega C), q) \leq \frac{\varepsilon}{2}$. Therefore for every A -formula $\varphi(x, \bar{a})$ with $\varphi \in \Sigma$ such that $q(x) \vdash \varphi(x, \bar{a})$, we have that $\models \varphi(f, \bar{a}) \leq \frac{\varepsilon}{2}$. The fact

$$\forall x \left(\chi(x, \bar{a}, \bar{c}) < \frac{1}{2} \rightarrow \varphi(x, \bar{a}) \leq \frac{\varepsilon}{2} \right)$$

is a closed formula satisfied by \bar{c} . By indiscernibility this is true of every tuple $\bar{c}' \in C$ with the same order type as \bar{c} , but this is inconsistent with the modified condition (*) (replacing A with A_ω) above, since this implies that $\{\chi(x, \bar{a}, \bar{c}') < \frac{1}{2}, \psi_{\chi(\cdot, \bar{a}, \bar{c}')}(\bar{c}', \bar{a}_{\chi(\cdot, \bar{a}, \bar{c}')})) > \frac{\varepsilon}{2}\}$ is inconsistent (by setting $\varphi(x, \bar{a})$ to $\psi_{\chi(\cdot, \bar{a}, \bar{c}')}(\bar{c}', \bar{a}_{\chi(\cdot, \bar{a}, \bar{c}')}))$). Therefore no such e can exist and for any type $r \in S_1(A_\omega)$ realized in \mathfrak{N} , $d_{\Delta, A_\omega}(q, r) \geq \frac{\varepsilon}{2}$. Therefore \mathfrak{N} is not Δ -saturated and T is not weakly Δ - \aleph_1 -categorical.

But this contradicts our assumption, so there is no largest uncountable κ such that T is weakly Δ - κ -categorical and in fact T is weakly Δ - κ -categorical for all uncountable κ . □

	ω_1	Lip- ω_1	GH- ω_1	None
ω	Trivial	Unknown	Unknown	Trivial
Lip- ω	Ex. 6.6.36	Ex. 6.6.42	Unknown	Ex. 6.6.39
GH- ω	Ex. 6.6.37	Ex. 6.6.43	Ex. 6.6.44	Ex. 6.6.38
None	Trivial	Ex. 6.6.40	Ex. 6.6.41	Trivial

Table 7: Known combinations of separable and inseparable ordinary, Lipschitz, and Gromov-Hausdorff categoricity for metric space theories.

Question 6.6.32. *If T is weakly Δ - κ -categorical for some $\kappa \geq \aleph_1$, does it follow that T is Δ - κ -categorical?*

In the particular case of a discrete theory with a stratified language we can resolve this positively.

Proposition 6.6.33. *Suppose Δ is a distortion system for some complete discrete theory T equivalent to some stratified language \mathcal{L} , and suppose that $\mathfrak{M}, \mathfrak{N} \models T$ are Δ -saturated models of the same cardinality, then they are Δ -approximately isomorphic.*

Proof. For each $n < \omega$, \mathfrak{M} and \mathfrak{N} are saturated as \mathcal{L}_n -structures, therefore they are isomorphic as \mathcal{L}_n -structures. Thus they are Δ -approximately isomorphic. \square

Corollary 6.6.34. *Suppose Δ is a distortion system for some complete discrete theory T equivalent to some stratified language \mathcal{L} . If T is Δ - κ -categorical for some uncountable κ then for every uncountable λ , T is Δ - λ -categorical.*

6.6.3 Some Examples and the Relationship between Different Notions of Categoricity

This section is a case study of the relationship between ordinary categoricity and Lipschitz and Gromov-Hausdorff approximate categoricity in the theories of metric spaces.

The results are summarized in the Table 7, where ‘ $\kappa \in \{\omega, \omega_1\}$ ’ means κ -categorical, ‘Lip- κ ’ means Lip- κ -categorical and not κ -categorical, ‘GH- κ ’ means GH- κ -categorical and not Lip- κ -categorical, and ‘none’ means not GH- κ -categorical. It’s not hard to prove that δ_{Lip} uniformly dominates δ_{GH} and that therefore Lip- κ -categoricity implies GH- κ -categoricity. The boxes labeled ‘Trivial’ are trivial in the sense that it is very easy to encode discrete structures in finite languages as metric spaces [Han], and to verify that such structures fall in the corresponding groups here. Of course we haven’t proven that Δ - ω_1 -categoricity is equivalent to Δ - κ -categoricity for all uncountable κ , but needless to say if we had a counterexample we would have mentioned it by now.

‘Unknown’ indicate combinations that are not currently known to be possible. There seems to be a general phenomenon where the combination of ordinary ω -categoricity and strictly approximate ω_1 -categoricity is impossible. In the case of theories with a $\{0, 1\}$ -valued metric (although allowing $[0, 1]$ -valued predicates), ω -categoricity implies that the \emptyset -type spaces are all finite, so any such theory is interdefinable with a purely discrete theory. That said a distortion system for such a theory could still be non-trivial if the theory does not admit quantifier elimination to a finite language. For any ω -categorical discrete theory that admits quantifier elimination to a finite language, all distortion systems are uniformly equivalent to isomorphism, so here we clearly have that Δ - ω_1 -categoricity implies ω_1 -categoricity. This suggests a pair of purely discrete questions.

Question 6.6.35. (i) *Does there exist a discrete theory T in a stratified language \mathcal{L} such that T is ω -categorical but only approximately ω_1 -categorical?*

(ii) *Does there exist a discrete theory T in a stratified language $\mathcal{L} = \bigcup_{i < \omega} \mathcal{L}_i$ such that T is ω -categorical, $T \upharpoonright \mathcal{L}_i$ is ω_1 -categorical for every $i < \omega$, but T is not*

ω_1 -categorical?

Note that a positive answer for part (ii) would imply a positive answer for part (i). An example of (ii) would have to be rather strange. It is easy to show that such a T cannot have any Vaughtian pairs, so T must fail to be ω -stable. Since it is ω -categorical this would imply that it is strictly stable. There is really only one known strictly stable ω -categorical theory, constructed by Hrushovski, but it has a finite language, whereas an example of what we need would necessarily have an infinite language.

Now we turn to the examples in the chart. The following two examples are very similar. The idea is to encode a sequence of constants in a structureless set in increasingly ‘harder to detect’ ways.

Example 6.6.36. *A metric space theory that is strictly Lip- ω -categorical and ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $\omega \times \{0, 1\}$ with $d((i, j), (k, \ell)) = 1$ if $i \neq k$ and $d((i, 0), (i, 1)) = \frac{1}{2} + 2^{-i-2}$. Let $T = \text{Th}(\mathfrak{M})$. \square

Example 6.6.37. *A metric space theory that is strictly GH- ω -categorical and ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $\omega \times \{0, 1\}$ with $d((i, j), (k, \ell)) = 1$ if $i \neq k$ and $d((i, 0), (i, 1)) = 2^{-i-1}$. Let $T = \text{Th}(\mathfrak{M})$. \square

The next two examples are also similar to each other.

Example 6.6.38. *A metric space theory that is strictly GH- ω -categorical and not GH- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $[0, \frac{1}{2}] \times \{0, 1\}$ with $d((x, i), (y, j)) = 1$ if $x \neq y$ and $d((x, 0), (x, 1)) = x$. Let $T = \text{Th}(\mathfrak{M})$. \square

Example 6.6.39. *A metric space theory that is strictly Lip- ω -categorical and not GH- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $[\frac{1}{4}, \frac{1}{2}] \times \{0, 1\}$ with $d((x, i), (y, j)) = 1$ if $x \neq y$ and $d((x, 0), (x, 1)) = x$. Let $T = \text{Th}(\mathfrak{M})$. \square

The following Examples 6.6.40 and 6.6.41 are the prototypes for the subsequent Examples 6.6.42, 6.6.43, and 6.6.44. The idea is to encode as a metric space a structure that is \mathbb{Z} -chains with maps into $[-1, 1]$ of the form $\cos(n + \theta)$. Any two such chains are ‘approximately isomorphic’ regardless of their values of θ , since 1 radian is an irrational rotation.

Example 6.6.40. *A metric space theory that is not GH- ω -categorical and is strictly Lip- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $\mathbb{Z} \times \{0, 1\}$ with $d((n, i), (m, j)) = 1$ if $|n - m| > 1$, $d((n, i), (n + 1, j)) = \frac{1}{2}$, and $d((n, 0), (n, 1)) = \frac{1}{4} + \frac{1}{8} \cos(n)$. Let $T = \text{Th}(\mathfrak{M})$. \square

Example 6.6.41. *A metric space theory that is not GH- ω -categorical and is strictly GH- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $\mathbb{Z} \times \{0, 1\}$ with $d((n, i), (m, j)) = 1$ if $|n - m| > 1$, $d((n, i), (n + 1, j)) = \frac{1}{2}$, and $d((n, 0), (n, 1)) = \frac{1}{8} + \frac{1}{8} \cos(n)$. Let $T = \text{Th}(\mathfrak{M})$. \square

Now we modify the previous examples with ‘increasingly hard to detect’ constants, analogous to Examples 6.6.36 and 6.6.37. This forces the separable model to have infinitely many \mathbb{Z} -chains.

Example 6.6.42. *A metric space theory that is strictly Lip- ω -categorical and strictly Lip- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $\mathbb{N} \times \mathbb{Z} \times 0, 1$ with

- $d((a, n, i), (b, m, j)) = 1$ if $a \neq b$ or if $a = b$ and $|n - m| > 1$,
- $d((a, n, i), (a, n + 1, j)) = \frac{1}{2}$,
- $d((a, n, 0), (a, n, 1)) = \frac{1}{4} + \frac{1}{8} \cos(n)$ if $n \neq 0$, and
- $d((a, 0, 0), (a, 0, 1)) = \frac{1}{4} + \frac{1}{8} + 2^{-a-4}$.

Let $T = \text{Th}(\mathfrak{M})$. □

Example 6.6.43. *A metric space theory that is strictly GH- ω -categorical and strictly Lip- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $(\mathbb{N} \times \mathbb{Z} \times 0, 1) \cup (\mathbb{N} \times \{0\} \times \{2\})$ with

- $d((a, n, i), (b, m, j)) = 1$ if $a \neq b$ or if $a = b$ and $|n - m| > 1$,
- $d((a, n, i), (a, n + 1, j)) = \frac{1}{2}$,
- $d((a, n, 0), (a, n, 1)) = \frac{1}{4} + \frac{1}{8} \cos(n)$ if $n \neq 0$,
- $d((a, 0, 0), (a, 0, 1)) = \frac{1}{4} + \frac{1}{8}$,

- $d((a, 0, 0), (a, 0, 2)) = \frac{1}{4} + \frac{1}{8} + 2^{-a-4}$, and

- $d((a, 0, 1), (a, 0, 2)) = 2^{-a-4}$.

Let $T = \text{Th}(\mathfrak{M})$. □

Example 6.6.44. *A metric space theory that is strictly GH- ω -categorical and strictly GH- ω_1 -categorical.*

Description. Let \mathfrak{M} be a metric space whose universe is $(\mathbb{N} \times \mathbb{Z} \times 0, 1) \cup (\mathbb{N} \times \{0\} \times \{2\})$ with

- $d((a, n, i), (b, m, j)) = 1$ if $a \neq b$ or if $a = b$ and $|n - m| > 1$,

- $d((a, n, i), (a, n + 1, j)) = \frac{1}{2}$,

- $d((a, n, 0), (a, n, 1)) = \frac{1}{8} + \frac{1}{8} \cos(n)$ if $n \neq 0$,

- $d((a, 0, 0), (a, 0, 1)) = \frac{1}{8} + \frac{1}{8}$,

- $d((a, 0, 0), (a, 0, 2)) = \frac{1}{8} + \frac{1}{8} + 2^{-a-4}$, and

- $d((a, 0, 1), (a, 0, 2)) = 2^{-a-4}$.

Let $T = \text{Th}(\mathfrak{M})$. □

If we try to use this construction to fill in the missing Lip- ω -categorical and GH- ω_1 -categorical square by taking Example 6.6.41 and adding encoded constants in the style of Example 6.6.42, what we get is a theory that is only GH- ω -categorical.

Note that these last 5 are also examples showing that a Δ - ω_1 -categorical theory need not be unidimensional, since they contain many orthogonal types. It's even possible to

modify this idea to get a discrete theory in a stratified language that is approximately uncountably categorical and yet not unidimensional (rather than using an irrational rotation of the circle along \mathbb{Z} -chains, use an ‘irrational rotation’ of the 2-adic integers along \mathbb{Z} -chains, or, roughly equivalently, the structure \mathbb{N} together with predicates U_i for each $i < \omega$ such that $U_i(n)$ is true if and only if the i th binary digit of n is 1), but curiously these examples seem to be limited to having trivial geometry in their types. There are also examples of strictly approximately uncountably categorical discrete theories with non-trivial geometries (such as the theory of the vector space \mathbb{F}_p^ω together with a sequence of predicates encoding projections onto the first ω \mathbb{F}_p factors, with the obvious stratification), but these seem to be unidimensional. This suggests a question.

Question 6.6.45. *If T is a discrete theory in a stratified language \mathcal{L} which is approximately uncountably categorical and T has minimal types with non-trivial geometry, does it follow that T is unidimensional?*

Another natural question, which is related to Questions 6.6.35 and 6.6.45, arises from the observation that all of these examples are superstable.

Question 6.6.46. *If T is (weakly) Δ - κ -categorical for some uncountable κ , does it follow that T is superstable?*

Appendix A

A Compendium of Relevant Topological Facts

Here we will collect various topological facts which we have used throughout the body of the thesis.

A.1 (Pseudo-)metric Spaces

Fact A.1.1. *For any pseudo-metric space X , if $\{f_i\}_{i \in I}$ is a family of ω -uniformly continuous functions $f_i : X \rightarrow \mathbb{R}$, for some fixed modulus ω , and $\inf_{i \in I} f(a)$ exists for some $a \in X$, then $x \mapsto \inf_{i \in I} f_i(x)$ is well-defined everywhere and is an ω -uniformly continuous function. The same holds for sup.*

Fact A.1.2. *If ρ is a pseudo-metric on X and $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a sub-additive (i.e. $\alpha(r + s) \leq \alpha(r) + \alpha(s)$) non-decreasing function satisfying $\alpha(0) = 0$, then $\alpha \circ \rho$ is a pseudo-metric on X . In particular, $\lceil \rho \rceil$ is a pseudo-metric on X , and for any $a \geq 0$, $\min\{\rho, a\}$ is a pseudo-metric on X .*

If $\{\rho_i\}_{i \in I}$ is a family of pseudo-metrics on X such that for any $x, y \in X$, the set $\{\rho_i(x, y)\}_{i \in I}$ is bounded, then $\sup_{i \in I} \rho_i$ is a pseudo-metric on X .

If $\{\rho_i\}_{i \in I}$ is a family of pseudo-metric on X such that for any $x, y \in X$, $\sum_{i \in I} \rho_i(x, y)$

exists, then $\sum_{i \in I} \rho_i$ is a pseudo-metric on X .

More generally if $\{\rho_i\}_{i \in I}$ is a family of extended pseudo-metrics on X , then $\sup_{i \in I} \rho_i$ and $\sum_{i \in I} \rho_i$ are extended pseudo-metrics on X .

Fact A.1.3. Every closed subset of a metric space is G_δ . Every open subset is F_σ .

Fact A.1.4 (Lebesgue's Number Lemma). If X is a compact metric space, then for any finite open cover \mathcal{U} of X , there is an $\varepsilon > 0$ such that for any $x \in X$, $B_{\leq \varepsilon}(x) \subseteq U$ for some $U \in \mathcal{U}$.

Fact A.1.5. If (X, d) is a complete metric space, then either $|X| = \#^{\text{dc}} X$ or $|X| = (\#^{\text{dc}} X)^{\aleph_0}$.

Fact A.1.6. In any metric space X , for any $A, B \subseteq X$, the Hausdorff distance between A and B is equal to any of the following:

- $\inf\{\varepsilon > 0 : A \subseteq B^{<\varepsilon}, B \subseteq A^{<\varepsilon}\},$
- $\max\{\sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(a, b)\},$ and
- $\|d_{\text{inf}}(x, A) - d_{\text{inf}}(x, B)\|_\infty = \sup_{x \in X} |d_{\text{inf}}(x, A) - d_{\text{inf}}(x, B)|,$

where $\inf \emptyset = \text{diam}(X)$ and $\sup \emptyset = 0$ (in particular $\text{diam}(\emptyset) = 0$).

Furthermore, if X is a complete metric space then $\{F \subseteq X : F \text{ closed}\}$ is a complete (extended) metric space under d_H .

Recall that a modulus is a continuous function $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ which is non-decreasing and satisfies $\omega(0) = 0$. Moduli are often a convenient way of packaging ' ε - δ information.'

Lemma A.1.7. *Let $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a bounded function with $f(0) = 0$ and $f(x) \rightarrow 0$ as $x \rightarrow 0$. There exists a convex downward (and therefore sub-additive) modulus g such that $f \leq g$.*

Proof. Let $r > 0$ be such that for all $x \geq 0$, $f(x) < r$. Let

$$g(x) := \inf\{mx + c : m, c \geq 0, (\forall y \geq 0)my + c \geq f(y)\}.$$

Since $f(x)$ is bounded, $g(x)$ is finite for every x . We clearly have that $g \geq f$, and since $g(x)$ is an infimum of a family of convex downwards, non-increasing, non-negative functions, we have that $g(x)$ is convex downwards, non-increasing, and non-negative, so we just need to show that $g(0) = 0$ and that $g(x)$ is continuous.



For any $\varepsilon > 0$, there is a $\delta > 0$, such that for all $x \in [0, \delta]$, $f(x) < \varepsilon$. This implies that $\frac{r}{\delta}y + \varepsilon \geq f(y)$ for all $y \geq 0$, so $g(0) \leq \varepsilon$. Since we can do this for every $\varepsilon > 0$, we have that $g(0) = 0$. Furthermore, this implies that for all $x \in [0, \frac{\delta}{r}\varepsilon]$, $g(x) \leq 2\varepsilon$, so $g(x)$ is continuous at 0.

Fix $s > 0$. For any $x \geq s$, we have that

$$g(x) = \inf\left\{mx + c : m, c \geq 0, m \leq \frac{r}{s}, (\forall y \geq 0)my + c \geq f(y)\right\}.$$

Therefore, by Fact A.1.1, $g(x)$ is $\frac{r}{s}$ -Lipschitz on $[s, \infty)$, and therefore continuous. Since we can do this for any $s > 0$, $g(x)$ is continuous. \square

Proposition A.1.8. *Let $P(\varepsilon, \delta)$ be a statement such that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that $P(\varepsilon, \delta)$ holds. Furthermore suppose that if $P(\varepsilon, \delta)$ holds, then for any $\varepsilon', \delta' > 0$ with $\varepsilon' \geq \varepsilon$ and $\delta' \leq \delta$, $P(\varepsilon', \delta')$ holds as well.*

Figure 11: $\alpha(\varepsilon)$ in Proposition A.1.8 part (i)Figure 12: $\gamma(\delta)$ and $\beta(\delta)$ in Proposition A.1.8 part (ii)

(i) There exists a modulus α such that for every $\varepsilon > 0$, $P(\varepsilon, \alpha(\varepsilon))$ holds.

(ii) There exists a convex downwards (and therefore sub-additive) modulus β such that for every sufficiently small $\delta > 0$, $P(\beta(\delta), \delta)$ holds.

Proof. (i) For each $n < \omega$, find δ_n such that $P(2^{-n}, \delta_n)$ holds, $\delta_n \leq 2^{-n}$, and if $n > 0$, $\delta_n < \delta_{n-1}$. Let ω be defined in a piecewise way. Let $\alpha(0) = 0$, and if $x \geq 1$, let $\alpha(1) = \delta_1$. Otherwise find the unique $n < \omega$ such that $2^{-n-1} \leq x < 2^{-n}$. Let $\alpha(x) = \delta_{n+2} + 2^{n+1}(\delta_{n+1} - \delta_{n+2})(x - 2^{-n-1})$ (i.e. a linear interpolation between $(2^{-n-1}, \delta_{n+2})$ and $(2^{-n}, \delta_{n+1})$). (See Figure 11.)

By construction, α is a modulus. For any $\varepsilon > 0$, if $\varepsilon \geq 1$, then $\alpha(\varepsilon) = \delta_1$, so $P(\varepsilon, \alpha(\varepsilon))$ holds. Otherwise if $2^{-n-1} \leq \varepsilon < 2^{-n}$, then $\omega(x) \leq \delta_{n+1}$, so since $P(2^{-n-1}, \delta_{n+1})$ holds, we have that $P(\varepsilon, \alpha(x))$ holds as well.

(ii) Since $\delta_n < \delta_{n-1}$ for every $n < \omega$, α is invertible on $[0, 1]$. Let γ be the inverse of α on $[0, 1]$ extended by $\gamma(x) = 1$ wherever this is not defined. For any $\delta > 0$ with $\delta \leq \delta_1$, we have that $\alpha(\gamma(\delta)) = \delta$, so $P(\gamma(\delta), \alpha(\gamma(\delta))) \equiv P(\gamma(\delta), \delta)$ holds. Applying Lemma A.1.7 to γ gives the required β . (See Figure 12.) \square

Corollary A.1.9. *Let X and Y be pseudo-metric spaces, and let $f : X \rightarrow Y$ be a function. If*

- Y is bounded and
- for every $\varepsilon > 0$, there exists a $\delta > 0$ such that for any $x_0, x_1 \in X$, if $d^X(x_0, x_1) < \delta$, then $d^Y(f(x_0), f(x_1)) \leq \varepsilon$,

then f is uniformly continuous.

Proof. Assume that $f : X \rightarrow Y$ is a function between two metric spaces such that for every $\varepsilon > 0$, there exists a $\delta > 0$ such that for all $x_0, x_1 \in X$, if $d^X(x_0, x_1) \leq \delta$, then $d^Y(f(x_0), f(x_1)) \leq \varepsilon$. Let $P(\varepsilon, \delta)$ state that for all $x_0, x_1 \in X$, if $d^X(x_0, x_1) \leq \delta$, then $d^Y(f(x_0), f(x_1)) \leq \varepsilon$.

This implies that we have a modulus β such that for any sufficiently small $\delta > 0$, $P(\beta(\delta), \delta)$ holds. So we have for $d^X(x_0, x_1)$ sufficiently small that $d^Y(f(x_0), f(x_1)) \leq \beta(d^X(x_0, x_1))$. Let $r > 0$ be chosen such that if $d^X(x_0, x_1) < r$, then $d^Y(f(x_0), f(x_1)) \leq \beta(d^X(x_0, x_1))$. Now $\gamma(x) = \frac{x + \beta(x)}{r + \beta(r)} \max\{\text{diam}(Y), 1\}$ works as a modulus witnessing that f is uniformly continuous. \square

Note that without the boundedness assumption, Corollary A.1.9 may fail. The standard inclusion of \mathbb{N} into \mathbb{R} with \mathbb{N} given the discrete metric and \mathbb{R} given the standard metric is not uniformly continuous according to our Definition 0.3.9, but is according to the second condition in Corollary A.1.9.

A.2 Compact Hausdorff Spaces

It is a well known fact that in arbitrary topological spaces, the convergence of sequences is not enough to capture the entire topology. There are two generalizations of sequences that are commonly used, nets and filters.

Definition A.2.1. A *directed set* I is a partially ordered set such that for any $a, b \in I$ there exists a $c \in I$ such that $a \leq c$ and $b \leq c$.

In a topological space X , a *net* is a function, $f : I \rightarrow X$, from some directed set I into X . A net *converges* to a point $x \in X$ if for every open neighborhood $U \ni x$, there exists an $a \in I$ such that for all $b \geq a$, $f(b) \in U$.

A filter \mathcal{F} on X converges to a point $x \in X$ if for every open neighborhood $U \ni x$, $U \in \mathcal{F}$. ◁

A word of warning about nets: The term ‘subnet’ is not defined in the obvious way. $g : J \rightarrow X$ is a *subnet* of $f : I \rightarrow X$ if there exists an order preserving, cofinal map $h : J \rightarrow I$ such that $g = f \circ h$. In particular, it is not true that a sequence in a compact Hausdorff space always has a convergent sub-sequence.

Fact A.2.2. *A topological space is Hausdorff if and only if every net or filter converges to at most one point.*

For any topological space X the following are equivalent:

- X is compact.
- Every ultrafilter on X converges.
- Every net in X has a convergent subnet.

In particular, a space is compact Hausdorff if and only if every ultrafilter converges to a unique point.

Fact A.2.3 (Dini’s Theorem). *If X is a compact space and $\{f_i\}_{i < \omega}$ is a sequence of continuous functions $f_i : X \rightarrow \mathbb{R}$ such that*

- $f_i(x)$ is non-decreasing in i for each x ,
- $\{f_i\}_{i < \omega}$ converges pointwise to some $f : X \rightarrow \mathbb{R}$, and
- f is a continuous function,

then $\{f_i\}_{i < \omega}$ converges uniformly.

Fact A.2.4 (Stone–Weierstrass Theorem for Cubes). *For any ordinal k , any sequence $\{I_n\}_{n < k}$ of compact intervals, if we write A for the set of functions $f : \mathbb{R}^k \rightarrow \mathbb{R}$ generated by \max , \min , $+$, rational scaling, the constant 1, and the variables x_0, x_1, \dots , then the set $\{f \upharpoonright \prod_{n < k} I_n : f \in A\}$ is dense in the collection of continuous real valued functions on $\prod_{n < k} I_n$ under the uniform norm.*

Fact A.2.5 (Stone–Weierstrass Theorem). *For any compact Hausdorff space X , if A is a set of continuous functions closed under \max and \min such that*

for any pair of distinct points $x, y \in X$, any $r, s \in \mathbb{R}$ and any $\varepsilon > 0$, there is an $f \in A$ such that $|f(x) - r| < \varepsilon$ and $|f(y) - s| < \varepsilon$,

then A is dense in the set of continuous functions on X under the uniform norm.

Fact A.2.6 (Urysohn’s Lemma). *If X is a normal topological space (in particular, a compact Hausdorff space), and $F, G \subseteq X$ are disjoint closed sets, then there exists a continuous function $f : X \rightarrow \mathbb{R}$ such that $F \subseteq f^{-1}(0)$ and $G \subseteq f^{-1}(1)$.*

Corollary A.2.7. *If X is a compact Hausdorff space and $F, U \subseteq X$ are closed and open, respectively, such that $F \subseteq U$, then there exists a closed G_δ set G such that $F \subseteq \text{int } G \subseteq G \subseteq U$.*

Fact A.2.8 (Tietze Extension Theorem). *If X is a normal topological space (in particular, a compact Hausdorff space), $F \subseteq X$ is a closed set, and $f : F \rightarrow \mathbb{R}$ is a continuous function, then there exists a continuous function $g : X \rightarrow \mathbb{R}$ such that $f = g \upharpoonright F$.*

Fact A.2.9. *In a compact Hausdorff space X , a closed set F is the zeroset of a continuous function if and only if it is G_δ (i.e. a countable intersection of open sets).*

Fact A.2.10. *If X is a compact Hausdorff space, then any strictly coarser topology on X is not Hausdorff and any strictly finer topology on X is not compact.*

Fact A.2.11. *If X is a compact space, Y is a Hausdorff space, and $f : X \rightarrow Y$ is a continuous injection, then X is homeomorphic to its image under f . In particular, a continuous bijection from one compact Hausdorff space to another is always a homeomorphism.*

Recall that the *weight* of a topological space X , written $\text{wt}(X)$, is the least cardinality of a basis of open sets for X .

Fact A.2.12. *A compact Hausdorff space X is metrizable if and only if it is second countable (i.e. $\text{wt}(X) = \aleph_0$).*

Fact A.2.13. *Every compact Hausdorff space X is homeomorphic to a closed subset of $[0, 1]^{\text{wt}(X)}$.*

Fact A.2.14. *Every compact Hausdorff space X is homeomorphic to a quotient of a totally disconnected compact Hausdorff space Y with $\text{wt}(X) = \text{wt}(Y)$.*

If X is second countable and non-empty, Y can be taken to be 2^ω .

Fact A.2.15 (Union of Compactly Many Closed Sets is Closed, [Esc09]). *If X is a compact space and $F \subseteq X \times Y$ is a closed set, then the set $\{y \in Y : (\exists x \in X) \langle x, y \rangle \in F\}$ is closed in Y .*

Fact A.2.16. *If X is a compact Hausdorff space and U_0, \dots, U_{n-1} is a finite open cover of X , then there exists F_0, \dots, F_{n-1} , a sequence of G_δ closed sets, such that $F_i \subseteq U_i$ for each $i < \omega$, and $\text{int } F_0, \dots, \text{int } F_{n-1}$ is a cover of X .*

Proof. Perform the following operation on each U_i in turn: $G_i = X \setminus \bigcup_{k < n, k \neq i} U_k$ is a closed subset of U_i . Let F_i be a closed G_δ set such that $G_i \subseteq \text{int } F_i \subseteq F_i \subseteq U_i$. Replace U_i with $\text{int } F_i$. \square

Appendix B

Continuous Analogs of Classical Discrete Results

B.1 Some Results that Translate Directly

Here we collect classical discrete model theory results (other than the compactness theorem) that translate directly to continuous logic with minimal changes in the proof (at least conceptually).

Lemma B.1.1. *For any signature \mathcal{L} , every real \mathcal{L} -formula is logically equivalent to an unnested real \mathcal{L} -formula.*

Proof. First we will show that for any atomic real $\mathcal{L}(\bar{x})$ -formula $P\bar{t}$, there is an unnested real $\mathcal{L}(\bar{x})$ -formula $\varphi_{P\bar{t}}(\bar{x})$ logically equivalent to $P\bar{t}$ and with $I(\varphi_{P\bar{t}}) \subseteq I(P)$. We then get the full result by systematically replacing every instance of $P\bar{t}$ in any real \mathcal{L} -formula with $\varphi_{P\bar{t}}$. (Renaming bound variables as needed to avoid capture. The condition on $I(\varphi_{P\bar{t}})$ ensures that this procedure is well defined.)

Given an atomic \mathcal{L} -formula $P\bar{t}$, construct a tree $T \subseteq \omega^{<\omega}$ whose root is labeled with $P\bar{t}$. For each node σ , labeled with either $P\bar{s}$, $f\bar{s}$, or x . If it is one of the first two cases and \bar{s} is non-empty, add children $\sigma \frown i$ to T for each s_i and label each child with s_i . By the definition of term, this tree will always be well-founded. For each $\sigma \in T \setminus \{\emptyset\}$, if σ

is labeled with a non-variable term, let x_σ be a fresh variable symbol of the same sort as the term labeling σ . If σ is labeled with a variable v , let $x_\sigma = v$.

Now, assign \mathcal{L} -formula ψ_σ to each $\sigma \in T$ by the following:

- If σ is a leaf, it must be labeled by a variable or a constant, s , so let $\psi_\sigma = d(x_\sigma, s)$.
- If σ is not a leaf and is also not \emptyset and it is labeled with $f s_0 \dots s_n$, let

$$\psi_\sigma = \left[\inf_{x_{\sigma \frown 0} \dots x_{\sigma \frown n}} d(x_\sigma, f x_{\sigma \frown 0} \dots x_{\sigma \frown n}) + \omega_f (\max\{\psi_{\sigma \frown 0}, \dots, \psi_{\sigma \frown n}\}) \right]_0^{\text{db}(S(x_\sigma))}.$$

- If σ is not a leaf and is also not \emptyset and it is labeled with $f \bar{s}$ such that $a(f)$ is infinite, let

$$\psi_\sigma = \left[\inf_{x_{\sigma \frown 0} \dots} d(x_\sigma, f x_{\sigma \frown 0} \dots) + \omega_f \left(\sup_{i < \omega} \frac{\text{db}(S(x_{\sigma \frown 0}))}{2^i \text{db}(S(x_{\sigma \frown i}))} \psi_{\sigma \frown i} \right) \right]_0^{\text{db}(S(x_\sigma))}.$$

- If $a(P)$ is finite, let

$$\psi_\emptyset = \inf_{x_0 \dots x_n} P(x_0, \dots, x_n) + \omega_P (\max\{\psi_0, \dots, \psi_n\}),$$

otherwise let

$$\psi_\emptyset = \inf_{x_0 \dots} P(x_0, \dots) + \omega_P \left(\sup_{i < \omega} \frac{\text{db}(S(x_0))}{2^i \text{db}(S(x_i))} \psi_i \right).$$

Finally, let $\varphi_{P\bar{i}} = [\psi_\emptyset]_{\inf I(P)}^{\sup I(P)}$. Note that $\varphi_{P\bar{i}}$ is unnested by construction.

Now we want to argue that for any $\sigma \in T \setminus \{\emptyset\}$, for any \mathcal{L} -structure \mathfrak{M} and any variable assignment ι whose domain contains all the relevant variables, $\psi_\sigma^{\mathfrak{M}}(\iota) =$

$d(\iota(x_\sigma), s^{\mathfrak{M}}(\iota))$, where s is the term labeling σ . If σ is a leaf then this is obvious. Assuming we've shown this for all of the children of some node σ which is labeled by $f\bar{s}$, we have that for any $\bar{a} \in \mathfrak{M}$ of the same sorts as $x_{\sigma \frown 0} x_{\sigma \frown 1} \dots$, if we let χ_σ be the formula inside the quantifier of the definition of ψ_σ , then $\chi_\sigma^{\mathfrak{M}}(\iota[\frac{\bar{a}}{x_{\sigma \frown 0} \dots}]) \geq d(\iota(x_\sigma), f^{\mathfrak{M}}(\iota[\frac{\bar{a}}{x_{\sigma \frown 0} \dots}]))$, because f is required to be ω_f -uniformly continuous. So since we clearly have $\psi_\sigma^{\mathfrak{M}}(\iota) \leq d(\iota(x_\sigma), f^{\mathfrak{M}}(\iota))$, by the induction hypothesis, we in fact have $\psi_\sigma^{\mathfrak{M}}(\iota) = d(\iota(x_\sigma), f^{\mathfrak{M}}(\iota))$, as required.

Finally, the same argument with P shows that in any \mathcal{L} -structure \mathfrak{M} and for any variable assignment ι whose domain is large enough, $\psi_\emptyset^{\mathfrak{M}}(\iota) = \varphi_{P\bar{t}}^{\mathfrak{M}}(\iota) = (P\bar{t})^{\mathfrak{M}}(\iota)$, as required, so $\varphi_{P\bar{t}}$ and $P\bar{t}$ are logically equivalent and have the same free variables. We also clearly have $I(\varphi_{P\bar{t}}) \subseteq I(P)$, by construction. \square

Note that it is not generally true that every restricted real formula is logically equivalent to an unnested restricted real formula, unless every function symbol is Lipschitz.

Proposition B.1.2 (The Löwenheim-Skolem Theorem). *Let \mathfrak{M} be a pre-structure. For any $A \subseteq \mathfrak{M}$ there is an elementary sub-pre-structure $\mathfrak{N}_0 \preceq \mathfrak{M}$ such that $N_0 \supseteq A$ and $|\mathfrak{N}_0| \leq \aleph_0 + |\mathcal{L}| + |A|$.*

If \mathfrak{M} is a structure, then there exists an elementary sub-structure $\mathfrak{N} \preceq \mathfrak{M}$ such that $N \supseteq A$ and $\#^{\text{dc}}\mathfrak{N} \leq \aleph_0 + |\mathcal{L}| + |A|$.

Proof. Follow the standard proof, but use open \mathcal{L} -formulas. The second part follows from the first. \square

Omitting types is a slight exception to the stated purposed of this appendix in that the resulting statement of the theorem is different in an important way, but the proof is essentially the same.

Proposition B.1.3 (Omitting Types). *Let \mathcal{L} be a countable signature, and let T be an \mathcal{L} -theory.*

- *Let $\{\Sigma_i(x)\}_{i < \omega}$ be a sequence of partial types such that for each $i < \omega$, $\text{int}_{S_x(T)}[\Sigma] = \emptyset$. There exists a countable pre-model $\mathfrak{M} \models T$ that omits each Σ_i .*
- *Let $\{p_i(x)\}_{i < \omega}$ be a sequence of complete types such that for each $i < \omega$, p_i is not d -atomic in $S_x(T)$. There exists a separable model $\mathfrak{N} \models T$ that omits each p_i .*

Proof. The proof of the first part is the same as the classical proof.

The second part follows from the first part and the following observation: $p_i(x)$ is not d -atomic if and only if there is an $\varepsilon > 0$ such that $B_{\leq \varepsilon}(p_i)$ (which is a partial type, i.e. a closed set of types) has empty topological interior. Finally the types realized in the completion of a pre-model are in the metric closure of the types realized in the pre-model itself, so p_i must be omitted in the completion of the pre-model from the first part. □

Lemma B.1.4. *For any type space $S_n(T)$, if d -atomic types are dense in $S_n(T)$, then for any definable set $D \subseteq S_n(T)$, atomic types are dense in D as well.*

Proof. Fix an open set U such that $D \cap U$ is non-empty. Find an open set V such that $\text{cl } V \subseteq U$ and $D \cap V$ is non-empty. Find $\varepsilon_0 > 0$ small enough that $(\text{cl } V)^{\leq 3\varepsilon_0} \subseteq U$.

The set $D^{< \varepsilon} \cap V$ is non-empty by construction. Let p_0 be a d -atomic type in $D^{< \varepsilon} \cap V$. At any stage $k < \omega$, given d -atomic $p_k \in D^{< 2^{-k}\varepsilon}$, consider the set $D^{< 2^{-k-1}\varepsilon} \cap B_{< 2^{-k}\varepsilon}(p_k)$. Since this is a superset of $D \cap B_{< 2^{-k}\varepsilon}(p_k)$, this is non-empty. Since p_k is d -atomic, this is an open set. Therefore we can find a d -atomic type $p_{k+1} \in D^{< 2^{-k-1}\varepsilon} \cap B_{< 2^{-k}\varepsilon}(p_k)$.

The sequence $\{p_k\}_{k < \omega}$ is by construction a Cauchy sequence. Let p_ω be its limit. $d(p_\omega, D) = 0$ by construction, so $p_\omega \in D$. Also by construction, we have that $d(p_0, p_\omega) \leq$

$\sum_{k<\omega} 2^{-k}\varepsilon = 2\varepsilon < 3\varepsilon$, and so since $p_0 \in V$ we have that $p_\omega \in U$. Since the limit of any Cauchy sequence of d -atomic types is d -atomic, we have that p_ω is d -atomic.

Since we can do this for every open U such that $D \cap U$ is non-empty, we have that d -atomic types are dense in D . \square

Proposition B.1.5. *For a countable complete theory T , the following are equivalent:*

(i) *T has an atomic model (i.e. every n -type realized for every $n < \omega$ is d -atomic).*

(ii) *T has a prime model.*

(iii) *d -atomic types are dense in $S_n(T)$ for every $n < \omega$.*

Proof. (i) \Rightarrow (iii). Obvious.

(iii) \Rightarrow (ii). If d -atomic types are dense in $S_n(T)$ for every $n < \omega$, then by Lemma B.1.4 we have that if $p(\bar{x})$ is a d -atomic type over finitely many variables, then for any open formula $U(\bar{x}, y)$ such that $p(\bar{x}) \cup \{U(\bar{x}, y)\}$ is consistent, there is a type $q(\bar{x}, y)$ extending p such that q is d -atomic (in $S_{\bar{x}y}(T)$). The same argument as in discrete logic allows us to build a countable pre-model \mathfrak{M} with an enumeration $\{a_i\}_{i<\omega}$ such that for each $n < \omega$, $\text{tp}(a_0 \dots a_{n-1})$ is d -atomic (in $S_n(T)$). Lemma 2.5.6 now implies that \mathfrak{M} , and therefore $\overline{\mathfrak{M}}$, elementarily embeds into any model of T . Therefore $\overline{\mathfrak{M}}$ is prime, and T has a prime model.

(ii) \Rightarrow (i). The prime model must be separable. By the omitting types theorem, every type realized in it must be d -atomic, so it is atomic. \square

Proposition B.1.6 (The Engeler–Ryll–Nardzewski–Svenonius Theorem). *Let T be a complete theory with non-compact models in a countable signature. The following are equivalent:*

- T is ω -categorical (any two separable models are isomorphic).
- For every n , $S_n(T)$ is metrically compact.
- For every n , every $p \in S_n(T)$ is d -atomic.

Proof. Theorem 6.6.3 is a more general version of this. \square

Proposition B.1.7. *Let $D(\bar{x})$ be a \mathfrak{C} -definable set (where \mathfrak{C} is the monster). Fix a small set A . If $\{D(\mathfrak{C})\}$ is invariant under $\text{Aut}(\mathfrak{C}/A)$, then there is an A -definable set $E(\bar{x})$ such that $\llbracket D(\bar{x}) \rrbracket_{\mathfrak{C}} = \llbracket E(\bar{x}) \rrbracket_{\mathfrak{C}}$.*

Proof. Since $\llbracket D(\bar{x}) \rrbracket$ is A -invariant, we clearly have that $\llbracket D(\bar{x}) \rrbracket^{<\varepsilon}$ is A -invariant for every $\varepsilon > 0$ as well. A typical argument shows that if $X \subseteq S_{\bar{x}}(\mathfrak{C})$ is A -invariant and $f : S_{\bar{x}}(\mathfrak{C}) \rightarrow S_{\bar{x}}(A)$ is the natural restriction map, then there is some $Y \subseteq S_{\bar{x}}(A)$ such that $X = f^{-1}(Y)$. Since $f(S_{\bar{x}}(\mathfrak{C}) \setminus \llbracket D(\bar{x}) \rrbracket^{<\varepsilon})$ is closed by compactness, this implies that $f(\llbracket D(\bar{x}) \rrbracket^{<\varepsilon})$ is open for every $\varepsilon > 0$. Let $\llbracket E(\bar{x}) \rrbracket = f(\llbracket D(\bar{x}) \rrbracket)$. Since f is 1-Lipschitz, we have that $\llbracket E(\bar{x}) \rrbracket \subseteq f(\llbracket D(\bar{x}) \rrbracket^{<\varepsilon}) \subseteq \text{int}\llbracket E(\bar{x}) \rrbracket^{<\varepsilon}$ for every $\varepsilon > 0$, so $E(\bar{x})$ is definable. We have $\llbracket D(\bar{x}) \rrbracket_{\mathfrak{C}} = \llbracket E(\bar{x}) \rrbracket_{\mathfrak{C}}$ by construction. \square

Proposition B.1.8 (Craig Interpolation Theorem). *Let \mathcal{L}_0 be a signature, and let \mathcal{L}_1 and \mathcal{L}_2 be two expansions of it such that $\mathcal{L}_1 \cap \mathcal{L}_2 = \mathcal{L}_0$. Furthermore, let V_1 and V_2 be sets of variables such that $V_1 \cap V_2 = V_0$. Let F be a closed $\mathcal{L}_1(V_1)$ -formula, and let U be an open $\mathcal{L}_2(V_2)$ -formula such that $F \models U$ (i.e. for any $\mathcal{L}_1 \cup \mathcal{L}_2$ -structure \mathfrak{M} and $\iota : V_1 \cup V_2 \rightarrow M$ such that $\mathfrak{M} \models F(\iota)$, we have $\mathfrak{M} \models U(\iota)$). There exists an open $\mathcal{L}_0(V_0)$ -formula W and a closed $\mathcal{L}_0(V_0)$ -formula G such that $F \models W \models G \models U$.*

Proof. Let $r_1 : S_{V_1}(\mathcal{L}_1) \rightarrow S_{V_0}(\mathcal{L}_0)$ and $r_2 : S_{V_2}(\mathcal{L}_2) \rightarrow S_{V_0}(\mathcal{L}_0)$ be the compositions of the reduct maps and the projection maps. Assume that $r_1(\llbracket F \rrbracket) \cap r_2(\llbracket \neg U \rrbracket)$ is non-empty,

and let p_0 be a complete $\mathcal{L}_0(V_0)$ -type in the intersection. Let p_1 be an extension of p_0 to a complete $\mathcal{L}_1(V_1)$ -type, and let p_2 be an extension of p_0 to a complete $\mathcal{L}_2(V_2)$ -type. Pick κ such that $\text{cf}(\kappa)$ is much larger than $|\mathcal{L}_1| + |\mathcal{L}_2| + |V_1| + |V_2|$. For $i < 2$, let \mathfrak{M}_i be an \mathcal{L}_i -structure that is a κ -special model of the closed sentences in p_i , and let ι_i be a V_i -assignment such that $\text{tp}(\iota_i) = p_i$. By specialness, there is an isomorphism $f : \mathfrak{M}_1 \upharpoonright \mathcal{L}_0 \rightarrow \mathfrak{M}_2 \upharpoonright \mathcal{L}_0$ such that $f \circ (\iota_1 \upharpoonright V_0) = \iota_2 \upharpoonright V_0$. Using this isomorphism we can extend \mathfrak{M}_1 to an $\mathcal{L}_1 \cup \mathcal{L}_2$ -structure \mathfrak{M}_{12} . If we let ι_{12} be the $V_1 \cup V_2$ -assignment for \mathfrak{M}_{12} extending ι_1 and ι_2 , we have that $(\mathfrak{M}_{12}, \iota_{12})$ contradicts that $F \models U$.

Therefore $r_1(\llbracket F \rrbracket)$ and $r_2(\llbracket \neg U \rrbracket)$ are disjoint. So there exists a real $\mathcal{L}_0(V_0)$ -formula φ such that $r_1(\llbracket F \rrbracket) \subseteq \llbracket \varphi = 0 \rrbracket$ and $r_2(\llbracket \neg U \rrbracket) \subseteq \llbracket \varphi = 1 \rrbracket$. Now $W = (\varphi < \frac{1}{2})$ and $G = (\varphi \leq \frac{1}{2})$ are the required formulas. \square

Proposition B.1.9 (Beth Definability Theorem). *Let T be a (possibly incomplete) \mathcal{L} -theory. Suppose that P is a predicate symbol not in \mathcal{L} and that $T' \supseteq T$ is an $\mathcal{L} \cup \{P\}$ -theory such that for every model $\mathfrak{M} \models T$, there is a unique expansion of \mathfrak{M} to a model of T' , then there exists an \mathcal{L} -formula $\varphi(\bar{x})$, with \bar{x} a tuple of variables of the same sorts as $a(P)$, such that $T' \models (\forall \bar{x})\varphi(\bar{x}) = P(\bar{x})$.*

Proof. Let P_0 and P_1 be two predicate symbols with the same arity and modulus of continuity as P , and let T'_0 and T'_1 be T' with each instance of P replaced with P_0 and P_1 , respectively. By assumption, we have that for any $r < s$, the theory $T'_0 \cup T'_1 \cup \{\exists \bar{x} P_0(\bar{x}) \leq r \wedge P_1(\bar{x}) \geq s\}$ is inconsistent. Therefore by compactness, there must be closed sentences F_0 and G_1 such that $F_0 \in T'_0$, $G_1 \in T'_1$, and $\{F_0, G_1, P_0(\bar{x}) \leq r, P_1(\bar{x}) \geq s\}$ is inconsistent. If we let F_1 be F_0 with all instances of P_0 replaced with P_1 , and likewise let G_0 be G_1 with all instance of P_1 replaced with P_0 , then we have that $T'_i \models F_i \wedge G_i$ and that $F_i \wedge G_i$ is an

\mathcal{L}_i -sentence for both $i < 2$. Now we have that $\{F_0 \wedge G_0, F_1 \wedge G_1, P_0(\bar{x}) \leq r, P_1(\bar{x}) \geq s\}$ is still inconsistent, so in particular, $F_0 \wedge G_0 \wedge P_0(\bar{x}) \leq r \models \neg(F_1 \wedge G_1) \vee P_1(\bar{x}) < s$. Note that if F and G are F_0 and G_0 with P_0 replaced by P , then we have that $T' \models F \wedge G$ as well.

By the Craig interpolation theorem, there is an open $\mathcal{L}(\bar{x})$ -formula $O_{r,s}(\bar{x})$ and a closed $\mathcal{L}(\bar{x})$ -formula $C_{r,s}(\bar{x})$ such that

$$F_0 \wedge G_0 \wedge P_0(\bar{x}) \leq r \models O_{r,s}(\bar{x}) \models C_{r,s}(\bar{x}) \models (F_1 \wedge G_1) \rightarrow P_1(\bar{x}) < s.$$

Therefore, by symmetry, these statements apply with F_i, G_i, P_i replaced with F, G, P and we have that for any tuple \bar{a} in a model of T' if $P(\bar{a}) < s$, then $O_{r,s}(\bar{a})$ holds for any $r < s$ with $P(\bar{x}) < r < s$. On the other hand, for any \bar{b} that satisfy $O_{r,s}(\bar{b})$, we automatically have that $P(\bar{b}) < s$. Therefore $P(\bar{x}) < s$ is logically equivalent to $\bigwedge_{k < \omega} O_{s-2^{-k}, s}(\bar{x})$. A similar argument gives that for any r , $P(\bar{x}) > r$ is logically equivalent to $\bigwedge_{k < \omega} \neg C_{s, s+2^{-k}}(\bar{x})$. Therefore the value of $P(\bar{c})$ for a given \bar{c} is entirely determined by $\text{tp}(\bar{c}) \upharpoonright \mathcal{L}$, and furthermore, the function mapping types in $S_n(T)$ to the corresponding value of P is continuous. Therefore there is an \mathcal{L} -formula $\varphi(\bar{x})$ such that $T' \models (\forall \bar{x}) \varphi(\bar{x}) = P(\bar{x})$. \square

B.2 Two Proofs of Morley's Theorem

In this section we will present two different proofs of Morley's theorem in continuous logic: a transcription of the proof in [BY05] to the more specific formalism of continuous logic and an adaptation of the classical proof in [CK90]. There is a fair amount of overlap

between these two proofs, which is presented in the first subsection.

B.2.1 Common Material

Definition B.2.1. A theory T is *totally transcendental* (or *t.t.*) if for every $\varepsilon > 0$ there does not exist an ε -perfect tree of closed formulas, where an ε -perfect tree is a family $\{F_\sigma\}_{\sigma \in 2^{<\omega}}$ of closed formulas such that $F_{\sigma \frown i} \subseteq \text{int}(F_\sigma) \neq \emptyset$ and $d_I(F_{\sigma \frown 0}, F_{\sigma \frown 1}) := \inf\{d(p, q) : p \in F_{\sigma \frown 0}, q \in F_{\sigma \frown 1}\} > \varepsilon$. \triangleleft

Definition B.2.2. A theory T is κ stable if for every A with $|A| \leq \kappa$, we have that $\#^{\text{dc}}S_1(A) \leq \kappa$. \triangleleft

Proposition B.2.3. *If T is ω -stable, then T is t.t.*

Proof. Given an ε -perfect tree, $\{F_\sigma\}_{\sigma \in 2^{<\omega}}$, we can construct a ε -perfect tree of closed formulas, $\{\llbracket \varphi_\sigma \leq \frac{1}{2} \rrbracket\}_{\sigma \in 2^{<\omega}}$, with φ_σ restricted formulas.

To see this let $\varphi_\emptyset(x) = d(x, x)$. Then for each σ , we can find $\varphi_{\sigma \frown 0}$ such that $F_{\sigma \frown 0} \subseteq \llbracket \varphi_{\sigma \frown 0} < \frac{1}{2} \rrbracket$ and $\llbracket \varphi_{\sigma \frown 0} \leq \frac{1}{2} \rrbracket \subseteq \text{int}(F_\sigma) \setminus F_{\sigma \frown 1}^{\leq \varepsilon}$, by compactness, and then $d_I(\llbracket \varphi_{\sigma \frown 0} \leq \frac{1}{2} \rrbracket, F_{\sigma \frown 1}) > \varepsilon$, so we can find $\varphi_{\sigma \frown 1}$ such that $F_{\sigma \frown 1} \subseteq \llbracket \varphi_{\sigma \frown 1} < \frac{1}{2} \rrbracket$ and $\llbracket \varphi_{\sigma \frown 1} \leq \frac{1}{2} \rrbracket \subseteq \text{int}(F_\sigma) \setminus \llbracket \varphi_{\sigma \frown 0} \leq \frac{1}{2} \rrbracket^{\leq \varepsilon}$. By induction this gives the required ε -perfect tree.

Since each φ_σ is a restricted formula, the whole tree involves only a countable set of formulas A . $\#^{\text{dc}}S_1(A) \geq 2^\omega$, because for each path $\alpha \in 2^\omega$, the set $G_\alpha = \bigcap_{i < \omega} \llbracket \varphi_{\alpha \upharpoonright i} \leq \frac{1}{2} \rrbracket$ is non-empty by compactness, and for any $\alpha, \beta \in 2^\omega$ with $\alpha \neq \beta$, $d_I(G_\alpha, G_\beta) > \varepsilon$. So by picking a type in each path we witness that $\#_{>\varepsilon}^{\text{ent}}S_1(A) \geq 2^\omega$. Hence the density character must be at least as large. \square

Proposition B.2.4. *If T is t.t., then for every $\kappa \geq |T|$, T is κ -stable.*

Proof. Assume that for some $\kappa \geq |T|$, T is not κ -stable. Let A be a set of parameters such that $|A| \leq \kappa$ and $\#^{dc}S_1(A) > \kappa$. Since $\kappa \geq |T|$, $S_1(A)$ has a topological base of cardinality κ . Let $\varepsilon > 0$ be such that the metric entropy satisfies $\#_{>\varepsilon}^{\text{ent}}S_1(A) \geq \kappa^+$, which must exist because κ^+ is a regular cardinal. Furthermore note that there must be a maximal ($> \varepsilon$)-separated set with cardinality κ^+ , again because it is a regular cardinal.

Let $X = S_1(A) \setminus \{U \subseteq S_1(A) : U \text{ open, } \#_{>\varepsilon}^{\text{ent}}U \leq \kappa\}$. Clearly X is closed. It must be non-empty because otherwise $S_1(A)$ would be covered by open sets with metric entropy $\leq \kappa$. By compactness there would be a finite sub-cover, so by the pigeonhole principle one of these sets must contain a ($> \varepsilon$)-separated set with cardinality κ^+ , which is a contradiction.

Now if V is any open-in- X set, it must satisfy $\#_{>\varepsilon}^{\text{ent}}V \geq \kappa^+$. Otherwise if we extend it to an open set U such that $U \cap X = V$, then U would be the union of a set V with $\#_{>\varepsilon}^{\text{ent}}V \leq \kappa$ and κ -many open sets $\{U_i\}_{i < \kappa}$ satisfying $\#_{>\varepsilon}^{\text{ent}}U_i \leq \kappa$, implying that U is an open set that should have been removed when we made X , which is a contradiction.

Now construct a ε -perfect tree by induction. Let $F_\sigma = X$. For each $\sigma \in 2^{<\omega}$, by induction $\text{int}_X F_\sigma \neq \emptyset$, so since $\#_{>\varepsilon}^{\text{ent}}\text{int}_X F_\sigma \geq \kappa^+ > 2$, we can find $p, q \in \text{int}_X F_\sigma$ such that $d(p, q) > \varepsilon$. Since $p \notin B_{\leq \varepsilon}(q)$, it has a closed neighborhood $F_{\sigma \frown 0}$ disjoint from $B_{\leq \varepsilon}(q)$. Since $q \notin F_{\sigma \frown 0}^{\leq \varepsilon}$, it has a closed neighborhood $F_{\sigma \frown 1}$ disjoint from $F_{\sigma \frown 0}^{\leq \varepsilon}$.

So by induction you get a ε -perfect tree, contradicting that T is t.t. □

Definition B.2.5. A model $\mathfrak{M} \models T$ is κ -saturated if for every $A \subseteq \mathfrak{M}$ with $|A| \leq \kappa$, \mathfrak{M} realizes every type in $S_1(A)$. It is saturated if it is $|\mathfrak{M}|$ -saturated. ◁

Proposition B.2.6. *If T is κ -stable, then for every regular cardinal λ with $|T| + \aleph_1 \leq \lambda \leq \kappa$, T has a model of density character κ which is λ -saturated.*

Proof. Proof is the same as the proof in discrete logic. There is some subtlety with ω -saturated structures (in particular ω -stability only guarantees the existence of ‘approximately ω -saturated’ separable structures), but it washes out for uncountable cardinalities. \square

Proposition B.2.7. *Two saturated models of the same density character are isomorphic.*

Proof. Same as in discrete logic. Do back-and-forth using dense sub-pre-structures (you won’t get a bijection between the dense sub-pre-structures, rather you’ll construct a metric on their union that makes them both dense, giving an isomorphism between the structures). \square

Proposition B.2.8. *Let T be a complete countable theory with non-compact models. For any $\kappa \geq \omega$, T has a model with density character κ that only realizes separably many types over any countable set.*

Proof. Take any model \mathfrak{A} of T . Let $Q \subseteq A$ be an infinite ($\geq \varepsilon$)-separated set. Add a distance predicate for Q to the language. Recast \mathfrak{A} as a discrete structure \mathfrak{A}^{dis} (the particulars of how you do this don’t matter, all that matters is that the language is countable). Skolemize the theory, and find an indiscernible sequence among the realizations of Q . Run the discrete Ehrenfeucht-Mostowski argument over the indiscernible sequence in Q to get a $\mathfrak{B}^{dis} \equiv \mathfrak{A}$ which only realizes countably many types over any countable set. Reinterpret \mathfrak{B}^{dis} as a metric pre-structure \mathfrak{B}_0 . One can check the following;

- $\mathfrak{B}_0 \equiv \mathfrak{A}$.
- If for any countable $X \subseteq B_0$ and any $a, b \in B_0$, $a \equiv_X b$ as elements of the discrete structure, then $a \equiv_X b$ as elements of the metric pre-structure.

Therefore \mathfrak{B}_0 is a pre-model elementarily equivalent to \mathfrak{A} which realizes only countably many types over countable sets. Since $|\mathfrak{B}_0| \leq \kappa$, we have that $\#^{dc}\mathfrak{B}_0 \leq \kappa$, but since \mathfrak{B}_0 also has a $(\geq \varepsilon)$ -separated set, Q , of size κ , we have that $\kappa \leq \#_{>\varepsilon}^{\text{ent}}\mathfrak{B}_0 \leq \#^{dc}\mathfrak{B}_0 \leq \kappa$, so $\#^{dc}\mathfrak{B}_0 = \kappa$.

Taking the completion $\mathfrak{B} = \overline{\mathfrak{B}_0}$ gives the required metric structure. \square

Proposition B.2.9. *If T is a complete countable theory with non-compact models such that T is κ -categorical for some $\kappa \geq \aleph_1$, then T is ω -stable.*

Proof. We know that T has a model of density character κ that only realizes separably many types over countable sets, so since there is only one model of density character κ and any set of types of density character $\leq \kappa$ can be realized in a model with density character κ , we must have $\#^{dc}S_1(A) \leq \omega$ for every A with $|A| \leq \omega$. \square

Corollary B.2.10. *For $\kappa \geq \aleph_1$, a countable theory T is κ -categorical if and only if every model of cardinality κ is saturated.*

Definition B.2.11. Let $x \in X$, with (X, d) a topometric space. x is (d, ε) -atomic if $x \in \text{int}B_{\leq \varepsilon}(x)$. \triangleleft

Note that x is d -atomic if it is (d, ε) -atomic for every $\varepsilon > 0$.

Proposition B.2.12. *Let T be a t.t. theory. Fix A , $n < \omega$, and $X \subseteq S_n(A)$ a closed subset.*

(i) *For every $\varepsilon > 0$, (d, ε) -atomic-in- X points are dense in X .*

(ii) *d -atomic points are dense in X .*

Proof. (i). Assume that for some $\varepsilon > 0$, (d, ε) -atomic-in- X points are not dense in X . Then we can find a non-empty open-in- X set $U \subseteq X$ such that no point in U is (d, ε) -atomic-in- X . Furthermore this is true for any open-in- X set $V \subseteq U$.

Claim: This implies that $\#_{>\varepsilon}^{\text{ent}} V \geq \omega$ for any such V .

Proof of claim: Assume that $\#_{>\varepsilon}^{\text{ent}} V < \omega$. Then let $\{a_i\}_{i < k}$ be a maximal $(> \varepsilon)$ -separated subset of V . By maximality, $\bigcup_{i < k} B_{\leq \varepsilon}(a_i) \supseteq V$, so for each a_i , we have $a_i \in V \setminus \bigcup_{j \neq i} B_{\leq \varepsilon}(a_j) \subseteq B_{\leq \varepsilon}(a_i)$, which implies that $a_i \in \text{int}_X B_{\leq \varepsilon}(a_i)$, so a_i is (d, ε) -atomic-in- X , contradicting our assumption.

Now by the same argument as in one of the previous propositions, since we have $\#_{>\varepsilon}^{\text{ent}} V \geq \omega > 2$ for every non-empty open-in- X set $V \subseteq U$, we can form a ε -perfect tree, contradicting that T is t.t.

(ii). For each $k < \omega$, let $U_k = \bigcup \{\text{int}_X B_{\leq \varepsilon}(a) : a \text{ } (d, \varepsilon)\text{-atomic-in-}X\}$. This is a topologically open dense subset of X . By the Baire category theorem we have that $Q = \bigcap_{k < \omega} U_k$ is non-empty and dense. For any $a \in Q$, for every k , there exists an a_k such that $a \in \text{int}_X B_{\leq 2^{-k}}(a_k) \subseteq B_{\leq 2^{-k+1}}(a)$, so $a \in \text{int}_X B_{\leq 2^{-k+1}}(a)$ for every $k < \omega$ and a is d -atomic.

Therefore d -atomic-in- X points are dense in X . □

Note that the proof doesn't really require that T be t.t., just that $S_n(A)$ contain no ε -perfect trees, which is a topometric property of $S_n(A)$. In particular the same proof works in small theories for the type spaces $S_n(T)$. (Although note that in continuous logic a small theory may fail to be small after the addition of finitely many new constants.)

Proposition B.2.13. *If T is t.t., then it has atomic models over any set, i.e. for any set A there is a model $\mathfrak{M} \supseteq A$ such that for any $\mathfrak{N} \supseteq A$, $\mathfrak{M} \preceq \mathfrak{N}$ with an A -elementary*

embedding.

Proof. Proceed in the same manner as the normal proof: construct a pre-structure by transfinite induction. The only tricky part is that a type can be d -atomic over a countable set of parameters, but it is true that if $\text{tp}(b/A_{c_{<\omega}})$ is d -atomic and for each $i < \omega$, $\text{tp}(c_i/A_{c_{<i}})$ is d -atomic, then $\text{tp}(b/A)$ is d -atomic. The easiest way to see this is in terms of the characterization that a type is d -atomic if and only if it's realized in every model containing the parameter set (or rather the countable subset of parameters over which it is d -atomic, but since the construction of an atomic model is well-ordered, each element is d -atomic over a countable set of the previous elements of the construction). Also note that a metric limit of d -atomic types is d -atomic. \square

B.2.2 Ben Yaacov's Proof after Morley

This section contains a transcription of Ben Yaacov's proof of Morley's theorem [BY05], originally in the context of CATs, to the less general framework of continuous logic.

ε -Cantor-Bendixson Analysis

Clearly you could define analogous Cantor-Bendixson ranks to the one defined here for any given counting notion and there are at least nine distinct natural counting notions, but, to avoid clutter and because we really only need one notion of rank, we're not going to define Cantor-Bendixson rank for topometric spaces in full generality (note that what we define here as ε -Cantor-Bendixson rank is the same as what is called (f, ε) -Cantor-Bendixson rank in [BY08c]:

Definition B.2.14. (i) For a topometric space (X, d) , ε -Cantor-Bendixson derivative

X'_ε is given by $X \setminus \bigcup\{U : \text{open and } \#_\varepsilon^{\text{ent}}U < \omega\}$.

(ii) The ε -Cantor-Bendixson sequence, $X_\varepsilon^{(\alpha)}$ is given by:

- $X_\varepsilon^{(0)} = X$,
- $X_\varepsilon^{(\lambda)} = \bigcap_{\beta < \lambda} X_\varepsilon^{(\beta)}$ for λ a limit,
- $X_\varepsilon^{(\alpha+1)} = \left(X_\varepsilon^{(\alpha)}\right)'_\varepsilon$, and
- $X_\varepsilon^{(\infty)} = \bigcap_{\alpha \in \text{Ord}} X_\varepsilon^{(\alpha)}$.

(iii) For any point $x \in X$, the ε -Cantor-Bendixson rank of x , $CB_\varepsilon(x)$ is $\sup\{\alpha : x \in X_\varepsilon^{(\alpha)}\}$.

(iv) For any set $A \subset X$, the ε -Cantor-Bendixson rank of A , $CB_\varepsilon(A)$ is $\sup\{CB_\varepsilon(x) : x \in A\}$ or ∞ if no such supremum exists. By convention $CB_\varepsilon(\emptyset) = -1$.

(v) For a set $A \subset X$ the ε -Cantor-Bendixson degree of A , $CBd_\varepsilon(A)$ is $\#_\varepsilon^{\text{ent}}\left(A \cap X_\varepsilon^{(CB_\varepsilon(A))}\right)$.

◁

Some things to note:

- For any point $x \in X$, $CB_\varepsilon(x) = \min\{CB_\varepsilon(U) : x \in U, U \text{ open}\}$, but also for any basis of open sets \mathcal{B} , $CB_\varepsilon(x) = \min\{CB_\varepsilon(U) : x \in U \in \mathcal{B}\}$.
- $X_\varepsilon^{(\alpha)}$ is always closed.
- If $\#_\varepsilon^{\text{ent}}U$ is finite then $\#_\varepsilon^{\text{ent}}V$ will be finite for any $V \subseteq U$ in some basis of open sets, therefore $X_\varepsilon^{(\alpha)}$ always stabilizes before κ^+ where κ is the minimum cardinality of a basis of open sets in X .

- If A is compact then $CBd_\varepsilon(A)$ is always finite and positive: Positive because $A \cap X_\varepsilon^{(\alpha)}$ form a descending sequence of closed sets and by compactness its intersection, and in particular $A \cap X_\varepsilon^{(CB_\varepsilon(A))}$, is always non-empty. Finite because it can be covered by open sets with finite $\#_\varepsilon^{\text{ent}}$ and by compactness there is a finite sub-cover, so $\#_\varepsilon^{\text{ent}} \left(A \cap X_\varepsilon^{(CB_\varepsilon(A))} \right)$ must be finite as well, otherwise there would be an infinite ($> \varepsilon$)-separated set and by the pigeonhole principle one of the open sets in the finite sub-cover would contain infinitely many of these elements, contradicting the finiteness of $\#_\varepsilon^{\text{ent}}$ for that set.

The following is essentially one direction of [BY08c, Prop. 3.17]. The converse is true as well, but we won't need it here.

Fact B.2.15. *For a compact topometric space (X, d) , if X contains no ε -perfect tree, then every $x \in X$ has ordinal ε -Cantor-Bendixson rank.*

Proof. If X has unranked points for some $\varepsilon > 0$, then $X_\varepsilon^{(\infty)} \neq \emptyset$. $X_\varepsilon^{(\infty)}$ has the property that any open $U \subset X_\varepsilon^{(\infty)}$ (open in the subspace topology) has $\#_\varepsilon^{\text{ent}} U \geq \omega$. Therefore we can always find two points $x, y \in U$ with $d(x, y) > \varepsilon$. By the same argument as in the proof of Proposition B.2.4 we can get a ε -perfect tree in $X_\varepsilon^{(\infty)}$. \square

Lemma B.2.16. (i) *If every $x \in X$ for some locally compact topometric space (X, d) has ordinal ε -Cantor-Bendixson rank, then (d, ε) -isolated points (i.e. $x \in B_{\leq \varepsilon}(x)^\circ$) are dense in X .*

(ii) *If every $x \in X$ has ordinal ε -Cantor-Bendixson rank for every $\varepsilon > 0$, then d -isolated points (i.e. $x \in B_{\leq \varepsilon}(x)^\circ$ for every $\varepsilon > 0$) are dense in X .*

Proof. (i) If some non-empty open set $U \subset X$ remains unchanged in $X'_{>\varepsilon}$, then it will be stable for the entire ε -Cantor-Bendixson sequence and no point in it will be ranked.

Therefore every non-empty open $U \subset X$ loses a point in the first ε -Cantor-Bendixson derivative. For each U let $V \subset U$ be open such that $\#_\varepsilon^{\text{ent}} V < \omega$. Let x_0, \dots, x_n be a maximal ($> \varepsilon$)-separated subset of V . Now consider $V^\dagger = V \setminus \bigcup_{i < n} B_{\leq \varepsilon}(x_i)$, this is an open set since closed balls are topologically closed, furthermore $V^\dagger \subseteq B_{\leq \varepsilon}(x_0)$, otherwise x_0, \dots, x_n would not be maximal, therefore x_0 is contained in V and therefore U and is (d, ε) -isolated.

(ii) Let $U \subset X$ be open. By assumption there exists some x_0 which is $(d, 2^0)$ -isolated. Furthermore x_0 has a compact sub-neighborhood F_0 such that $F_0 \subseteq B_{\leq 2^0}(x_0)$. Proceed by induction. For each n let x_n be a $(d, 2^{-n})$ -isolated point contained in F_{n-1}° , and let F_n be a compact sub-neighborhood of $B_{\leq 2^{-n}}(x_n)$. By compactness $\bigcap_{n < \omega} F_n$ is non-empty. Furthermore since the F_n are contained in a descending sequence of closed 2^{-n} -balls, every sequence of points $a_n \in F_n$ forms a Cauchy sequence, so the intersection consists of a single point, x . By construction every closed ball centered at x contains an open neighborhood of x , so x is d -isolated. \square

ε -Morley Rank and Degree

Definition B.2.17. Let $B \subseteq S_n(A)$ be an arbitrary subset:

- (i) The ε -Morley rank of B , $MR_\varepsilon(B)$ is $CB_\varepsilon(\pi^{-1}(B))$ where $\pi : S_n(\mathfrak{C}) \rightarrow S_n(A)$ is the natural projection and \mathfrak{C} is the monster model.
- (ii) The ε -Morley degree of B , $Md_\varepsilon(B)$ is $CBd_\varepsilon(\pi^{-1}(B))$.
- (iii) The ε -Morley (rank, degree) of B , $M(R, d)_\varepsilon(B)$ is $(MR_\varepsilon(B), Md_\varepsilon(B))$. \triangleleft

The ε -Morley (rank, degree) is always ordered lexicographically (rank then degree).

Note that this is a well-ordering, even though the degree as we've defined it may turn out to be infinite for arbitrary sets.

We'll need the following fact:

Lemma B.2.18. *If $F \subseteq S_n(A)$ is a closed set, then*

$$MR_\varepsilon(F) = \min\{MR_\varepsilon(\llbracket\varphi \leq r\rrbracket) \mid \varphi \text{ a restricted } \mathcal{L}_A\text{-formula, } r \in \mathbb{Q}, \varphi^p \leq r\}.$$

Proof. First note that clearly $MR_\varepsilon(F) \leq \min\{MR_\varepsilon(\llbracket\varphi \leq r\rrbracket) \mid \varphi \text{ a } \mathcal{L}_A\text{-formula, } r \in \mathbb{Q}, \varphi^p \leq r\} = \alpha$. Now assume that $MR_\varepsilon(F) < \alpha$, so there exists an open set $U \subseteq S_n(\mathfrak{C})$ such that $\pi^{-1}(F) \subseteq U$, where $\pi : S_n(\mathfrak{C}) \rightarrow S_n(A)$ is the natural surjective projection map, and $MR_\varepsilon(U) < \alpha$. Now let $\{V_i\}_{i \in I}$ be the collection of all open subsets of $S_n(A)$ such that $F \cap \overline{V_i} = \emptyset$. Since type space is compact Hausdorff we have that $\bigcup V_i = S_n(A) \setminus F$. By continuity each $\pi^{-1}(V_i)$ is open, and furthermore they cover $S_n(\mathfrak{C}) \setminus U$ (because they cover $S_n(\mathfrak{C}) \setminus \pi^{-1}(F)$, which is a superset). By compactness there is a finite sub-cover, call it $\pi^{-1}(V_0), \dots, \pi^{-1}(V_{k-1})$. Then $G = \overline{V_0} \cup \dots \cup \overline{V_{k-1}}$ has $F \cap G = \emptyset$, so by compactness there must exist a restricted \mathcal{L}_A -formula φ and rational number r such that $F \subseteq \llbracket\varphi \leq r\rrbracket \subseteq S_n(A) \setminus G$. Then by construction $\pi^{-1}(F) \subseteq \pi^{-1}(\llbracket\varphi \leq r\rrbracket) \subseteq U$, so we have $MR_\varepsilon(\llbracket\varphi \leq r\rrbracket) \leq MR_\varepsilon(U) < \alpha$, which is a contradiction, therefore we have that $MR_\varepsilon(F) = \min\{MR_\varepsilon(\llbracket\varphi \leq r\rrbracket) \mid \varphi \text{ a restricted } \mathcal{L}_A\text{-formula, } r \in \mathbb{Q}, \varphi^p \leq r\}$. \square

Analogously to the discrete case, a ranked type is determined (although with some error) by any closed formula (and in fact any subset of type space whatsoever) it satisfies which has the same rank and degree:

Proposition B.2.19. *If $B \subseteq S_n(A)$ is an arbitrary subset, $p_i \in B$ are complete n -types*

over A such that $MR_\varepsilon(B) = MR_\varepsilon(p_i) = \alpha$ and $\sum_i Md_\varepsilon(p_i) > Md_\varepsilon(B)$, then for some $i \neq j$, $d(p_i, p_j) \leq \varepsilon$. In particular any collection of $Md_\varepsilon(B) + 1$ many p_i is sufficient.

Proof. Assume for the sake of contradiction that $d(p_i, p_j) > \varepsilon$ for all $i \neq j$, then $d(\pi^{-1}(p_i), \pi^{-1}(p_j)) > \varepsilon$ as well. Let $d_i = Md_\varepsilon(p_i)$ and for each i let $q_{i0}, q_{i1}, \dots, q_{i(d_i-1)}$ be a maximal ($> \varepsilon$)-separated subset of $\pi^{-1}(p_i)$ such that each q_{in} has ε -Morley rank α . Then $\{q_{in}\}_{n < d_i}$ is a ($> \varepsilon$)-separated subset of $\pi^{-1}(B)$ with each element of rank α and with size strictly larger than the ε -Morley degree of B , which is a contradiction, therefore for some $i \neq j$, $d(p_i, p_j) \leq \varepsilon$. \square

Corollary B.2.20. *If $B \subseteq S_n(A)$ is an arbitrary subset, $p, q \in B$ are complete n -types over A and $M(R, d)_\varepsilon(B) = M(R, d)_\varepsilon(p) = M(R, d)_\varepsilon(q) = (\alpha, d)$, then $d(p, q) \leq \varepsilon$.*

Approximate Morley Sequences

Now comes the core technical hiccup in the proof. We can't actually form Morley sequences, so we have to settle for arbitrarily good approximations. We also need to throw some extra scaffolding into the construction of the Morley sequence:

Proposition B.2.21 ([BY05]). *Let T be a countable complete ω -stable continuous first-order theory with monster model \mathfrak{C} . Let $D \subset \mathfrak{C}$ be a small set of parameters with $|D| = \lambda$, and let $\{a_i\}_{i < \lambda^+} \subset \mathfrak{C}$ be a sequence of elements. Then for any $\varepsilon > 0$ there exists a sub-sequence $\{a_{i(j)}\}_{j < \lambda^+}$ which is 'approximately indiscernible' in the following sense: for any two increasing multi-indices $I, J \in [\lambda^+]^{<\omega}$ of the same length k , $d(tp(a_{i(I)}/D), tp(a_{i(J)}/D)) < \varepsilon$, where $a_{i(I)} = a_{i(I_0)}a_{i(I_1)}\dots a_{i(I_{k-1})}$ (and likewise for J).*

Proof. Let $\llbracket \varphi \leq r \rrbracket$ with restricted $\mathcal{L}_{\mathfrak{C}}$ -formula φ and rational r be satisfied by λ^+ many elements of $\{a_i\}_{i < \lambda^+}$ such that $\llbracket \varphi \leq r \rrbracket$ has minimal ε -Morley (rank, degree) equal to

(α, d) among such formulas. Note that $\llbracket d(x, x) \leq 1 \rrbracket$ is satisfied by λ^+ many such elements, so the collection of such formulas is non-empty. By adding finitely many new elements to D we may assume that φ is a restricted \mathcal{L}_D -formula.

Now construct $\{a_{i(j)}\}$ inductively along with a sequence of sets D_j , where $D_{-1} = D$ and $D_{<j} = \bigcup_{k < j} D_k$. Also let $a_{i(<j)} = \{a_{i(k)}\}_{k < j}$.

At stage j , find a_i with index greater than any chosen so far such that $\varphi(a_i) \leq r$ and $M(R, d)_\varepsilon(\text{tp}(a_i/D_{<j}a_{i(<j)})) = (\alpha, d)$. Such an a_i must exist, otherwise every a_i (with index greater than any chosen so far) must satisfy some closed formula $\llbracket \psi_i \leq r_i \rrbracket$ with $M(R, d)_\varepsilon(\llbracket \psi_i \leq r_i \rrbracket) < (\alpha, d)$ by Lemma B.2.18, but there are only λ many such formulas (because we only considered restricted formulas), so for some $\llbracket \psi_i \leq r_i \rrbracket$, λ^+ many elements of $\{a_i\}_{i < \lambda^+}$ satisfy the formula (because λ^+ is a regular cardinal), which is a contradiction, therefore some such a_i must exist. Let $a_{i(j)} = a_i$.

To build D_j , for each pair of increasing multi-indices $I, J \leq j$ of length k such that at least one of I or J contains j (so that we don't construct duplicates), find a pair of k -tuples $\bar{b}_{I,J}$ and $\bar{b}_{J,I}$ satisfying the conditions

- $a_{i(I)}\bar{b}_{I,J} \equiv_D \bar{b}_{J,I}a_{i(J)}$ and
- $d^{\mathfrak{c}}(a_{i(I)}, \bar{b}_{I,J}) = d^{\mathfrak{c}}(\bar{b}_{J,I}, a_{i(J)}) = d(\text{tp}(a_{i(I)}/D), \text{tp}(a_{i(J)}/D))$.

Finally add $\bar{b}_{I,J}$ and $\bar{b}_{J,I}$ to $D_{<j}$ to get D_j . (Essentially: $\bar{b}_{I,J}$ is an automorphic image of $a_{i(J)}$ taking it as near as possible to $a_{i(I)}$ while fixing D and $\bar{b}_{J,I}$ is the corresponding pre-image of $a_{i(I)}$.) Note that at each stage $|D_j| = \lambda$.

Lastly we need to verify the approximate indiscernibility condition. Assume that for some k we've show that for any two increasing multi-indices I, J of length $k - 1$ that $d(\text{tp}(a_{i(I)}/D), \text{tp}(a_{i(J)}/D)) < \varepsilon$ (which is trivially satisfied when $k = 1$). Note that by

construction since $a_{i(I)}\bar{b}_{I,J} \equiv_D \bar{b}_{J,I}a_{i(J)}$ we have that $S_1(Da_{i(I)}\bar{b}_{I,J})$ and $S_1(D\bar{b}_{J,I}a_{i(J)})$ are ‘isomorphic as type spaces’ or more precisely related by some automorphism h of \mathfrak{C} that takes $\bar{b}_{J,I}a_{i(J)}$ to $a_{i(I)}\bar{b}_{I,J}$ and fixes D . Note that for any $a_{i(j)}$, and any $D' \subset D_{<j}a_{i(<j)}$ such that $D \subseteq D'$, we still have $M(R, d)_\varepsilon(\text{tp}(a_{i(j)}/D')) = (\alpha, d)$, because removing parameters can only increase ε -Morley (rank, degree) (since the preimage in $S_1(\mathfrak{C})$ is a larger set) but $\text{tp}(a_{i(j)}/D') \in \llbracket \varphi < 1 \rrbracket$, so the ε -Morley (rank, degree) can be no larger than (α, d) , so it must still be (α, d) . Therefore by Corollary B.2.20 for any $p, q \in S_1(Da_{i(I)}\bar{b}_{I,J}) \cong S_1(D\bar{b}_{J,I}a_{i(J)})$ which are contained in $\llbracket \varphi < 1 \rrbracket$ and have ε -Morley (rank, degree) (α, d) , we have $d(p, q) < \varepsilon$. Therefore if $I < \ell$ and $J < m$, then we have that $d(\text{tp}(a_{i(\ell)}/Da_{i(I)}\bar{b}_{I,J}), \text{tp}(h(a_{i(m)})/Da_{i(I)}\bar{b}_{I,J})) < \varepsilon$. Therefore there’s another automorphism g of \mathfrak{C} , fixing $Da_{i(I)}\bar{b}_{I,J}$ such that $d^\mathfrak{C}(a_{i(\ell)}, g(h(a_{i(m)}))) < \varepsilon$, and so we have $d^\mathfrak{C}(a_{i(I)}a_{i(\ell)}, g(h(a_{i(J)}a_{i(m)}))) < \varepsilon$. Therefore for any two increasing multi-indices K, L of length k we have $d(\text{tp}(a_{i(K)}/D), \text{tp}(a_{i(L)}/D)) < \varepsilon$. So by induction we have this condition for all $k < \omega$. \square

Corollary B.2.22. *Let T be a countable complete ω -stable continuous first-order theory, D a small set of parameters with $|D| = \lambda$, and $\{a_i\}_{i < \lambda^+}$, a sequence of elements which is $(\geq \varepsilon)$ -separated. Then there exists a descending nested sequence of subsequences $\mathcal{I}_n = \{a_{i(j,n)}\}_{j < \lambda^+}$ such that $\mathcal{I}_{n+1} \subset \mathcal{I}_n$ and such that the sequence of λ^+ -types $\text{tp}(\mathcal{I}_n/D)$ forms a Cauchy sequence in $S_{\lambda^+}(D)$ whose limit is the type of a non-constant D -indiscernible sequence of length λ^+ .*

Proof. By iterating the previous proposition we get the descending nested sequence of subsequences of $\{a_i\}_{i < \lambda^+}$. The condition that the sequence is $(\geq \varepsilon)$ -separated is

preserved by passing to subsequences, so the limiting type is $(\geq \varepsilon)$ -separated coordinate-wise and therefore non-constant. \square

Morley's Theorem for Continuous Logic

Proposition B.2.23. *Let T be a countable complete ω -stable continuous first-order theory and $\mathfrak{M} \models T$ a non-saturated model with density character $\kappa \geq \aleph_1$. Then for any countable $C \subset M$ and any $\mu \geq \aleph_1$, there exists $\mathfrak{N} \models T$ containing C with $\#^{dc} N = \mu$ such that \mathfrak{N} is not \aleph_1 -saturated. So in particular there exists non-saturated models of every uncountable density character.*

Proof. Let \mathfrak{M} be a non-saturated model of T of density character κ . Since \mathfrak{M} is non- μ -saturated it is not approximately μ -saturated and there exists some $A \subset M$ with $|A| = \lambda < \mu$ and some $p \in S_1(A)$ such that for every $a \in M$, $d(p, \text{tp}(a/A)) > \varepsilon$ for some $\varepsilon > 0$, or in other words $\text{tp}(a/A) \notin B_{\leq \varepsilon}(p)$. We can assume without loss of generality that A contains C .

Since \mathfrak{M} has density character κ , there is some $\delta > 0$ such that $\#_{\delta}^{\text{ent}} M \geq \lambda^+$, therefore we can find a $(\geq \delta)$ -separated sequence $\{a_i\}_{i < \lambda^+}$ of size λ^+ (since λ^+ is a regular cardinal), and by Corollary B.2.22 we get a descending nested sequence of subsequences $\{a_i\}_{i < \lambda^+} \supseteq \mathcal{I}_0 \supseteq \mathcal{I}_1 \supseteq \dots$ such that $\text{tp}(\mathcal{I}_n/A)$ forms a Cauchy sequence in $S_{\lambda^+}(A)$, limiting to some $\text{tp}(\{b_i\}_{i < \lambda^+}/A)$ (with b_i from the Monster model) which is $(\geq \delta)$ -separated coordinatewise.

Now for a given increasing multi-index $I < \lambda^+$, let $a_{I,n}$ be the sub-tuple of \mathcal{I}_n whose elements are the I_0 th, I_1 st, I_2 nd, and so on elements of \mathcal{I}_n (and importantly not the corresponding indices in $\{a_i\}_{i < \lambda^+}$, so for a fixed I , $a_{I,n}$ can vary wildly as n changes).

Consider each restricted \mathcal{L}_A -formula $\varphi(x, \bar{y})$ such that for some increasing multi-index

I , and all sufficiently large n , $\varphi(x, a_{I,n}) < \frac{1}{4}$ is satisfied in \mathfrak{M} . Call the set of all such formulas Σ . Then since $\mathfrak{M} > \varepsilon$ -omits p , for any $c \in \mathfrak{M}$ satisfying $\varphi(c, a_{I,n}) < \frac{1}{4}$, we have that $d(p, \text{tp}(c/A)) > \varepsilon$. Therefore there exists a pair of \mathcal{L}_A -formulas, $\psi_\varphi(x)$ and $\chi_\varphi(x)$ such that $\psi_\varphi(c) = 0$, $\chi_\varphi^p = 0$, and $\llbracket \psi_\varphi(x) \leq \frac{1}{2} \rrbracket$ and $\llbracket \chi_\varphi(x) \leq \frac{1}{2} \rrbracket$ are $> \frac{\varepsilon}{2}$ -inconsistent. Therefore by uniform continuity and approximate indiscernibility we have that for all I and sufficiently large n , $\inf_x \varphi(x, a_{I,n}) \vee \psi_\varphi(x) < \frac{1}{4}$. Therefore $T \models \inf_x \varphi(x, b_I) \vee \psi_\varphi(x) \leq \frac{1}{4}$.

Let $B_0 = C$. Iteratively construct a sequence of countable subsets of A , B_k , like so: for each k let B_{k+1} be B_k together with the parameters used in $\psi_\varphi(x)$ and $\chi_\varphi(x)$ for all \mathcal{L}_{B_k} -formulas in Σ . Let $B = \bigcup_{k < \omega} B_k$, and note that this is a countable set of parameters containing C . Let $q = p \upharpoonright B$.

Now let $\{c_i\}_{i < \mu}$ be a B -indiscernible sequence with the same Ehrenfeucht–Mostowski type as $\{b_i\}_{i < \lambda^+}$. Let \mathfrak{A} be a model containing $B \cup \{c_i\}_{i < \mu}$, and let \mathfrak{N}_0 be a sub-pre-model constructed over $B \cup \{c_i\}_{i < \mu}$. By downward Löwenheim-Skolem we can assume that \mathfrak{N}_0 has density character μ . The claim is that $\mathfrak{N}_0 (> \frac{\varepsilon}{4})$ -omits q and is thus not \aleph_1 -saturated. To see this assume that \mathfrak{N}_0 has some element a such that $d(q, \text{tp}(a/B)) \leq \frac{\varepsilon}{4}$, then since $\text{tp}(a/B \cup \{c_i\}_{i < \mu})$ is principal there exists some \mathcal{L}_B -formula and c_I such that $\varphi(a, c_I) = 0$ and for all b , $\varphi(b, c_I) < \frac{1}{4}$ implies $d(\text{tp}(a/B), \text{tp}(b/B)) \leq \frac{\varepsilon}{4}$. But by construction then there exists $\psi_\varphi(x)$ and $\chi_\varphi(x)$ such that $\varphi(b, c_I) < \frac{1}{4}$, $\psi_\varphi(b) < \frac{1}{4}$, but this implies that $d(\text{tp}(b/B), q) > \frac{\varepsilon}{2}$, but this is impossible because $d(q, \text{tp}(b/B)) \leq d(q, \text{tp}(a/B)) + d(\text{tp}(a/B), \text{tp}(b/B)) \leq \frac{\varepsilon}{2}$. Therefore $\mathfrak{N}_0 (> \frac{\varepsilon}{4})$ -omits q and $\mathfrak{N} = \overline{\mathfrak{N}_0} (\geq \frac{\varepsilon}{4})$ -omits q and is the required non- \aleph_1 -saturated model of T with density character μ . □

A more careful argument should be able to get a sharper ($\geq \varepsilon$)-omitting result, or

something like that.

Corollary B.2.24 ([BY05, SU11]). *Let T be a countable complete continuous first-order theory which is κ -categorical for some $\kappa \geq \aleph_1$, then T is μ -categorical for all $\mu \geq \aleph_1$.*

Proof. By Proposition B.2.1, it's sufficient to show that if T is ω -stable and has a non-saturated model of density character μ for some $\mu \geq \aleph_1$, then it has a non-saturated model of density character κ for every $\kappa \geq \aleph_1$, which is a direct consequence of Proposition B.2.23. \square

B.2.3 A Proof after Chang and Keisler

This section contains a novel adaptation of the proof of Morley's theorem presented in [CK90] to the context of continuous logic.

More general forms of the following two results were originally shown in [BY10b]. Here we are presenting slightly simpler proofs in the special cases that are relevant.

Lemma B.2.25. *Let $F, G \subseteq S_n(A)$ be closed sets such that $d_{inf}(F, G) > \varepsilon > 0$. Then there exists a continuous metrically 1-Lipschitz function (and equivalently an A -definable 1-Lipschitz predicate) $P : S_n(A) \rightarrow [0, \varepsilon]$ such that $P(F) = \{0\}$ and $P(G) = \{\varepsilon\}$.*

Proof. Let $R = \{r_i\}_{i < \omega}$ be an enumeration of a countable dense subset of $(0, \varepsilon)$.

Construct a sequence of pairs of closed sets $\{F_i, G_i\}_{i < \omega}$ by induction on i , where by induction we will ensure that for each $i, j < \omega$ such that $r_i < r_j$, $d_{inf}(F_i, G_j) > r_j - r_i$, $d_{inf}(F, G_i) > r_i$, $d_{inf}(G, F_i) > \varepsilon - r_i$, and $F_i \cup G_i = S_n(A)$.

At stage $i < \omega$, construct F_i and G_i like so: Let $K = F^{\leq r_i} \cup \bigcup \{F_j^{\leq r_i - r_j} : j < i, r_j < r_i\}$ and $L = G^{\leq \varepsilon - r_i} \cup \bigcup \{G_j^{\leq r_j - r_i} : j < i, r_i < r_j\}$. These are finite unions of closed sets and

are therefore closed. By the induction hypothesis and the triangle inequality, $K \cap L = \emptyset$, so since $S_n(A)$ is compact, we can find F_i and G_i such that $K \subseteq \text{int}F_i$, $L \subseteq \text{int}G_i$, and $F_i \cup G_i = S_n(A)$. (We could even make them closed formulas of restricted closed formulas, if we were so inclined.)

Now finally let $P(x) = \inf\{r_i : x \in F_i\}$ or ε if that set is empty. First, to see that $P(x)$ is continuous (which is really the same argument as in the normal proof of Urysohn's Lemma), note that for each $s \in [0, \varepsilon]$, the set $\llbracket P < s \rrbracket$ is equal to $\bigcup\{S_n(A) \setminus G_i : r_i < s\}$, because if $P(x) < s$, then for some $r_i < s$, $x \in F_i$, so for any r_j with $r_i < r_j \leq s$, $x \notin G_j$. This is a union of open sets and is therefore open. The same argument gives that $\llbracket P > s \rrbracket$ is open, therefore P is topologically continuous.

To see that P is metrically 1-Lipschitz, note that if $P(x) = s < t = P(y)$, then for any $\delta > 0$, we can find $i, j < \omega$ such that $s < r_i < r_j < t$ and such that $|s - r_i| < \delta$ and $|r_j - t| < \delta$, so since $x \in F_i$ and $y \in F_j$ we have $d(x, y) > r_j - r_i$. Since we can do this for arbitrarily small $\delta > 0$, we get that $d(x, y) \geq t - s$, which is precisely the 1-Lipschitz condition. \square

Note that $P(\bar{x}, \bar{a})$ being 1-Lipschitz in \bar{x} does not imply that $P(\bar{x}, \bar{b})$ is Lipschitz in \bar{x} for $\bar{b} \neq \bar{a}$ or that $P(\bar{x}, \bar{y})$ is Lipschitz in $\bar{x}\bar{y}$. It also doesn't imply that $P(x, \bar{a})$ can be uniformly approximated by Lipschitz restricted formulas with uniformly bounded Lipschitz constants. It is also not clear when we can accomplish the same as the lemma given $d_{inf}(F, G) = \varepsilon > 0$.

Corollary B.2.26. *The family of continuous Lipschitz functions is dense in $C(S_n(A), \mathbb{R})$ under the uniform norm.*

Proof. The family of continuous Lipschitz functions satisfy the requirements of the lattice

version of the Stone-Weierstrass theorem (Fact A.2.5), so they are dense in $C(S_n(A), \mathbb{R})$.

□

Corollary B.2.27. *For $p, q \in S_n(A)$,*

$$\begin{aligned} d(p, q) &= \sup\{|P(p) - P(q)| : P \text{ an } A\text{-definable } 1\text{-Lipschitz predicate}\} \\ &= \sup\left\{\frac{1}{L(P)}|P(p) - P(q)| : P \text{ an } A\text{-definable Lipschitz predicate}\right\}, \end{aligned}$$

where $L(P)$ is the Lipschitz constant of P (with the understanding that $\frac{0}{0} = 0$). Furthermore, if $S_n(A)$ is topologically separable, then there is a countable family $\{P_i\}_{i < \omega}$ of 1-Lipschitz A -definable predicates such that the supremum can be taken over $\{P_i\}_{i < \omega}$.

Proposition B.2.28. *Let T be a countable complete t.t. theory with non-compact models. Let $\mathfrak{M} \models T$ be inseparable. Then \mathfrak{M} has arbitrarily large elementary extensions, \mathfrak{N} , such that for any countable $A \subset \mathfrak{M}$, if $p \in S_1(A)$ is realized in \mathfrak{N} , then for every $\varepsilon > 0$, there is a $q \in S_1(A)$ with $d(p, q) \leq \varepsilon$ such that \mathfrak{M} realizes q .*

Proof. Since \mathfrak{M} is inseparable, there is some $\eta > 0$ such that $\#_{\geq \eta}^{\text{ent}} \mathfrak{M} \geq \aleph_1$. Let $Q \subseteq M$ be a maximal ($\geq \eta$)-separated set of cardinality \aleph_1 (which exists, since \aleph_1 is a regular cardinal). Note that since Q is a maximal ($\geq \eta$)-separated set, it forms a ($< \eta$)-cover of \mathfrak{M} .

Consider Q as a subset of $S_1(\mathfrak{M})$, and let X be its set of condensation points (i.e. the set of all points p such that for all open $U \ni p$, $|Q \cap U| \geq \aleph_1$). Clearly this is a closed set. It must be non-empty, because if it were empty then by compactness there would be a finite cover \mathcal{U} of $S_1(\mathfrak{M})$ such that for each $U \in \mathcal{U}$, $|Q \cap U| \leq \omega$, contradicting the fact that Q is uncountable. Finally no $a \in \mathfrak{M} \subset S_1(\mathfrak{M})$ is an element of X , because

for any $a \in \mathfrak{M}$ there is a $b \in Q$ such that $d(a, b) < \eta$, so a has an open neighborhood containing only one element of Q .

By the above proposition, X has a d -isolated-in- X point t . Let a be a realization of t , and let \mathfrak{N} be atomic over $M \cup \{a\}$. Let $A \subset \mathfrak{M}$ be countable, and let $p \in S_1(A)$ be a type realized in \mathfrak{N} . We want to argue that for every $\varepsilon > 0$, there is a $q \in S_1(A)$ with $d(p, q) < \varepsilon$ such that q is realized in \mathfrak{M} . Since T is countable and A is countable, $C(S_n(A), \mathbb{R})$ is separable under the uniform norm, so in particular there is a countable family $\{P_i\}_{i < \omega}$ of A -definable 1-Lipschitz $[0, 1]$ -valued predicates such that for all $q_0, q_1 \in S_1(A)$, $d(q_0, q_1) = \sup_{i < \omega} |P_i(q_0) - P_i(q_1)|$. Since p is realized in \mathfrak{N} , it must be d -isolated over $M \cup \{a\}$. Let $\chi(x, a)$ be the M -definable distance predicate for p . We have that $\mathfrak{N} \models \sup_x |P_i(x) - P_i(p)| \div \chi(x, a)$ for all $i < \omega$ (this is just saying that P_i is 1-Lipschitz, so things near p need to take on values for P_i close to the value of P_i at p). Note also that $\mathfrak{N} \models \inf_x \chi(x, a)$. These predicates are both in the type $\text{tp}(a/M) = t(y)$.

Since t is d -isolated-in- X , we can find a definable predicate R such that $R(t) = 0$ and for all $q \in X$, $d(t, q) \leq R(q)$ (by the prop above we can find $f : X \rightarrow [0, 1]$ satisfying this property on X , then using the Tietze extension theorem we can extend f to the required R , which is a definable predicate because it is continuous on type space). Note that for any $\varepsilon > 0$, for any closed $F \subseteq \llbracket R \leq \varepsilon \rrbracket$, either $d(t, F) \leq \varepsilon$ or $|F \cap Q| \leq \omega$. (If $|F \cap Q| \geq \aleph_1$, then by compactness, $F \cap X \neq \emptyset$, but since $F \subseteq \llbracket R \leq \varepsilon \rrbracket$, this implies that $F \cap X \subseteq \llbracket R \leq \varepsilon \rrbracket \cap X \subseteq B_{\leq \varepsilon}(t) \cap X$.)

Choose $\varepsilon > 0$. The open formulas $U_i = \llbracket \sup_x |P_i(x) - P_i(p)| \div \chi(x, y) < \frac{\varepsilon}{2} \rrbracket$ and $U_{-1} = \llbracket \inf_x \chi(x, y) < \frac{\varepsilon}{2} \rrbracket$ are open neighborhoods of t , so since the formulas in the formulas U_i for $i > -1$ are uniformly Lipschitz there is a $\delta > 0$ so that $d(t, S_1(A) \setminus U_i) \geq \delta$. So by the comment above, for each U_i , we have that $|(\llbracket R \leq \frac{\delta}{2} \rrbracket \setminus U_i) \cap Q| \leq \omega$, which implies that

$|(\llbracket R < \frac{\delta}{2} \rrbracket \setminus U_i) \cap Q| \leq \omega$ (since that is a subset), but for each $-1 \leq i < \omega$, $\llbracket R < \frac{\delta}{2} \rrbracket \cap U_i$ is an open neighborhood of t , and thus has $|Q \cap \llbracket R < \frac{\delta}{2} \rrbracket \cap U_i| \geq \aleph_1$. By counting we get that $|Q \cap \llbracket R < \frac{\delta}{2} \rrbracket \cap \bigcap_{-1 \leq i < \omega} U_i| \geq \aleph_1$, so in particular $|Q \cap \bigcap_{-1 \leq i < \omega} U_i| \geq \aleph_1$, and there is some $b \in Q$ such that $\mathfrak{M} \models \sup_x |P_i(x) - P_i(p)| \div \chi(x, b) < \frac{\varepsilon}{2}$, for each $i < \omega$, and $\mathfrak{M} \models \inf_x \chi(x, b) < \frac{\varepsilon}{2}$.

Let $c \in \mathfrak{M}$ be such that $\mathfrak{M} \models \chi(c, b) < \frac{\varepsilon}{2}$. For each $i < \omega$ we have that $\mathfrak{M} \models |P_i(c) - P_i(p)| \div \chi(c, b) < \frac{\varepsilon}{2}$ and therefore $\mathfrak{M} \models |P_i(c) - P_i(p)| < \varepsilon$, so by the corollary above we have that $d(p, \text{tp}(c/A)) \leq \varepsilon$, thus $q = \text{tp}(c/A)$ is the required type for $\varepsilon > 0$. \square

The following is an adaptation of [CK90, Thm. 7.1.1].

Proposition B.2.29. *Let T be a countable complete theory with non-compact models such that for some $\kappa \geq \aleph_1$, T has a non-saturated model of density character κ , then T has a non-saturated model of density character \aleph_1 .*

Proof. Let \mathfrak{A} be a non-saturated model of density character $\kappa \geq \aleph_1$. Let $E \subset A$ be some set of parameters with cardinality $|E| < \kappa$ such that \mathfrak{A} ($> \varepsilon$)-omits some type $p \in S_1(E)$ (i.e. for any type q realized by \mathfrak{A} in $S_1(E)$, $d(p, q) > \varepsilon$). By the downward Löwenheim-Skolem theorem we may assume that the density character of \mathfrak{M} is $|E|^+$. Let \mathfrak{A}_0 be a dense pre-structure with cardinality $|E|^+$.

Recast \mathfrak{A}_0 as a discrete structure, with a predicate for every open or closed formulas $\llbracket \varphi \triangleleft r \rrbracket$ with φ a restricted formula over \emptyset , r a rational number, and $\triangleleft \in \{<, \leq\}$ (so there are only countably many such predicates). Add the following things to make the structure \mathfrak{A}^* :

- For some $\varepsilon > 0$, \mathfrak{A}^* has a ($> \varepsilon$)-separated set of cardinality $|E|^+$, since $|E|^+$ is a regular cardinal. Let Q be a predicate selecting out a ($> \varepsilon$)-separated set of

cardinality $|E|^+$.

- Let f be a function that is a bijection between A_0 and $Q^{\mathfrak{A}^*}$.
- Let E be a predicate selecting out the set E .
- For each restricted formula $\varphi(x, \bar{e})$ with $\bar{e} \in E$, let $s_{\varphi, \bar{e}}$ be the smallest number such that $B_{\leq \varepsilon/2}(p) \subseteq \llbracket \varphi(x, \bar{e}) \leq s_{\varphi, \bar{e}} \rrbracket$. Pick some set $U \subset A$ with cardinality $|E|$ and associate each element bijectively with some formula of the form $\llbracket \varphi(x, \bar{e}) \leq r \rrbracket$ with r a rational number greater than or equal to $s_{\varphi, \bar{e}}$. (All such rational numbers should appear in the list.)
- Let S be a binary predicate such that for any $a \in A_0$ and $c \in U$, $\mathfrak{A}^* \models S(a, c)$ if and only if $\mathfrak{A}_0 \models \varphi(x, \bar{e}) \leq r$ where $\llbracket \varphi(x, \bar{e}) \leq r \rrbracket$ is the formula corresponding to c .
- Let R be a binary predicate such that for any $c \in U$ and $e \in E$, $\mathfrak{A}^* \models R(c, e)$ if and only if e is one of the constants occurring in the formula corresponding to c .

Now by Theorem 3.2.11 in Chang and Keisler, \mathfrak{A}^* has a countable elementary substructure \mathfrak{B} with an elementary extension \mathfrak{C} such that $U^{\mathfrak{B}} = U^{\mathfrak{C}}$, $E^{\mathfrak{B}} = E^{\mathfrak{C}}$, and $|C| = \aleph_1$. In particular, note that no non-standard elements of U show up in \mathfrak{C} ; they still correspond to actual restricted formulas.

Now let \mathfrak{C}^\dagger be the metric pre-structure corresponding to \mathfrak{C} .

First note that $\#^{dc} \mathfrak{C}^\dagger = \aleph_1$. It has a dense subset of cardinality \aleph_1 and a ($> \varepsilon$)-separated set of cardinality \aleph_1 , so its density character is \aleph_1 . (In principle the dense subset of cardinality \aleph_1 might contain many elements with infinitesimal distances between them. So merely having a dense subset in the discrete structure of cardinality \aleph_1

isn't enough to ensure that the density character is \aleph_1 . Conversely, in general a structure can have a $(> \varepsilon)$ -separated set of cardinality \aleph_1 but have a larger density character, although it can't happen in this particular case.)

Secondly note that the set $E^{\mathfrak{C}} \subset E^{\mathfrak{A}^*}$ is a countable set of parameters and $U^{\mathfrak{C}} \subset U^{\mathfrak{A}^*}$ defines a consistent partial type over $E^{\mathfrak{C}}$ (we know it is consistent because it is a restriction of $B_{\leq \varepsilon/2}(p) \subseteq S_1(A)$). The predicate R ensures that the parameters needed for any formula in $U^{\mathfrak{C}}$ are actually in $E^{\mathfrak{C}}$, since they are algebraic over it.

Finally the theory of \mathfrak{C} knows that no element of \mathfrak{C} simultaneously satisfies all formulas of the form $\llbracket \varphi(x, \bar{e}) \leq r \rrbracket$ encoded by elements of $U^{\mathfrak{C}}$, since the same is true in \mathfrak{A}^* . This partial type is a superset of $B_{\leq \varepsilon/2}(p \upharpoonright E^{\mathfrak{C}})$, so \mathfrak{C}^\dagger ($\geq \frac{\varepsilon}{2}$)-omits $p \upharpoonright E^{\mathfrak{C}}$ and therefore so does its completion, $\overline{\mathfrak{C}^\dagger}$, which is the required non-saturated model of density character \aleph_1 . \square

Theorem B.2.30. *If T is a countable complete theory with non-compact models which is κ -categorical for some $\kappa \geq \aleph_1$, then it is λ -categorical for every $\lambda \geq \aleph_1$.*

Proof. Every model of T of density character κ is saturated, so by Proposition B.2.28, every model of T of density character \aleph_1 is saturated and then by Proposition B.2.29 every model of T of density character λ for any $\lambda \geq \aleph_1$ is saturated. Therefore T is λ -categorical for every $\lambda \geq \aleph_1$. \square

B.3 Encoding Structures as Metric Spaces

It is a well known fact [Hod93] that any discrete structure with finite signature can be encoded as a graph in a particularly strong way:

Fact B.3.1. *For any finite signature \mathcal{L} there is a sentence χ in a language with a single binary predicate such that every model of χ is a graph and the class of models of χ is bi-interpretable with the class of \mathcal{L} -structures with more than one element. Furthermore this bi-interpretation preserves embeddings and is computable in the sense that presentations of models of χ are uniformly computable from presentations of the corresponding \mathcal{L} -structure and vice versa.*

This immediately implies that the set of validities involving a single binary predicate is undecidable, even though monadic first-order logic, which involves only unary predicates, is decidable. This is in contrast to the situation in continuous first-order logic, introduced in [BYBHU08]. There is an easy encoding of a graph (V, E) as a metric space (V, d) wherein

$$d(x, y) = \begin{cases} 0 & x = y \\ \frac{1}{2} & xEy \\ 1 & \text{otherwise} \end{cases} .$$

So the set of continuous validities in the empty signature is undecidable for any reasonable notion of computable continuous formulas. Moreover, discrete structures can be encoded as metric spaces, in light of Fact B.3.1.

The proof of Fact B.3.1 uses a ‘tag construction,’ in which each tuple x_0, x_1, \dots, x_{k-1} related by some relation P is connected by a tag which is engineered to distinguish each x_i and to be distinguishable from tags corresponding to relations other than P . This particular construction does not generalize in any satisfactory way to metric structures, but the main result of this section is a generalization of Fact B.3.1:

Theorem B.3.2. *For any countable metric signature \mathcal{L} and $r > 0$, there is a theory T*

in the empty signature such that the class of models of T is bi-interpretable with the class of \mathcal{L} -structures with diameter $\geq r$. This bi-interpretation preserves embeddings and d -finiteness of types. If the original structure is not strongly infinite dimensional, then the interpreted structure will also not be strongly infinite dimensional. Furthermore, the bi-interpretation is computable in the sense that presentations of models of T are uniformly computable from presentations of the corresponding \mathcal{L} -structure and vice versa.

There are some improvements in Theorem B.3.2 over Fact B.3.1, namely that the encoding works in the empty signature—which is largely cosmetic—and that we can encode countable signatures rather than just finite ones. *d -finiteness of types* is a technical niceness condition introduced in [BYU07] that will be discussed below. Strong infinite dimensionality is relevant from the point of view of computable structure theory, as the continuous degree of a point in a finite dimensional or weakly infinite dimensional metric space is always total [KP14]. These two concepts play no essential role in the construction, although they do motivate a particular choice in it, namely using a disjoint union construction rather than a product construction.

The restriction that the metric structures have diameter uniformly bounded below is the necessary analog of the ‘more than one element’ restriction. A simple compactness argument shows that we could never have uniform bi-interpretability between a single elementary class of metric spaces and the class of all \mathcal{L} -structures of positive diameter. In both the discrete case and the metric case we could avoid this non-uniformity by appending a new sort to every structure that always contains precisely two elements distance 1 apart. Also it should be noted that this is a non-issue from a computable structure theory point of view.

Finally there is the issue of finite axiomatizability, which the generalization loses,

although, as will be discussed in Section B.4, there is no clear analog of finite axiomatizability in continuous logic.

B.3.1 Preliminaries

For the purposes of this section we will assume that all signatures are relational and all predicates are $[0, 1]$ -valued. We also need to present some unsurprising definitions. Note that we will not be allowing the imaginary **2**.

For computable metric signatures, obviously we should require that the predicate ranges and maximum sort diameters be uniformly computable before recasting in the form above (although really all we need are uniformly computable upper and lower bounds), in order to ensure that we can uniformly compute presentations of recast structures from presentations of the original structures.

Definition B.3.3. A *computable metric signature* is a metric signature such that \mathcal{S} and \mathcal{P} are computable subsets of ω , a is a computable function which is total on \mathcal{P} , and $P \mapsto \Delta_P$ is a uniformly computable family of total computable functions. \triangleleft

We should be clear about what a computable metric structure is.

Definition B.3.4. Given a computable metric signature \mathcal{L} , a *computable \mathcal{L} -structure* is an \mathcal{L} -structure whose universes are a uniformly computable family of computable metric spaces (in the sense of [Wei00]) and whose predicate interpretations are all uniformly computable functions. \triangleleft

Because we are using these notions in a detailed way, we will modify and expand Definition 1.3.2 for the purposes of this section.

- Definition B.3.5.** (i) For a metric signature \mathcal{L} , a *finitary \mathcal{L} -formula* is an expression of the form $\sum_{n < \omega} 2^{-(n+1)} \varphi_n$, with φ_n a sequence of $[0, 1]$ -valued restricted \mathcal{L} -formulas such that the entire sequence contains finitely many free variables. Such a formula has a syntactic modulus of uniform continuity of $\sum_{n < \omega} 2^{-(n+1)} \Delta_{\varphi_n}$.
- (ii) An *ω -infinitary \mathcal{L} -formula* is an expression of the same form allowing possibly infinitely many free variables.¹
- (iii) An *\mathcal{L} -formula* is either a finitary or an ω -infinitary \mathcal{L} -formula.
- (iv) A *computable \mathcal{L} -formula* is an \mathcal{L} -formula such that the sequence of formulas φ_n is computable (the φ_n are required to be restricted formulas and can therefore be encoded by natural numbers). \triangleleft

With some straightforward work, one can show the following fact.

Fact B.3.6. *If $\{\varphi_n\}_{n < m}$ is a uniformly computable sequence of \mathcal{L} -formulas for some $m \leq \omega$ and $F : [0, 1]^m \rightarrow [0, 1]$ is a computable function, then $F(\bar{\varphi})$ is logically equivalent to a computable \mathcal{L} -formula. Furthermore the equivalent formula is uniformly computable in $\{\varphi_n\}_{n < m}$ and F .*

B.3.2 Expansions

We need to specify a few notions of expansions and interdefinability in continuous logic.

Definition B.3.7. (i) For a given metric signature \mathcal{L} and a finitary \mathcal{L} -formula $\varphi(\bar{x})$,
 a *definitional expansion* of \mathcal{L} by φ is a metric signature \mathcal{L}^* containing the same

¹Such an expression has a uniformly computable syntactic modulus of uniform continuity in terms of the appropriate metric on ω -tuples, but it is somewhat more complicated to state.

sorts as \mathcal{L} and a single new predicate symbol P with $a(\varphi) = a(P)$ and $\Delta_\varphi = \Delta_P$. For \mathfrak{A} an \mathcal{L} -structure, the corresponding \mathcal{L}^* -structure \mathfrak{A}^* is given by interpreting P as φ . We also refer to iterated definitional expansions as definitional expansions.

- (ii) An \mathcal{L} -structure, \mathfrak{A} , and a \mathcal{K} -structure, \mathfrak{B} , are *interdefinable* if there are definitional expansions \mathfrak{A}^* and \mathfrak{B}^* which make them isomorphic up to relabeling of sorts and predicate symbols. (We allow metrics to be relabeled.) An elementary class C_0 of \mathcal{L} -structures and an elementary class C_1 of \mathcal{K} -structures are *interdefinable* if there are functors $F : C_0 \rightarrow C_1$ and $G : C_1 \rightarrow C_0$ given by uniform definitional expansions and relabelings which form an equivalence of categories, where we treat C_0 and C_1 as categories with elementary embeddings as morphisms (i.e. $F \circ G$ and $G \circ F$ are both naturally isomorphic to the identity functor²).

Note that we aren't requiring that the syntactic moduli of continuity match. \triangleleft

Lemma B.3.8. (i) *For any metric signature \mathcal{L} (not necessarily countable), there is a metric signature \mathcal{K} which is interdefinable with an imaginary expansion of \mathcal{L} such that \mathcal{K} has a uniform bound of 2 on the arities of its predicate symbols. For computable signatures, the signature \mathcal{K} is uniformly computable from \mathcal{L} , and presentations of \mathcal{L} -structures can be uniformly converted into corresponding presentations of \mathcal{K} -structure and vice versa.*

- (ii) *There is a \mathcal{K} -theory $T_{\mathcal{L}}$, uniformly computable from \mathcal{L} , such that the models of $T_{\mathcal{L}}$ are precisely the interpretations of \mathcal{L} -structures.*

²That is to say, for each structure \mathfrak{A} in C_0 , there is a designated isomorphism $\alpha_{\mathfrak{A}} : \mathfrak{A} \rightarrow G \circ F(\mathfrak{A})$ such that for any $\mathfrak{B} \in C_0$ and any elementary map $f : \mathfrak{A} \preceq \mathfrak{B}$, $\alpha_{\mathfrak{B}} \circ f = (G \circ F(f)) \circ \alpha_{\mathfrak{A}}$. And likewise for C_1 .

Proof. (i) For each predicate symbol p , we can define a unary formula on the sort $\prod_{O \in a(p)} O$ in the obvious way. These, together with projection maps between product sorts and the original \mathcal{L} -sorts, are clearly enough to define any predicate originally definable in an \mathcal{L} -structure in a completely uniform way. Since the projection maps are encoded as 2-ary predicates, we have the required arity bound. This procedure is also clearly uniformly computable, both for signatures and presentations of structures.

(ii) All that $T_{\mathcal{L}}$ needs to say is that the predicates corresponding to projection maps are actually projection maps and that the product sorts are products of the sorts they project onto. \square

Definition B.3.9. For any metric signature \mathcal{L} with designated home sort H and any real number r satisfying $0 < r \leq 1$, $C_{\mathcal{L},r}$ is the class of \mathcal{L} -structures \mathfrak{A} satisfying $\text{diam}(H^{\mathfrak{A}}) \geq r$. \triangleleft

The following lemma is the source of all non-uniformity relative to r in the entire construction and is analogous to the fact that a discrete structure with only one element cannot interpret any structure with more than one element. It could be avoided by appending the imaginary sort $\mathbf{2}$ and treating that as the designated home sort H .

Lemma B.3.10. *Let X be a compact metric space. For structures in the class $C_{\mathcal{L},r}$, there is a uniformly definable imaginary Y such that for any $\mathfrak{A} \in C_{\mathcal{L},r}$, $Y^{\mathfrak{A}} \cong X$, with each point of $Y^{\mathfrak{A}}$ and every continuous function $(Y^{\mathfrak{A}})^n \rightarrow [0, 1]$ uniformly \emptyset -definable.*

Proof. Let x_0, x_1, y_0, y_1 be variables in H , and consider the \mathcal{L} -formula

$$\rho(x_0, x_1, y_0, y_1) = \frac{1}{r} |d(x_0, x_1) - d(y_0, y_1)| \downarrow 1.$$

This is a pseudo-metric on H^2 . H^2/ρ contains more than one point for any $\mathfrak{A} \in C_r$, because of the diameter requirement. In particular it has a definable subset consisting of the ρ -equivalence classes of pairs satisfying $d(x_0, x_1) = 0$ and pairs satisfying $d(x_0, x_1) \geq r$, with each of those points being \emptyset -definable by the formulas $\frac{1}{r}d(x_0, x_1)$ and $1 - \frac{1}{r}d(x_0, x_1)$, respectively. Let D denote this definable set.

This definable set corresponds to the imaginary **2**, so the rest follows from Proposition 3.4.1. In order to verify computability we need to verify that there is a uniform computable definition of the Hilbert cube, $[0, 1]^\omega$, with the standard metric, but in order to accomplish this all we need to do is fix the standard computable surjection from C to $[0, 1]$ and then pass to the ω -power. Then any computable compact metric space X has a computable embeddings into $[0, 1]^\omega$ and total computable functions on X have uniformly computable extensions to all of $[0, 1]^\omega$ by Fact B.3.16, since infima of total computable functions on computable compact metric spaces are computable. \square

There are some potential subtleties involving uniform computability of formulas defining computable compact imaginaries and computable predicates on them. In the current context we only need Lemma B.3.10 for a small handful of very specific tame compact metric spaces, so we'll deal with computability on a case-by-case basis.

Lemma B.3.11. *For any $C_{\mathcal{L}, r}$, with $r > 0$, if $\{O_n^{\mathfrak{A}}\}_{n < k}$ is a finite collection of sorts of diameter ≤ 1 , then the disjoint union $U = \bigsqcup_{n < k} O_n$ with metric $d(x, y) = 1$ for $x \in O_n$ and $y \in O_m$ (with $n \neq m$) and $d(x, y) = d_{O_n}(x, y)$ for $x, y \in O_n$ is a uniformly definable imaginary in $C_{\mathcal{L}, r}$. Furthermore the formulas defining U are uniformly computable in \mathcal{L} , r , and the list of sorts.*

Proof. By Lemma B.3.10, the discrete space $\Delta_k = \{0, \dots, k-1\}$, with the metric $\delta(x, y) =$

1 if $x \neq y$, is uniformly an imaginary of $C_{\mathcal{L},r}$ (although in particular we don't have to go through Cantor space, and we can realize Δ_k as a quotient of some $\Delta_{2^\ell} = (\Delta_2)^\ell$ in a uniformly computable way). Furthermore we can arrange that each element of Δ_k is definable.

Define a formula $\rho(\bar{x}, \bar{y})$ on $\Delta_k \times \prod_{n < k} O_n$ by

$$\rho(\bar{x}, \bar{y}) = \left(\delta(x_0, y_0) + \min_{n < k} (d(x_{n+1}, y_{n+1}) + \delta(x_0, n)) \right) \downarrow 1.$$

Checking definitions gives that $\Delta_k \times \prod_{n < k} O_n / \rho$ is the required imaginary. This formula is also clearly uniformly computable. \square

B.3.3 Countable Disjoint Unions of Sorts

We already constructed the countable disjoint union imaginary in Corollary 3.4.9, but we will have to repeat some of that construction here, as we need to track computability issues.

A common trick in discrete logic is merging a finite collection of sorts by taking the disjoint union and adding unary predicates selecting out each sort. This can't be extended to infinitely many predicates without changing the category of models; by compactness there will be models with elements not in any given sort. The added flexibility of continuous logic allows us to do this with countably many sorts at once without changing the category of models. Specifically we can arrange it so that any sequence of types that ought to limit to an 'unsorted' type is shunted into a single unique overflow point. This is very similar to the emboundment method used in [BY08a] to treat unbounded metric structures.

It should be noted that if \mathcal{L} has finitely many sorts and (possibly infinitely many) predicates with uniformly bounded arity, this section can be skipped and the construction in Theorem B.3.23 will work directly.

Definition B.3.12. Let $\{O_n\}_{n<\omega}$ be a countable sequence of \mathcal{L} -sorts. For any \mathcal{L} -structure \mathfrak{A} , the (*countable*) *metric disjoint union* of $\{O_n\}_{n<\omega}$ is a metric structure with the set

$$U^{\mathfrak{A}} = \{*\} \cup \bigsqcup_{n<\omega} O_n^{\mathfrak{A}}$$

as its universe, where $*$ is a single new point. The metric disjoint union is written $\bigsqcup_{n<\omega}^* O_n$.

To define the metric on $U^{\mathfrak{A}}$, let $x, y \in O_n^{\mathfrak{A}}$, $z \in O_m^{\mathfrak{A}}$, with $n \neq m$. Then we have

$$\begin{aligned} d_U^{\mathfrak{A}}(x, y) &= 2^{-n} d_{O_n^{\mathfrak{A}}}(x, y), \\ d_U^{\mathfrak{A}}(x, z) &= |2^{-n} - 2^{-m}|, \text{ and} \\ d_U^{\mathfrak{A}}(x, *) &= 2^{-n}, \end{aligned}$$

where the other values are determined by symmetry. We will prove in Proposition B.3.13 that this defines a complete metric space.

A predicate on some $O_{n_1}^{\mathfrak{A}} \times \cdots \times O_{n_k}^{\mathfrak{A}}$ is extended to a predicate on $U^{\mathfrak{A}}$ by setting its value to 1 (i.e. ‘false’) when the input is not part of its domain.

Finally we add a distance predicate for the set $\{*\}$ (recall that we have restricted ourselves to relational languages, so we can’t use a constant). ◁

Proposition B.3.13. (i) *The countable metric disjoint union, $U = \bigsqcup_{n<\omega}^* O_n$, of a sequence $\{O_n\}_{n<\omega}$ of \mathcal{L} -sorts is well-defined, i.e. the metric given in the definition*

is actually a complete metric.

- (ii) The predicates interpreted on it are uniformly continuous. If they are Lipschitz in the original signature, they will still be Lipschitz (although possibly with a different Lipschitz constant).
- (iii) For any fixed \mathcal{L} and r , the countable metric disjoint union is isomorphic to a uniformly definable imaginary for all $\mathfrak{A} \in C_{\mathcal{L},r}$. The relevant formulas and the map of presentations $\mathfrak{A} \mapsto U^{\mathfrak{A}}$ are uniformly computable from the sequence $\{O_n\}_{n < \omega}$, the signature \mathcal{L} , and the real number r , so in particular if those are all computable, then the relevant formulas and the map of presentations are computable.
- (iv) Each O_n as a subset of U is a definable subset of U and (considering U as an imaginary sort) there is a definable bijection between O_n as a sort and O_n as a definable subset of U . The relevant formulas are uniformly definable in \mathcal{L} and computable.
- (v) For a fixed sequence $\mathcal{S} = \{O_n\}_{n < \omega}$ of \mathcal{L} -sorts with $O_0 = H$, the designated home sort, there is a signature $\mathcal{L}_{\mathcal{S}}$ and a theory $T_{\mathcal{S}}$, both uniformly computable from \mathcal{L} and \mathcal{S} , and a $\mathcal{L}_{\mathcal{S}}$ -sentence $\Xi_{\mathcal{S}}$, such that the models of $T_{\mathcal{S}} \cup \{\Xi_{\mathcal{S}} \geq r\}$ are precisely the same as reducts to the sort $\bigsqcup_{n < \omega}^* O_n$ of structures in $C_{\mathcal{L},r}$.

Proof. (i) The expression given for d clearly obeys all metric space axioms besides the triangle inequality. The only unobvious case is the one consisting of two points in some O_n and a third point in some O_m with $n \neq m$ (where we let $O_\omega = \{*\}$ with the understanding that “ $2^{-\omega}$ ” = 0). Let $x, y \in O_n$ and $z \in O_m$ with $n \neq m$. By symmetry there are only 2 cases to check:

- $d(x, y) \leq 2^{-n}$ and $d(x, z) = d(y, z) = |2^{-n} - 2^{-m}| \geq 2^{-(n \downarrow m + 1)}$, so $d(x, z) + d(y, z) \geq 2^{-(n \downarrow m)} \geq 2^{-n} \geq d(x, y)$, and in this case the triangle inequality is obeyed.
- $d(x, z) = |2^{-n} - 2^{-m}|$, so $d(x, z) \leq d(x, y) + |2^{-n} - 2^{-m}| = d(x, y) + d(y, z)$.

So the triangle inequality is obeyed in all cases.

To see that the metric space is complete, note that any Cauchy sequence is either eventually contained in some O_n or limits to $*$.

(ii) If a predicate P on sort $O_{n_1} \times \cdots \times O_{n_k}$ has modulus of uniform continuity $\Delta_P(x)$, then the corresponding predicate on U is uniformly continuous with modulus of uniform continuity

$$\Delta_P^*(x) = (\Delta_P((2^N x) \downarrow 1) \uparrow (2^{N+1} x)) \downarrow 1,$$

where $N = n_0 \uparrow \dots \uparrow n_{k-1}$. Note that if P has Lipschitz constant L , then on U it will have Lipschitz constant $2^{N+1}L$, and in particular it will still be Lipschitz.

(iii) By Lemma B.3.10, the class $C_{\mathcal{L},r}$ has a uniformly definable imaginary isometric to the metric space (X, d) where $X = \{0\} \cup \{2^{-n} : n < \omega\}$ and d is the standard metric on \mathbb{R} . Let $W = X \times \prod_{n < \omega} O_n$ be the infinitary product sort.

(iv) For any $x \in U$, $d(x, O_n) = |d(x, *) - 2^{-n}|$.

Let $Q : X \rightarrow [0, 1]$ be the natural inclusion map, which is a definable predicate on X uniformly for all members of $C_{\mathcal{L},r}$. For each n , let

$$R_n(x) = 1 \div 2^{n+1} |Q(x) - 2^{-n}|,$$

i.e. R_n is a predicate on X which takes on the value 1 at 2^{-n} and 0 everywhere else.

Now define a pseudo-metric on W by

$$\rho(\bar{x}, \bar{y}) = |Q(x_0) - Q(y_0)| + \sum_{n < \omega} 2^{-n} R_n(x_0) R_n(y_0) d_{O_n}(x_{n+1}, y_{n+1}).$$

Although in principle this is $[0, 2]$ -valued, by construction it will only take on values in $[0, 1]$. Taking the quotient W/ρ will identify any two elements $\bar{a}, \bar{b} \in W$ if and only if $a_0 = b_0$ and either $a_0 = 0$ or $a_0 = 2^{-n}$ and $a_n = b_n$. So by making the identification of elements of the form $(2^{-n}, \dots, a_n, \dots)$ with a_n and elements of the form $(0, \dots)$ with $*$, we get a bijection between W/ρ and U , and by checking definitions we see that ρ induces the correct metric on U .

(v) The signature $\mathcal{L}_{\mathcal{J}}$ has a single sort and the same predicate symbols as \mathcal{L} with the same total arity along with a single new unary predicate symbol Q . For each predicate symbol P , the syntactic modulus of continuity is $(\Delta_P((2^N x) \downarrow 1) \uparrow (2^{N+1} x)) \downarrow 1$, where Δ_P is the syntactic modulus of continuity of P in \mathcal{L} , and $\Delta_Q(x) = x$.

Recall that for a given formula φ there is a closed sentence determined by φ that holds exactly when φ is the distance predicate of a definable set. (Specifically, the closed sentence given by unpacking the definition of $\text{DEF}x\varphi$.) Since these axioms are fixed sentences depending on φ which are ‘restricted relative to φ ,’ they are uniformly computable from φ .³ $T_{\mathcal{J}}$ has axioms saying that Q is a distance predicate and that the set defined by Q is non-empty and has diameter 0. Let $*$ be a constant referring to the unique point defined by Q . (We add this constant in order to make the following axioms easier to write down, but it is not strictly necessary.)

Let $f : [0, 1] \rightarrow [0, 1]$ be a total computable continuous function whose zeroset is

³Note that the procedure for translating a finitary topological formula to a condition involving a real valued formula is clearly uniformly computable.

precisely $X = \{0\} \cup \{2^{-n} : n < \omega\}$. $T_{\mathcal{L}}$ has the axioms

- $\forall x f(d(x, *)) = 0$ and
- $\exists x d(x, *) = r$, for each $r \in X$.

For each $r \in X$, $T_{\mathcal{L}}$ has the axiom $\text{DEF}_x |d(x, *) - r|$, stating that $|d(x, *) - r|$ is a distance predicate of a definable set.

By abuse of notation, label those definable sets O_n . The sentence $\Xi_{\mathcal{L}}$ is given by

$$\Xi_{\mathcal{L}} = \sup_{x, y \in O_0} d(x, y),$$

i.e. the diameter of O_0 . Finally there are axioms for each predicate symbol P and each incorrect sequence of input sorts saying that P takes on the value 1 on those inputs (which is expressible as the sorts are definable sets) and axioms stating that P obeys the modulus of continuity Δ_P relative to the rescaled metrics on the O_n . \square

The following proposition is clear by construction and in particular by part (iii) of Proposition B.3.13 above.

Proposition B.3.14. *If \mathcal{L} is a metric signature with countably many sorts and we let $U = \bigsqcup_{O \in \mathcal{S}}^* O$ be the imaginary disjoint union of all \mathcal{L} -sorts, then for all $\mathfrak{A} \in C_{\mathcal{L}, r}$, we have that \mathfrak{A} and $U^{\mathfrak{A}}$ have uniformly definable imaginary expansions which are uniformly interdefinable.*

Aside from the issue of topological dimension and continuous degrees of points in the structure discussed in the introduction, one of the mild technical advantages of a countable metric disjoint union over an ω -product is that parameters in non-trivial

ω -products tend to be poorly behaved in that they act like countable collections of parameters rather than finite collections of parameters. This general phenomenon of single parameters acting like countable collections of parameters can be blamed for many of the pathologies in continuous logic (e.g. pairs ab such that $\text{tp}(ab)$ is principal but $\text{tp}(a/b)$ is not, theories with exactly two separable models, small theories with only ‘approximately ω -saturated’ separable models, and ω -categorical theories which fail to be ω -categorical after naming an element). In [BYU07], Usvyatsov and Ben Yaacov introduced the notion of a d -finite type, which intuitively speaking characterizes when a finitary type actually behaves like a discrete finitary type, rather than a discrete ω -type. Uniform d -finiteness is a technical strengthening of d -finiteness that was needed in an analog of Lachlan’s theorem on the number of countable models of a superstable theory in [BYU07].

- Proposition B.3.15.** (i) Let $\bar{a} \in \bigsqcup_{n < \omega}^* O_n$ be an ℓ -tuple of elements not equal to $*$. For any set B of parameters, $\text{tp}(\bar{a}/B)$ is (uniformly) d -finite as a type in the correct product sort if and only if it is (uniformly) d -finite as a type in the sort $(\bigsqcup_{n < \omega}^* O_n)^\ell$. (Note that since $*$ $\in \text{dcl}(\emptyset)$, its type is always uniformly d -finite and adding it to a tuple preserves d -finiteness.)
- (ii) For any (locally) compact set $B \subset O_k^{\mathfrak{A}}$, the corresponding set in $\bigsqcup_{n < \omega}^* O_n^{\mathfrak{A}}$ is (locally) compact. (Although note that the countable metric disjoint union will typically fail to be locally compact at $*$.)
- (iii) For any topologically finite dimensional (resp. weakly infinite dimensional) set $B \subset O_k^{\mathfrak{A}}$, the corresponding set in $\bigsqcup_{n < \omega}^* O_n^{\mathfrak{A}}$ is finite dimensional (resp. weakly infinite dimensional). If each $O_k^{\mathfrak{A}}$ is finite dimensional, then $\bigsqcup_{n < \omega}^* O_n^{\mathfrak{A}}$ will be either finite

dimensional or weakly infinite dimensional and locally finite dimensional away from $$. If each $O_k^{\aleph_1}$ is weakly infinite dimensional, then $\bigsqcup_{n < \omega}^* O_n^{\aleph_1}$ is as well.*

Proof. These all follow from the fact that the natural inclusion maps $O_k \rightarrow \bigsqcup_{n < \omega}^* O_n$ are open, isometric-up-to-scaling, and bijections between definable sets. \square

In particular if T is ‘hereditarily ω -categorical’ (i.e. ω -categorical over every finite set of parameters) or has an exactly ω -saturated separable model, then $\text{Th}(\bigsqcup_{n < \omega}^* O_n)$ will as well [BYU07].

B.3.4 Making Everything Lipschitz

Ultimately we will need all of our predicate symbols to be Lipschitz since they will be encoded directly into a metric and metrics are always Lipschitz. There are a couple of ways to accomplish this. If the reader does not care about computability, this section can be skipped using the following Fact B.3.16. Also it should be noted that Fact B.3.16 does not rely on the signature in question being countable, but the result that we will use, Proposition B.3.21, does in general.

Fact B.3.16. *Let (X, d) be a metric space and $f : X \rightarrow [0, 1]$ be a uniformly continuous function. For each $0 < n < \omega$, let*

$$f_n(x) = \inf_y \left(\frac{1}{n} f(y) + d(x, y) \right) \downarrow 1.$$

Then $f_n(x)$ is a sequence of 1-Lipschitz functions such that $nf_n \rightarrow f$ uniformly as $n \rightarrow \infty$.

In general the transformation in Fact B.3.16 would cost a jump to compute on a given structure, i.e. if some degree \mathbf{a} computes a structure (M, P) with predicate P , then \mathbf{a}' will compute (M, P_0, P_1, \dots) with P_n given by the formula in Fact B.3.16. So to ensure that the construction is computable, we will have to use something else. We will use the fact that if α is a concave non-decreasing function such that $\alpha(0) = 0$, then for any metric d , $\alpha(d)$ is also a metric. If one of our predicates P has a concave non-decreasing modulus of uniform continuity, then this means that we can compose it with the metric to get a uniformly equivalent metric relative to which P is 1-Lipschitz.

The following is a fairly elementary real analytic fact, but we will include a proof for the sake of demonstrating that the procedure is computable. Note that we could avoid this lemma entirely if our moduli of uniform continuity were non-decreasing and sub-additive, which is often required and can always be arranged as shown in this lemma.

Lemma B.3.17. *Let $\delta : [0, 1] \rightarrow [0, 1]$ be a continuous function satisfying $\delta(0) = 0$. There is a continuous, concave, non-decreasing function $\alpha : [0, 1] \rightarrow [0, 1]$ satisfying $\alpha(0) = 0$ and $\alpha \geq \delta$. Furthermore, α is uniformly computable from δ .*

Proof. α will be the ‘non-decreasing convex hull of δ ,’ defined by the following formula:

$$\alpha(x) = \inf\{mx + b : 0 \leq m, b, (\forall y \in [0, 1])my + b \geq \delta(y)\}.$$

For each $n < \omega$, define

$$\alpha_n(x) = \inf\{mx + b : 0 \leq m, b, (\forall k \in \{0, 1, \dots, 2^n\})m(2^{-n}k) + b \geq \delta(2^{-n}k)\}.$$

When computing α_n , the largest m necessary is at most

$$m_n = 2^n \sup_{0 \leq k < 2^n} |\delta(2^{-n}(k+1)) - \delta(2^{-n}k)|,$$

and the largest b is always at most 1, so the computation of α_n amounts to minimizing a δ -computable linear function on a δ -computable bounded polytope, so the α_n are uniformly computable in δ [Wei00, Ch. 5]. Furthermore note that since each α_n is the infimum of a family of Lipschitz functions with uniformly bounded Lipschitz coefficients, α_n is Lipschitz and in particular continuous.

Now all we need to show is that α_n converges uniformly to α with a computable modulus of uniform convergence. For computability considerations, we will need the fact that the modulus of uniform continuity of a continuous function f on $[0, 1]$ is uniformly computable from f [Wei00, Ch. 6]. Let Δ_δ be the modulus of uniform continuity of δ . By replacing Δ_δ with $\sup_{0 \leq y \leq x} \Delta_\delta(y)$ (which is uniformly computable from Δ_δ , since $[0, x]$ is effectively compact uniformly in x), we may assume that Δ_δ is non-decreasing.

Now note that for each $n < \omega$, we have the following inequality:

$$\alpha_n \leq \alpha \leq \alpha_n + 2\Delta_\delta(2^{-n}). \quad (\star)$$

To see that this inequality is true, observe that for each interval $I = [2^{-n}k, 2^{-n}(k+1)]$, we must have

$$\begin{aligned} \delta(x) &\leq [\delta(2^{-n}k) \uparrow \delta(2^{-n}(k+1))] + \Delta(2^{-n}) \\ &\leq [\delta(2^{-n}k) \downarrow \delta(2^{-n}(k+1))] + 2\Delta(2^{-n}), \end{aligned}$$

for all $x \in I$. Therefore, if $m, b \geq 0$ satisfy the requirements in the infimum defining α_n , then for all $x \in I$,

$$mx + b \geq \delta(2^{-n}k) \uparrow \delta(2^{-n}(k+1)),$$

and thus (\star) follows, so we get that $\alpha_n \rightarrow \alpha$ uniformly as $n \rightarrow \infty$, and α is continuous. Furthermore, we clearly have a uniformly computable modulus of uniform convergence, so α is uniformly computable.

Finally note that α is concave and non-decreasing by construction (these are preserved by infima) and $\alpha(0) = 0$ since for every $\varepsilon > 0$, there is an $m > 0$ such that $mx + \varepsilon \geq \delta(x)$ for all $x \in [0, 1]$ by continuity of δ . \square

So as long as we have a single modulus of continuity that all relation symbols obey, we can find an inter-definable structure with a Lipschitz signature. We can always arrange this if our signature is countable.

Definition B.3.18 (Uniform uniform continuity). (i) A family of functions $f \in F$ on a metric space X is *uniformly uniformly continuous* or *u.u.c.* if there is a single modulus of uniform continuity valid for all $f \in F$.

(ii) A metric signature \mathcal{L} is *u.u.c.* if $\Delta_P = \Delta_Q$ for all predicate symbols P and Q . \triangleleft

Recall that two metric spaces (X_0, d_0) and (X_1, d_1) are bi-uniformly isomorphic if there is a uniformly continuous bijection $f : X_0 \rightarrow X_1$ with uniformly continuous inverse. Two metrics d_0, d_1 on the same space X are uniformly equivalent if (X, d_0) and (X, d_1) are bi-uniformly isomorphic under the identity map.

Lemma B.3.19. (i) If d is a $[0, 1]$ -valued metric and $\alpha : [0, 1] \rightarrow [0, 1]$ is a continuous, concave, non-decreasing function satisfying $\alpha(0) = 0$, then $\alpha(d) \uparrow d$ is a

metric that is uniformly equivalent to d .

(ii) *If (X, d) is a metric space with diameter ≤ 1 and $f_i : X \rightarrow [0, 1]$ for $i \in I$ is a family of u.u.c. functions with continuous, sub-additive, non-decreasing modulus of uniform continuity α , then $(X, \alpha(d) \uparrow d)$ is a metric space bi-uniformly isomorphic to (X, d) , such that the family $\{f_i\}_{i \in I}$ is 1-Lipschitz.*

Proof. (i) Concave functions are sub-additive. The pseudo-metric axioms are preserved under composition with sub-additive, non-decreasing functions which fix 0, so $\alpha(d)$ is a pseudo-metric. The maximum of two pseudo-metrics is still a pseudo-metric, so $\alpha(d) \uparrow d$ is a pseudo-metric. $\alpha(d) \uparrow d = 0$ if and only if $d = 0$, so it is actually a metric. $\alpha(d) \uparrow d$ and d are clearly uniformly equivalent.

(ii) This is immediate from (i). □

In the previous lemma we only need to take the maximum with d on the off chance that $\alpha = 0$. Ultimately there is no harm in doing so.

Lemma B.3.20. *If \mathcal{L} is a countable metric signature, then it is interdefinable with a u.u.c. metric signature \mathcal{K} . Furthermore if \mathcal{L} is computable, then we can take \mathcal{K} to be uniformly computable in \mathcal{L} .*

Proof. Let $\{P_i\}_{i < \omega} = \mathcal{P}$ be an enumeration of all the predicate symbols in \mathcal{L} (in any sort). For each $i < \omega$, let Q_i be the \mathcal{L} -formula $2^{-(i+1)}P_i$. The \mathcal{L} -formulas Q_i are u.u.c. with regards to the modulus of uniform continuity $\Delta = \sum_{i < \omega} 2^{-(i+1)}\Delta_{P_i}$. If we let \mathcal{K} be a metric signature with the same sorts as \mathcal{L} and predicate symbols for the Q_i , each with $\Delta_{Q_i} = \Delta$, then \mathcal{K} is the required metric signature.

The procedure described in Lemma B.3.17 is uniformly computable, so passing from \mathcal{L} to \mathcal{K} is uniformly computable as well. □

Proposition B.3.21. (i) *If \mathcal{L} is a countable metric signature, then it is interdefinable with a 1-Lipschitz metric signature \mathcal{K} , i.e. a signature such that $\Delta_P(x) = x$ for all predicate symbols P (although not for metrics, which are necessarily 2-Lipschitz). Furthermore \mathcal{K} is uniformly computable from \mathcal{L} .*

(ii) *There is a \mathcal{K} -theory $T_{\mathcal{L}}$ such that the models of $T_{\mathcal{L}}$ are precisely the interpretations of \mathcal{L} -structures as \mathcal{K} -structures. Furthermore $T_{\mathcal{L}}$ is uniformly computable from \mathcal{L} .*

Proof. (i) Aside from what we have already outlined in this section, the only subtlety is that the passage from d to $\alpha(d) \uparrow d$ may delete some information contained in d because of ‘clipping’ wherever α is locally constant (and therefore not locally invertible). To remedy this all we need to do is add, for each sort O , a new binary 1-Lipschitz predicate symbol $P_{d,O}$ whose interpretation is $\frac{1}{2}d_O$ before running the construction in this section. This does not prevent α from clipping the metric, but we lose no information since we can recover the original metric from this predicate.

(ii) $T_{\mathcal{L}}$ just needs to express that every predicate symbol is uniformly continuous with regards to the original metrics $d_O = 2P_{d,O}$ in the appropriate way, i.e. with axioms of the form

$$\forall xy |P(x) - P(y)| \leq \Delta_P(2P_{d,O}(x, y)),$$

and analogous axioms for predicates on more than one sort. □

B.3.5 Encoding in Metric Spaces

Most of the coding tricks used in the two following constructions boil down to the fact that if X and Y are metric spaces with diameter ≤ 1 , then for any 1-Lipschitz function $f : X \times Y \rightarrow [0, 1]$, you can extend the metrics on X and Y to $X \sqcup Y$ with

$d(x, y) = 2 + f(x, y)$ for $x \in X$ and $y \in Y$. After doing this, if X and Y are definable from the metric, we can recover f from the metric alone. The other fundamentally important thing is that since our metric structures have bounded diameter, we can add points at a larger diameter to ensure that they are \emptyset -definable in terms of the metric regardless of the content of the embedded metric structure.

For the sake of simplicity and to avoid writing a large number of fractions, we will write metrics with distances that are larger than 1. To bring this into line with the $[0, 1]$ -valued metric convention established at the beginning of the paper, divide all distances by 6.

Theorem B.3.22. *(i) If \mathcal{L} is a countable metric signature, then for any $C_{\mathcal{L}, r}$, there is a uniformly definable imaginary X such that for any $\mathfrak{A} \in C_{\mathcal{L}, r}$, \mathfrak{A} and the purely metric reduct $X_0^{\mathfrak{A}} = (X^{\mathfrak{A}}, d)$ are uniformly bi-interpretable in the sense that*

- *there are uniformly definable imaginary expansions of \mathfrak{A} and $X_0^{\mathfrak{A}}$ which are uniformly interdefinable, and*
- *there are uniformly definable bijections between the sorts of \mathfrak{A} and definable subsets of $X_0^{\mathfrak{A}}$, and $X_0^{\mathfrak{A}}$ is contained in the definable closure of the images of those bijections.*

Furthermore the interpretation preserves embeddings and (uniform) d -finiteness of types. If the original structure is not strongly infinite dimensional, then the interpreted structure will also not be strongly infinite dimensional. The interpretation preserves local compactness and local finite dimensionality away from a fixed compact \emptyset -definable set of bad points.

(ii) For any countable metric signature \mathcal{L} , there is a first-order theory $T_{\mathcal{L}}$ and a sentence Ξ such that for any $r \in (0, 1]$, the class of metric spaces of the form $X_0^{\mathfrak{A}}$ for $\mathfrak{A} \in C_{\mathcal{L}, r}$ is precisely the set of models of $T_{\mathcal{L}} \cup \{r \div \Xi\}$. If \mathcal{L} is a computable signature, then $T_{\mathcal{L}}$ is computable. Ξ does not depend on \mathcal{L} and is always computable. Furthermore there are computable mappings of presentations of \mathcal{L} -structures to presentations of models of $T_{\mathcal{L}}$ and vice versa (these mappings do not depend on r).

Proof. (i) By applying Lemma B.3.8, we may assume that \mathcal{L} has a uniform arity bound of 2. By applying Propositions B.3.14 and B.3.21, we may assume that \mathcal{L} has a single sort and is 1-Lipschitz. By recasting unary predicates P as binary predicates using $P(x, y) = P(x)$, we may assume that all predicates are binary.

Let $\{P_n\}_{n < \omega}$ be an enumeration of all predicates with $P_0(x, y) = \frac{1}{2}d(x, y)$.

$X^{\mathfrak{A}}$ will have the set $A \sqcup A \times \omega \sqcup \{\infty, t\}$ as its universe, where $A \times \omega \sqcup \{\infty\}$ will be a modified countable metric disjoint union, with overflow point ∞ , and t will be a tag to keep things straight. $X^{\mathfrak{A}}$ will have the unique metric defined by

- $d(x, y) = d^{\mathfrak{A}}(x, y)$, for $x, y \in A$,
- $d(x, (y, n)) = 2 + 2^{-n-1}d^{\mathfrak{A}}(x, y)$, for $x \in A$ and $(y, n) \in A \times \omega$,
- $d(x, \infty) = 2$ for $x \in A$,
- $d(x, t) = 5$, for $x \in A$,
- $d((x, n), (y, n)) = 2^{-n}d^{\mathfrak{A}}(x, y)$, for $(x, n), (y, n) \in A \times \omega$,
- $d((x, n), (y, n + 1)) = 2^{-n-1}(1 + P_n^{\mathfrak{A}}(x, y))$, for $(x, n), (y, n + 1) \in A \times \omega$,
- $d((x, n), (y, m)) = |2^{-n} - 2^{-m}|$, for $(x, n), (y, m) \in A \times \omega$ with $|n - m| > 1$,

- $d((x, n), \infty) = 2^{-n}$,
- $d((x, n), t) = 4 + 2^{-n-1}$, for $(x, n) \in A \times \omega$, and
- $d(\infty, t) = 4$.

All of the metric space axioms except for the triangle inequality are clearly obeyed by d . If all three points are in the same copy of A then the triangle inequality is obeyed, so we only need to check mixed triples. The majority of cases are mechanical to check, but there are a handful of tight or subtle cases that we will write out explicitly. Let $x, y \in A$ and $(z, n), (w, n), (u, n+1), (v, n+1), (s, m) \in A \times \omega$ with $|n - m| > 1$, where $n, m < \omega$. Also, recall that if $n \neq m$, then $|2^{-n} - 2^{-m}| \geq 2^{-n-1}$. Here are the cases we check explicitly:

- $d(x, (z, n)) \leq 2 + 2 \cdot 2^{-n-2} \leq d(x, (u, n+1)) + d((u, n+1), (z, n))$,
- $d((z, n), (u, n+1)) \leq 2^{-n-1}(1 + P_n(w, u) + d(z, w))$
 $\leq d((z, n), (w, n)) + d((w, n), (u, n+1))$,
- $d((z, n), (w, n)) \leq 2 \cdot 2^{-n-2} \leq d((z, n), (u, n+1)) + d((u, n+1), (w, n))$,
- $d((z, n), (v, n+1)) \leq 2^{-n-1}(1 + P_n(z, u) + d(u, v))$
 $= d((z, n), (u, n+1)) + d((u, n+1), (v, n+1))$, and
- $d((z, n), t) = 4 + 2 \cdot 2^{-n-2} \leq d((z, n), (u, n+1)) + d((u, n+1), t)$.

Just as in the proof of Proposition B.3.13, let $Y = 0 \cup \{2^{-n} : n < \omega\}$, and let $Q : Y \rightarrow [0, 1]$ be the natural inclusion map, which is a definable predicate on Y . For each n , let

$$R_n(x) = 1 \div 2^{n+1} |Q(x) - 2^{-n}|,$$

and define a pseudo-metric ρ on $Y \times A$ by

$$\beta_n(\bar{x}, \bar{y}) = R_n(x_0)R_{n+1}(y_0)P_n(x_{n+1}, y_{n+2}) + R_{n+1}(x_0)R_n(y_0)P_n(y_{n+1}, x_{n+2}) \text{ and}$$

$$\rho(\bar{x}, \bar{y}) = |Q(x_0) - Q(y_0)| + \sum_{n < \omega} 2^{-n} \left(R_n(x_0)R_n(y_0)d(x_{n+1}, y_{n+1}) + \frac{1}{2}\beta_n(\bar{x}, \bar{y}) \right).$$

Then $Y \times A/\rho$ will correspond to $A \times \omega \sqcup \{\infty\}$, where ∞ is the ρ -equivalence class of any element of the form $\langle 0, x \rangle$ for $x \in A$.

Recall that an element or set is definable if there is a formula which defines its distance predicate. If we have a $\{0, 1\}$ -valued indicator function, $\varphi(x)$, for the set $\varphi^{-1}(0)$, then that is even better and we can always define the distance to the set by $d(x, \varphi^{-1}(0)) = \inf_y d(x, y) + 5\varphi(y)$ if we need it. Once a point is definable we will freely use it as a constant to make the following formulas simpler. We will either find an indicator function or a distance predicate for a given set, whichever is easier to write down (although $0 \in I$ does not have a definable indicator function, as it is not co-definable).

First note that

$$\chi(x) = \sup_y d(x, y) \downarrow (4 \div d(x, y)) \downarrow 1$$

is 0 if and only if $x = t$ and 1 otherwise, because t is the only point for which there is no x with $2 \leq d(x, t) \leq 3$. So t is definable and we can use it as a constant to define distance predicates for A and each $A \times \{n\}$:

$$d(x, A) = \inf_y d(x, y) + 2|d(y, t) - 5| \text{ and} \tag{\dagger}$$

$$d(x, A \times \{n\}) = \inf_y d(x, y) + 2|d(y, t) - (4 + 2^{-n-1})|. \tag{\dagger}$$

These formulas work because of the definition of distance to t . $|d(y, t) - 5|$ and $|d(y, t) -$

$(4 + 2^{-n-1})$ roughly give the distances to A and $A \times \{n\}$ and then the method used in the proof of Proposition 9.19 in [BYBHU08] gives an exact distance predicate.

For each $n < \omega$, there is a definable bijection from A to $A \times \{n\}$ given by

$$d(y, f_n(x)) = 2^{n+1}(d(x, y) \div 2), \quad (\circ)$$

and so for any $n < \omega$, we can define P_n on A by

$$P_n(x, y) = 2^{n+1}(d(f_n(x), f_{n+1}(y)) \div 1).$$

So X is the required uniformly definable imaginary, which clearly preserves embeddings. The interpretation preserves (uniform) d -finiteness of types, lack of strong infinite dimensionality, local compactness, and local finite dimensionality by the same argument as in the proof of Lemma B.3.15 (specifically, the inclusion maps are open isometries-up-to-scaling).

The advertised set of bad points is $\{\infty, *\} \cup \bigcup_{n < \omega} (*, n)$. Since this is a closed compact set of \emptyset -definable points, it is algebraic over \emptyset .

(ii) $T_{\mathcal{L}}$ is a theory in the language of metric spaces of diameter 5. By Lemma B.3.8 and Propositions B.3.13 and B.3.21, we only need to construct $T_{\mathcal{L}}$ in the case where \mathcal{L} has one sort and is 1-Lipschitz.

$T_{\mathcal{L}}$ contains the axiom

- $\text{FUN}x (\inf_y d(x, y) + 5\chi(y))$

saying that the formula $\inf_y d(x, y) + 5\chi(y)$ is the distance predicate of a definable singleton. Since this axiom is a fixed sentence depending on φ (which is ‘restricted

relative to φ'), it is uniformly computable from φ . Let t denote the unique element of that definable singleton for the sake of making the following axioms simpler to write down. Let $f : [0, 5] \rightarrow [0, 1]$ be a computable total continuous function whose zeroset is precisely $Z = \{4 + 2^{-n-1} : n < \omega\} \cup \{4, 5\}$. $T_{\mathcal{L}}$ has the axioms

- $\forall x f(d(x, t)) = 0$ and
- $\exists x d(x, t) = r$ for $r \in Z$.

The first says that distances to t only take on values in Z . The second says that every distance in Z comes arbitrarily close to being attained (which for isolated points in Z , i.e. everything except 4, implies that the distance is exactly attained).

$T_{\mathcal{L}}$ also has axioms

- $\text{DEF} x d(x, A)$ and
- $\text{DEF} x d(x, A \times \{n\})$,

saying that the formulas (\dagger) are distance predicates to definable sets, i.e. that A and $A \times \{n\}$ are definable sets.

We also need axioms saying that the formula (\circ) defines isometries-up-to-scaling between A and $A \times \{n\}$. We can accomplish this with

- $(\forall x \in A)(\text{DEF} y \in A \times \{n\}) 2^{n+1}(d(x, y) \div 2)$ and
- $(\forall x \in A \times \{n\})(\text{DEF} y \in A) 2^{n+1}(d(x, y) \div 2)$.

And finally we need to actually assert that this function is an isometry-up-to-scaling, which can be done with

- $(\forall x_0, x_1 \in A)(\forall y_0, y_1 \in A \times \{n\})|d(x_0, x_1) - 2^n d(y_0, y_1)| \leq 2^{n+3}[(d(x_0, y_0) \div 2) + (d(x_0, y_1) \div 2)]$.

Furthermore we need axioms enforcing the definition of d given in part (i) of this proof other than the line involving P_n (which isn't determined by $T_{\mathcal{L}}$) and lines involving ∞ (which are automatically enforced by continuity). The distances between A and $A \times \{n\}$ are already enforced by the previous axioms. We need

- $(\forall x \in A)d(x, t) = 5$,
- $(\forall x \in A \times \{n\})(\forall y \in A \times \{m\})d(x, y) = |2^{-n} - 2^{-m}|$ for $|n - m| > 1$, and
- $(\forall x \in A \times \{n\})d(x, t) = 4 + 2^{-n-1}$.

For the P_n line we just need to enforce the lower bound of 2^{-n-1} and the upper bound of 2^{-n} and to ensure that P_n (which is definable from d since we can define the sets $A \times \{n\}$) obeys the correct modulus of uniform continuity (relative to the predicate $2P_0$, since the metric itself may have lost information because of clipping). This is accomplished by

- $(\forall x \in A \times \{n\})(\forall y \in A \times \{n+1\})2^{-n-1} \leq d(x, y) \leq 2^{-n}$ and
- $(\forall x_0, x_1, y_0, y_1 \in A)|P_n(x_0, x_1) - P_n(y_0, y_1)| \leq \Delta_{P_n}(2(P_0(x_0, y_0) \uparrow P_0(x_1, y_1)))$.

For those predicate symbols that were originally unary we need axioms enforcing that P_n only depends on one input, namely

- $(\forall x, y_0, y_1 \in A)P_n(x, y_0) = P_n(x, y_1)$

for each unary P_n . The existence of ∞ and its definability are implied by these other axioms (since the $A \times \{n\}$ form a Cauchy sequence of definable sets in the Hausdorff

metric whose diameters are limiting to 0 and a Hausdorff metric limit of definable sets is definable by Proposition 2.3.8). Finally Ξ is given by

- $\Xi = \sup_{x,y \in A} d(x, y),$

which is a real valued sentence that evaluates to the diameter of the set A . □

Assuming that the signature has finitely many sorts and a uniform arity bound (but maybe infinitely many predicate symbols) we can avoid the bad points entirely, but the construction is different. It is somewhat less delicate than the construction in Theorem B.3.22, so we'll only sketch the important specifics.

Theorem B.3.23. *If \mathcal{L} is a countable metric signature with finitely many sorts and a uniform arity bound, then the result of Theorem B.3.22 holds with no bad points, i.e. the bi-interpretation preserves local compactness and finite dimensionality everywhere.*

Proof. By applying Proposition B.3.21, we may assume that \mathcal{L} is 1-Lipschitz. Let $\{O_n\}_{n < k}$ be a finite list of all base sorts, and let $\{N_n\}_{n < \ell}$ be a finite list of all finitary product sorts of the form $\prod_{O \in a(P)} O$ for some predicate symbol P . The sort X will be constructed from a graph with the following nodes:

- For each $n < k$, a main copy of the sort O_n .
- For each $n < \ell$, a copy of $N_n = \prod_{O \in a(P)} O$ along with copies of each O in $a(P)$ (with multiplicity).
- For each $n < \ell$, a copy of $I = \{0\} \cup \{2^{-s} : s < \omega\}$.

Connections between the nodes will correspond to specific relationships being encoded in the metric:

- For each main copy of O_n and each copy of O_n associated to some N_m there is an edge. Call the associated copy O'_n . The metric between $x \in O_n$ and $y \in O'_n$ will be given by $d(x, y) = 2 + d_{O_n}(x, y)$, in order to encode a definable bijection between O_n and O'_n .
- For each N_m and associated O'_n there is an edge. If O'_n is the i th factor of N_m , then the metric between $\bar{x} \in N_m$ and $y \in O'_n$ will be given by $d(\bar{x}, y) = 2 + d_{O_n}(x_i, y)$, in order to encode a definable projection from N_m to O'_n .
- For each N_m and its associated copy I_m of I there is an edge. Let $\{P_n\}_{n < \omega}$ be a list of the predicates symbols on N_m . If $\bar{x} \in N_m$ and $2^{-n} \in I$, then $d(\bar{x}, 2^{-n}) = 2 + 2^{-n}P_n(\bar{x})$ and $d(\bar{x}, 0) = 2$. (This is where it's important that the predicate symbols be 1-Lipschitz. If P_n is not 1-Lipschitz, this formula cannot define a metric).

Let all other distances be 4. Finally add a single new point t , with distances to everything else between 5 and 6 chosen to make each node of the graph have a t -definable indicator function. Then using the same kind of formula as in the proof of Theorem B.3.22, t is \emptyset -definable, so each of the nodes in the graph is definable as well.

Note that for any function $\eta(x)$ taking on 0 on some copy of I and 1 everywhere else, the formula

$$\left(\eta(x) + 8 \sup_y d(x, y) \downarrow \left(\frac{1}{2} \div d(x, y) \right) \right) \downarrow 1$$

is $\{0, 1\}$ -valued and takes on the value 0 if and only if x is 1 (as an element of I).

Therefore we can use $1 \in I$ as a constant, and for each $n \leq \omega$, we can define a distance

predicate for $2^{-n} \in I$ (with $2^{-\omega} = 0$) by

$$d(x, 2^{-n}) = \inf_y d(x, y) + 2|d(y, 1) - (1 - 2^{-n})|.$$

So each point in each copy of I is \emptyset -definable.

Every point in $X^{\mathfrak{A}}$ is either an image of some Cartesian product of sorts in \mathfrak{A} or contained in a compact clopen definable set (either a copy of I or t). Finitary products preserve local compactness and local finite dimensionality, so in this construction there are no ‘bad points.’ \square

B.4 On ‘Finite Axiomatizability’ in Continuous Logic

The notion of finite axiomatizability is somewhat awkward in continuous logic. There are several possible definitions that suggest themselves, but none of them seem useful. This is the most literal transcription of the ordinary definition:

Definition B.4.1 (Finite axiomatizability version 1). theory T is *finitely axiomatizable* if and only if it is axiomatized by a finite collection of sentences. \triangleleft

Depending on what we mean by ‘sentence,’ every theory in a countable language is finitely axiomatizable in that continuous logic naturally has an infinitary conjunction of the form $\bigwedge_{n < \omega} 2^{-n} \varphi_n$, and we can just let φ_n be an enumeration of a countable dense subset of the logical consequences of T .

A sensible attempt to avoid this would be a definition like this:

Definition B.4.2 (Finite axiomatizability version 2). A theory T is *finitely axiomatizable* if and only if it is axiomatized by a finite collection of restricted sentences. \triangleleft

But this is arbitrary and fails to have any obvious meaningful semantic consequences.

We can try a more directly semantic definition like this:

Definition B.4.3 (Finite axiomatizability version 3). A theory T is *finitely axiomatizable* if and only if the class of models of T and its complement are both elementary. \triangleleft

This amounts to saying $[T] = \{T' \in S_0(\emptyset) : T' \vdash T, T' \text{ a complete theory}\}$ is a clopen subset of $S_0(\emptyset)$. The problem is that for any reasonable⁴ metric signature, $S_0(\emptyset)$ is connected, so the only finitely axiomatizable theories are the trivial theory and the inconsistent one. That said, ‘finite axiomatizability version 3’ relative to a theory can be non-trivial.

At this point we could argue that clopenness in type space is too strong of a condition in continuous logic. Definable sets do not correspond to clopen subsets of type space, but rather have a more subtle topometric characterization in terms of the d -metric: A closed set $D \subseteq S_n(T)$ is definable if and only if $D \subseteq \text{int}\{p \in S_n(T) : d(p, D) < \varepsilon\}$ for every $\varepsilon > 0$, where $\text{int}X$ is the topological interior of X . By analogy we could try a similar weakening of clopen as a basis for our definition of ‘finite axiomatizability,’ but the d -metric relies on T being a complete theory and for a complete theory $S_0(T)$ is trivial.

There are, however, contexts in which there is a meaningful non-trivial metric on

⁴If all function symbols in \mathcal{L} have concave moduli of continuity, then $S_0(\emptyset)$ can be continuously retracted to a point by scaling all non-metric relations to 0 and then scaling the metric to 0. On the other hand, if $\Delta_f(x) = x^2$ and the metric has diameter ≤ 1 , then the sentence

$$\inf_x d(x, f(x))$$

can only take on the values 0 or 1. Either there exists some x such that $d(x, f(x)) < 1$, in which case $x, f(x), f(f(x)), \dots$ converges to a fixed point of f , or for every x , $d(x, f(x)) = 1$.

$S_0(T)$ for an incomplete theory T . Specifically if we're examining a notion of approximate isomorphism (such as the perturbations in [BY08b] or the distortion systems we presented in Chapter 6), we get a metric on completions of T :

$$\rho(T_0, T_1) = \inf\{\varepsilon : \mathfrak{A} \models T_0, \mathfrak{B} \models T_1, \mathfrak{A}, \mathfrak{B} \text{ '}\varepsilon\text{-isomorphic'}\},$$

whatever 'ε-isomorphic' might mean. And in this case we get a weaker notion of finite axiomatizability:

Definition B.4.4 (Finite axiomatizability version 4). A theory T is *finitely axiomatizable relative to ρ* if there is a sentence χ such that $T \vdash \chi$ and for all complete theories T' , $T' \vdash \chi \leq \rho(T', [T])$, where $\rho(T', [T])$ is the point-set distance between T' and $[T]$. \triangleleft

This definition is equivalent to the topometric condition $[T] \subseteq \text{int}\{T' \in S_0(\emptyset) : \rho(T', [T]) < \varepsilon\}$ for every $\varepsilon > 0$. It should be noted that this is a proper generalization of version 3 in that we can take our notion of approximate isomorphism to be \mathfrak{A} and \mathfrak{B} are 0-isomorphic if they are isomorphic and 1-isomorphic if they are not.

This may be a reasonable definition in some context, although as discussed in [BY08b] the metrics ρ are generally much more poorly behaved than the d -metric. In any case it's unclear what one can do with this definition. To apply it to the results of Section B.3 we would need to choose a notion of approximate isomorphism before we could even ask the question of whether or not the theory $T_{\mathcal{L}}$ is 'finitely axiomatizable.'

B.5 On the Definition of NSOP in Continuous Logic

Here we (technically) resolve a question of Ben Yaacov's, presented in [Ben13, Quest. 4.14] by showing that the definition of NSOP present there does not correctly generalize NSOP in discrete logic and in fact is equivalent to stability. We also clarify the relationship between the definition of NSOP presented in [Ben13, Quest. 4.14] and the definition of NSOP presented in [CT16, Def. 2.1].

We will present a definition based on the definition presented in [Ben13, Quest. 4.14], which we will prove is equivalent. We will use the term *quasi-metric* instead of continuous pre-order, as we believe this makes the intuition in the argument clearer.

Definition B.5.1 (Nonstandard⁵). Given a set X , a *quasi-metric* ρ on X is a function $\rho : X^2 \rightarrow \mathbb{R}$ satisfying $\rho(x, x) = 0$, $\rho(x, y) \geq 0$, and $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$. \triangleleft

As mentioned in [Ben13], just as with definable pseudo-metrics, there is a canonical way to turn any given formula $\varphi(\bar{x}; \bar{z})$ into a definable quasi-metric, namely $q_\varphi(\bar{x}, \bar{y}) = \sup_{\bar{z}} \varphi(\bar{x}; \bar{z}) \dot{-} \varphi(\bar{y}; \bar{z})$. q_φ is always a quasi-metric, and if $\rho(x, y)$ is a quasi-metric, then $q_\rho(x, y) = \rho(x, y)$, by the reverse triangle inequality $\rho(x, y) \geq \rho(x, z) - \rho(y, z)$ and the fact that $\rho(x, y) \geq 0$. This is analogous to the construction in discrete logic of turning any given formula into a definable pre-order by setting $\bar{x} \sqsubseteq_\varphi \bar{y} = \forall \bar{z} \varphi(\bar{x}, \bar{z}) \rightarrow \varphi(\bar{y}, \bar{z})$.

Definition B.5.2. For any $r \in [0, 1)$, a $[0, 1]$ -valued definable quasi-metric $\rho(x, y)$ has an *r -chain of length n* if there exists a sequence $\{a_i\}_{i < n}$ such that $i < j$ implies $\rho(a_i, a_j) \leq r$ and $\rho(a_j, a_i) = 1$.

$\rho(x, y)$ has an *infinite r -chain* if it has an r -chain of length ω . \triangleleft

⁵Technically it should be *pseudo-quasi-metric*, but we will be using this term frequently.

Clearly if we have a definable quasi-metric with an infinite r -chain, then it has an instance of SOP in the sense of [Ben13, Quest. 4.14]. We will first prove the converse.

Proposition B.5.3. *If a formula $\varphi(\bar{x}, \bar{y})$ has the strict order property over the theory T in the sense of [Ben13, Quest. 4.14], then there is a definable quasi-metric $\rho(\bar{x}, \bar{y})$ that has an infinite r -chain for some $r \in [0, 1)$.*

Proof. Let $\{\bar{a}_i\}_{i < \omega}$ be a sequence witnessing that φ has the strict order property in the sense of [Ben13, Quest. 4.14], i.e. assume that the following hold:⁶

- $\varphi(\bar{x}, \bar{y}) = 0$,
- $\varphi(\bar{x}, \bar{z}) \leq \varphi(\bar{x}, \bar{y}) + \varphi(\bar{y}, \bar{z})$, and
- $\sup_{n < m} \varphi(\bar{a}_n, \bar{a}_m) < \inf_{n > m} \varphi(\bar{a}_n, \bar{a}_m)$.

If $\inf_{n > m} \varphi(\bar{a}_n, \bar{a}_m) > 0$, then pick some $r > 0$ with $\sup_{n < m} \varphi(\bar{a}_n, \bar{a}_m) < r < \inf_{n > m} \varphi(\bar{a}_n, \bar{a}_m)$. Otherwise, we need to modify φ so that it is still a witness but such that it admits such an r .

There is some $\varepsilon > 0$ such that $\inf_{n > m} \varphi(\bar{a}_n, \bar{a}_m) - \sup_{n < m} \varphi(\bar{a}_n, \bar{a}_m) > \varepsilon$. By uniform continuity, this implies that there is a $\delta > 0$ such that for $n \neq m$, $d(\bar{a}_n, \bar{a}_m) > \delta$. So, by replacing $\varphi(\bar{x}, \bar{y})$ with $\varphi(\bar{x}, \bar{y}) + u d(\bar{x}, \bar{y})$ for some sufficiently large $u > 0$, we get a formula with the strict order property over T which admits such an r .

Now set $\rho(\bar{x}, \bar{y}) = \frac{1}{r}[\varphi(\bar{x}, \bar{y})]_0^r$. It is not hard to check that ρ is a quasi-metric and, furthermore, that $\{\bar{a}_i\}_{i < \omega}$ is an infinite $(1 - \varepsilon) \uparrow 0$ -chain for $\rho(\bar{x}, \bar{y})$. \square

Now we will show that the parameter r is actually important and that an instance of an r -chain for a definable quasi-metric is roughly equivalent to an instance of ‘SOP $_{1/r}$ ’ in the following sense.

⁶[Ben13] does not require that the formula be non-negative, but we will show here that this assumption is unnecessary anyways.

Definition B.5.4 ([CT16, Def. 2.1]). For any $n \geq 3$, a formula $\varphi(x, y)$ has SOP_n if there is a sequence $\{a_i\}_{i < \omega}$ and a $\varepsilon > 0$ such that $\varphi(a_i, a_j) = 0$ for all $i < j$ and

$$\forall x_0 x_1 \dots x_{n-1} \bigvee_{i < n} \varphi(x_i, x_{i+1}) \geq \varepsilon$$

holds, where $x_n = x_0$. ◁

Proposition B.5.5. *For any $n \geq 3$ and $r < \frac{1}{n-1}$ if there is a definable quasi-metric with an infinite r -chain, then there is an instance of SOP_n .*

Proof. We will write this proof for 1-tuples. The proof for n -tuples is the same.

Let $\rho(x, y)$ be a quasi-metric, and let $\{a_i\}_{i < \omega}$ an infinite r -chain for $\rho(x, y)$. Consider the formula $\varphi(x, y) = (\rho(x, y) \div r) \uparrow (1 \div \rho(y, x))$. Clearly we have $\varphi(a_i, a_j) = 0$ for any $i < j$.

Let b_0, b_1, \dots, b_{n-1} be arbitrary, let

$$\chi(\bar{b}) = \varphi(b_0, b_1) \uparrow \varphi(b_1, b_2) \uparrow \dots \uparrow \varphi(b_{n-2}, b_{n-1}) \uparrow \varphi(b_{n-1}, b_0),$$

and fix $s > 0$ such that $\chi(\bar{b}) < s$. By the triangle inequality we have that

$$\begin{aligned} \rho(b_0, b_{n-1}) &\leq \rho(b_0, b_1) + \rho(b_1, b_2) + \dots + \rho(b_{n-2}, b_{n-1}), \\ &< \underbrace{(r + s) + (r + s) + \dots + (r + s)}_{n-1} = (n-1)(r + s), \end{aligned}$$

which implies that $\chi(\bar{b}) > 1 - (n-1)(r + s)$. Whenever $s < \frac{1-(n-1)r}{n}$ we have a contradiction, and therefore we must have $\forall \bar{b} \left(\chi(\bar{b}) \geq \frac{1-(n-1)r}{n} \right)$. Since $r < \frac{1}{n-1}$, and therefore $\frac{1-(n-1)r}{n} > 0$, we have an instance of SOP_n . ◻

We also have that the existence of a definable quasi-metric with an infinite $\frac{1}{2}$ -chain is equivalent to the existence of an unstable formula.

Definition B.5.6. A formula $\phi(\bar{x}, \bar{y})$ has the *order property* if there exists an infinite sequence $\{a_i\}_{i < \omega}$ and $r < s$ such that for any $i < k$, $\varphi(a_i, a_k) \leq r$ and $\varphi(a_k, a_i) \geq s$. \triangleleft

Clearly a definable quasi-metric with an infinite r -chain for any $r \in [0, 1)$ gives an instance of an unstable formula, so we only need to show the converse for $r = \frac{1}{2}$.

The idea of the following proof is much clearer in the discrete case. If we have a $\{0, 1\}$ -valued formula $\varphi(x, y)$ with the order property in a discrete theory, then a quasi-metric defined by $\rho(x, y) = 0$ if $x = y$, $\rho(x, y) = \frac{1}{2}$ if $x \neq y$ and $\varphi(x, y)$ holds, and $\rho(x, y) = 1$ if $x \neq y$ and $\varphi(x, y)$ does not hold. This is a definable quasi-metric with an infinite $\frac{1}{2}$ -chain. The majority of the difficulty of this proof is finding the right continuous analog of ρ and verifying that it is a quasi-metric.

Proposition B.5.7. *If a theory T has a formula with the order property, then it has a definable quasi-metric with an infinite $\frac{1}{2}$ -chain.*

Proof. Let $\varphi(\bar{x}, \bar{y})$ be an unstable formula, with witnessing sequence $\{\bar{a}_i\}_{i < \omega}$. By replacing $\varphi(\bar{x}, \bar{y})$ with a sufficiently close restricted formula, we may assume that φ is restricted and therefore Lipschitz. By replacing $\varphi(\bar{x}, \bar{y})$ with $\left[\frac{\varphi(\bar{x}, \bar{y}) - s}{r - s} \right]_0^1$, we may assume that $\varphi(\bar{x}, \bar{y})$ is $[0, 1]$ -valued and that $r = 0$ and $s = 1$. Note that φ is still Lipschitz. Let it be M -Lipschitz for some $M > 0$ (i.e. for any $\bar{x}\bar{y}$ and $\bar{z}\bar{w}$, $|\varphi(\bar{x}, \bar{y}) - \varphi(\bar{z}, \bar{w})| \leq Md(\bar{x}\bar{y}, \bar{z}\bar{w})$, where the metric on tuples is the max metric). By increasing M if necessary, we may assume that $M > 1$.

By uniform continuity of $\varphi(\bar{x}, \bar{y})$, there must be some $\varepsilon > 0$ such that for any $i \neq k$, $d(\bar{a}_i, \bar{a}_k) > \varepsilon$. By decreasing ε if necessary, we may assume that $\varepsilon < 1$. Consider the

formulas

$$\delta(\bar{x}, \bar{y}) = \left[\frac{M}{\varepsilon} d(\bar{x}, \bar{y}) \right]_0^1 \quad \text{and}$$

$$\rho(\bar{x}, \bar{y}) = [\delta(\bar{x}, \bar{y})]_0^{1/2} + \left(\delta(\bar{x}, \bar{y}) \div \frac{1}{2} \right) \downarrow \left(\frac{1}{2} \varphi(\bar{x}, \bar{y}) \right).$$

Note that when $\varphi(\bar{x}, \bar{y}) = 1$, $\rho(\bar{x}, \bar{y}) = \delta(\bar{x}, \bar{y})$, and, more generally, we always have $\rho(\bar{x}, \bar{y}) \leq \delta(\bar{x}, \bar{y})$.

We need to argue that $\rho(\bar{x}, \bar{y})$ is a quasi-metric. It clearly satisfies $\rho(\bar{x}, \bar{x}) = 0$ and $\rho(\bar{x}, \bar{y}) \geq 0$, so we only need to verify the triangle inequality.

Given \bar{x} , \bar{y} , and \bar{z} , if both $\delta(\bar{x}, \bar{y}) < \frac{1}{2}$ and $\delta(\bar{y}, \bar{z}) < \frac{1}{2}$ hold, then $\rho(\bar{x}, \bar{y}) = \delta(\bar{x}, \bar{y})$ and $\rho(\bar{y}, \bar{z}) = \delta(\bar{y}, \bar{z})$, and, since δ is a metric, we get that

$$\rho(\bar{x}, \bar{z}) \leq \delta(\bar{x}, \bar{z}) \leq \delta(\bar{x}, \bar{y}) + \delta(\bar{y}, \bar{z}) = \rho(\bar{x}, \bar{y}).$$

If both $\rho(\bar{x}, \bar{y}) \geq \frac{1}{2}$ and $\rho(\bar{y}, \bar{z}) \geq \frac{1}{2}$ hold, then we have that

$$\rho(\bar{x}, \bar{z}) \leq 1 = \frac{1}{2} + \frac{1}{2} \leq \rho(\bar{x}, \bar{y}) + \rho(\bar{y}, \bar{z}).$$

If $\rho(\bar{x}, \bar{y}) < \frac{1}{2}$ and $\rho(\bar{y}, \bar{z}) \geq \frac{1}{2}$ both hold, then we have that $\rho(\bar{x}, \bar{y}) = \delta(\bar{x}, \bar{y})$ and that $\delta(\bar{y}, \bar{z}) \geq \frac{1}{2}$, implying that $\rho(\bar{y}, \bar{z}) = \frac{1}{2} + (\delta(\bar{y}, \bar{z}) \div \frac{1}{2}) \downarrow (\frac{1}{2} \varphi(\bar{y}, \bar{z}))$.

If $\delta(\bar{x}, \bar{z}) < \frac{1}{2}$, then $\rho(\bar{x}, \bar{z}) < \frac{1}{2} \leq \rho(\bar{y}, \bar{z})$, and we're done, so assume that $\delta(\bar{x}, \bar{z}) \geq \frac{1}{2}$. Now we have that $\rho(\bar{x}, \bar{z}) = \frac{1}{2} + (\delta(\bar{x}, \bar{z}) \div \frac{1}{2}) \downarrow (\frac{1}{2} \varphi(\bar{x}, \bar{z}))$.

It is not hard to check that the quantity

$$\begin{aligned} & \left| \left(\frac{1}{2} + \left(\delta(\bar{x}, \bar{z}) \dot{\div} \frac{1}{2} \right) \downarrow \left(\frac{1}{2} \varphi(\bar{x}, \bar{z}) \right) \right) - \left(\frac{1}{2} + \left(\delta(\bar{y}, \bar{z}) \dot{\div} \frac{1}{2} \right) \downarrow \left(\frac{1}{2} \varphi(\bar{y}, \bar{z}) \right) \right) \right| \\ &= \left| \left(\delta(\bar{x}, \bar{z}) \dot{\div} \frac{1}{2} \right) \downarrow \left(\frac{1}{2} \varphi(\bar{x}, \bar{z}) \right) - \left(\delta(\bar{y}, \bar{z}) \dot{\div} \frac{1}{2} \right) \downarrow \left(\frac{1}{2} \varphi(\bar{y}, \bar{z}) \right) \right|, \end{aligned}$$

which is the same as $|\rho(\bar{x}, \bar{z}) - \rho(\bar{y}, \bar{z})|$, is no greater than

$$|\delta(\bar{x}, \bar{z}) - \delta(\bar{y}, \bar{z})| \uparrow \left| \frac{1}{2} \varphi(\bar{x}, \bar{z}) - \frac{1}{2} \varphi(\bar{y}, \bar{z}) \right|,$$

so since

$$\begin{aligned} |\delta(\bar{x}, \bar{z}) - \delta(\bar{y}, \bar{z})| &\leq \delta(\bar{x}, \bar{y}) = \rho(\bar{x}, \bar{y}) \text{ and} \\ |\varphi(\bar{x}, \bar{z}) - \varphi(\bar{y}, \bar{z})| &\leq Md(\bar{x}, \bar{y}) \downarrow 1 \leq \delta(\bar{x}, \bar{y}) = \rho(\bar{x}, \bar{y}), \end{aligned}$$

we have that $|\rho(\bar{x}, \bar{z}) - \rho(\bar{y}, \bar{z})| \leq \rho(\bar{x}, \bar{y})$, and so $\rho(\bar{x}, \bar{z}) \leq \rho(\bar{x}, \bar{y}) + \rho(\bar{y}, \bar{z})$.

The proof in the case where $\rho(\bar{x}, \bar{y}) \geq \frac{1}{2}$ and $\rho(\bar{y}, \bar{z}) < \frac{1}{2}$ is largely the same, and so $\rho(\bar{x}, \bar{y})$ is a quasi-metric. By construction we also have that for any $i < k$, $\delta(\bar{a}_i, \bar{a}_k) = 1$, and so $\rho(\bar{a}_i, \bar{a}_k) = \frac{1}{2} + \frac{1}{2} \varphi(\bar{a}_i, \bar{a}_k)$, implying that $\rho(\bar{a}_i, \bar{a}_k) = \frac{1}{2}$ and $\rho(\bar{a}_k, \bar{a}_i) = 1$. Hence, $\{\bar{a}_i\}_{i < \omega}$ is an infinite $\frac{1}{2}$ -chain for ρ , as required. \square

We have just established that the notion of SOP defined in [Ben13, Quest. 4.14] is equivalent to instability, which would seem to indicate that it is not quite the correct notion of NSOP in continuous logic. Literally interpreting the question [Ben13, Quest. 4.14], we get a negative answer for both part, as any randomization of an unstable NIP theory, as well as any randomization of an unstable NSOP theory, will have SOP as it

is defined in that question.

We will now give a more precise characterization of the relationship between SOP_n and the existence of infinite r -chains. The proof of the following theorem is a continuous modification of the proof of Proposition B.5.9, which is considerably simpler.

Proposition B.5.8. *For any $n \geq 3$, if there is an instance of SOP_n , then for any $r > \frac{1}{n}$, there is a definable quasi-metric with an infinite r -chain.*

Proof. We will write this proof for 1-tuples. Nothing changes in the n -tuple case.

Fix $r > \frac{1}{n}$. Let $\varphi(x, y)$ be the formula witnessing SOP_n . (By replacing $\varphi(x, y)$ with $\frac{1}{\varepsilon}\varphi(x, y) \downarrow 1$ we may assume that the ε in the definition of SOP_n is 1.)

We will construct a definable quasi-metric $\rho(x, y)$, which we will give informally first and then write down the explicit formula.

Informal Description. Find a $\delta > 0$ small enough that $\delta < 1$ and if $d(x, z) \leq \delta$ and $d(y, w) \leq \delta$, then $|\varphi(x, y) - \varphi(z, w)| < 1 - \frac{1}{nr}$.

To compute $\rho(x, y)$ we consider sequences of moves that cost us a certain amount. $\rho(x, y)$ will be the infimal cost (capped at 1):

- On the zeroth move we may jump from x to any z_0 at a cost of $\frac{1}{\delta}d(x, z_0)$.
- On move n , given z_{n-1} we may jump from z_{n-1} to any w_n at a cost of $\frac{1}{n} + \varphi(z_{n-1}, w_n)$ and then jump from w_n to any z_n at a cost of $\frac{1}{\delta}d(w_n, z_n)$.
- We can stop once we are at y at the end of a move. Alternatively if we have spent more than 1, we can just quit and take a total cost of 1.

End of Informal Description.

As a formula what we have is

$$\rho(x, y) = 1 \downarrow \frac{1}{\delta} d(x, y) \downarrow \inf_{\bar{z}, \bar{w}} \min_{0 < k < n} \left(\frac{k}{n} + \frac{1}{\delta} d(x, z_0) + \frac{1}{\delta} d(w_k, y) + \sum_{i=1}^k \varphi(z_{i-1}, w_i) + \frac{1}{\delta} d(w_i, z_i) \right).$$

To see that this is a quasi-metric all we need to do is verify that $\rho(x, x) = 0$, $\rho(x, y) \geq 0$, and $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$. $\rho(x, x) = 0$ and $\rho(x, y) \geq 0$ are clear. It is also clear that ρ is $[0, 1]$ -valued. So all we need to verify is the triangle inequality.

Assume that $\rho(a, b) < \alpha$ and $\rho(b, c) < \beta$. Let $a, z_0, w_1, z_1, \dots, w_k, b$ be a chain witnessing that $\rho(a, b) < \alpha$, and let $b, u_0, v_1, u_1, \dots, v_\ell, c$ be a chain witnessing that $\rho(b, c) < \beta$. I claim that

$$a, z_0, w_1, z_1, \dots, w_k, u_0, v_1, u_1, \dots, v_\ell, c$$

is a chain of length $k + \ell$ (i.e. a chain corresponding to a sequence of $k + \ell$ many moves) witnessing that $d(a, c) \leq \alpha + \beta$. If $k + \ell \geq n$, then $\alpha + \beta \geq 1$ anyways, so the inequality $d(a, c) \leq 1$ holds trivially. Otherwise if $k + \ell < n$, then we can fit these choices of the variables into the formula for $\rho(x, y)$. Note that $\frac{1}{\delta} d(w_k, u_0) \leq \frac{1}{\delta} d(w_k, b) + \frac{1}{\delta} d(b, u_0)$, so we get a term in the minimum in the definition of $\rho(x, y)$ such that for some choice of variables we have a chain with cost $< \alpha + \beta$ and we get $\rho(x, y) < \alpha + \beta$ as required.

Therefore we have $\rho(a, c) \leq \rho(a, b) + \rho(b, c)$ for all a, b, c , and $\rho(x, y)$ is a $[0, 1]$ -valued definable quasi-metric.

Now to see that we get an infinite r -chain, let $\{a_i\}_{i < \omega}$ be a sequence witnessing that $\varphi(x, y)$ is an instance of SOP_n . Assume that $\{a_i\}_{i < \omega}$ is indiscernible. Fix $i < j$, and

assume that $j > i + n$.

First we need to compute $\rho(a_i, a_j)$. Since $\varphi(a_i, a_j) = 0$, we can make the following sequence of moves, a_i, a_i, a_j, a_j , and we have that $\varphi(a_i, a_j) \leq \frac{1}{n}$.

Now we need to compute $\rho(a_j, a_i)$. Assume that $\rho(a_j, a_i) < 1 - \sigma$ for some $\sigma > 0$. Let $a_j, z_0, w_0, \dots, w_k, a_i$ be a chain witnessing $\rho(a_j, a_i) < 1 - \sigma$ (note that we must have $k < n$). Clearly we must have

$$d(a_j, z_0), d(w_1, z_1), \dots, d(w_k, a_i) < \delta(1 - \sigma) < \delta,$$

so by construction we must also have

$$\begin{aligned} |\varphi(z_0, w_1) - \varphi(a_j, w_1)| &< 1 - \frac{1}{nr} \text{ and} \\ |\varphi(z_{\ell-1}, w_\ell) - \varphi(w_{\ell-1}, w_\ell)| &< 1 - \frac{1}{nr}. \end{aligned}$$

Case 1: $\ell = n - 1$. Consider the values

$$\varphi(a_j, w_1), \varphi(w_1, w_2), \dots, \varphi(w_{\ell-1}, w_\ell), \varphi(w_\ell, a_i), \varphi(a_i, a_j).$$

Since $\varphi(x, y)$ is an instance of SOP_n , one of these must be equal to 1. We already know it can't be $\varphi(a_i, a_j)$, so it must be one of the others. This implies that one of the terms in the evaluation of $\rho(a_j, a_i)$ is $\geq 1 - (1 - \frac{1}{nr}) = \frac{1}{nr}$. So if $\sigma \leq 1 - \frac{1}{nr}$ we have a contradiction. Therefore $\rho(a_j, a_i) \geq \frac{1}{nr}$.

Case 2: $\ell < n - 1$. Consider the values

$$\varphi(a_j, w_1), \varphi(w_1, w_2), \dots, \varphi(w_{\ell-1}, w_\ell), \varphi(w_\ell, a_i), \varphi(a_i, a_{i+1}), \dots, \varphi(a_{i+n-\ell-1}, a_j).$$

The argument is essentially the same as in Case 1, just noting that all of the terms like $\varphi(a_{i+m}, a_{i+m+1})$ as well as $\varphi(a_{i+n-\ell-1}, a_j)$ must be $\leq \frac{1}{n}$. So once again we have that $\rho(a_j, a_i) \geq \frac{1}{nr}$.

So now we can consider the quasi-metric $\rho'(x, y) = nr\rho(x, y) \downarrow 1$, and we get that $\rho'(x, y)$ has an infinite r -chain. \square

The picture we arrive at is this. For any theory T , we can consider the set $C(T) = \{r \in [0, 1) : T \text{ has an inf. } r\text{-chain}\}$. By definition, this set is closed upwards.⁷ Combining Proposition B.5.5 and B.5.8, we get that $C(T)$ is always either

- \emptyset ,
- $[\frac{1}{n}, 1)$ for some $n \geq 2$,
- $[0, 1)$,
- $(\frac{1}{n}, 1)$ for some $n \geq 3$,
- $(0, 1)$, or

and that for any $n \geq 3$, $C(T) \supseteq (\frac{1}{n}, 1)$ if and only if T has an instance of SOP_n . Proposition B.5.7 tells us that $C(T) \supseteq [\frac{1}{2}, 1)$ if and only if T is unstable.⁸

For discrete theories we can get a more precise picture.

Proposition B.5.9. *If T is a discrete theory with an instance of SOP_n , then T has a definable quasi-metric with an infinite $\frac{1}{n}$ -chain.*

Proof. Let $\varphi(x, y)$ be a $\{0, 1\}$ -valued formula with SOP_n , i.e. one for which there exists a sequence $\{a_i\}_{i < \omega}$ with $\varphi(a_i, a_k) = 0$ for all $i < k$ but also for which $\{\varphi(x_0, x_1), \dots, \varphi(x_{n-2}, x_{n-1}), \varphi(x_{n-1}, x_0)\}$ is inconsistent.

⁷Although it is not hard to show that if you make the definition more restrictive and require that $\rho(\bar{x}, \bar{y}) = r$ instead of $\rho(\bar{x}, \bar{y}) \leq r$, the analog of $C(T)$ is still upwards closed. The easiest way to see this is to notice that for any $s > 0$, $\rho(x, y) \uparrow s\rho(y, x)$ is also a quasi-metric.

⁸Which arguably suggests that the ‘correct’ notion of NSOP_2 is just stability.

Let $\rho(x, y) = \frac{1}{n}L(x, y) \downarrow 1$, where $L(x, y)$ is the shortest length of a directed path from x to y using φ as an edge relation. Note that $\rho(x, y)$ is a formula despite the fact that $L(x, y)$ is not a formula.

$\rho(x, y)$ is clearly a quasi-metric. Now we have that for any $i < k$, $\rho(a_i, a_k) = \frac{1}{n}$, but $\rho(a_k, a_i) = 1$, so $\rho(x, y)$ has an infinite $\frac{1}{n}$ -chain. \square

It is also very easy to construct an example of a theory T with $C(T) = [0, 1)$ (the theory of $(\mathbb{Q}, <)$, for instance). This raises the following questions.

Question B.5.10. *For each $n \geq 3$, does there exist a (necessarily continuous) theory T such that $C(T) = (\frac{1}{n}, 1)$?*

Does there exist a theory T such that $C(T) = (0, 1)$? If there does, can such a theory be discrete?

Question B.5.11. *What is the relationship between theories with $C(T) = (0, 1)$ or $[0, 1)$ and SOP_∞ , as defined in [CT16]? What about SOP for discrete theories?*

Appendix C

Counterexamples in Continuous Logic

Here we collect some of the myriad counterexamples that exist in continuous logic.

Counterexample C.0.1. *For an example of a metric space where the metric entropy is not saturated consider the graph metric on the ‘complete ω -partite graph $K_{1,2,3,\dots}$,’ i.e. a metric space which is the disjoint union of sets X_n for $n < \omega$ satisfying $|X_n| = n$ such that the distance between any two points in the same set is 2 and the distance between any two points in different sets is 1. This space has $(\geq \frac{3}{2})$ -separated sets (which are also $(> \frac{3}{2})$ -separated) of size n for any $n < \omega$, specifically X_n , but it does not have one of size ω . This example can clearly be generalized to any limit cardinal.*

C.1 Topometry and Definability

Counterexample C.1.1. *A closed formula that is the graph of a function in some model of a theory, but not in others.*

Proof. Let \mathfrak{M} be a structure whose underlying metric space is \mathbb{R} with the metric $d(x, y) = |x - y| \downarrow 1$. Let the only non-metric predicate symbol in \mathfrak{M} ’s language be a binary 1-Lipschitz $[0, 1]$ -valued predicate, P . Let P be the distance predicate of the set defined

by $y = \sin(1/x)$. The set $\llbracket P(\mathfrak{M}) = 0 \rrbracket$ is of course the graph of a function, but it will fail to be the graph of a function in some elementary extension of \mathfrak{M} , because it is consistent that $P(x, y) = 0$ for all $y \in [-1, 1]$ for some x . \square

Counterexample C.1.2. *A definable partial function that does not extend to a definable total function on any definable set containing its domain.*¹

Proof. Let \mathfrak{M} be a two-sorted structure whose two sorts both have the underlying set $2 \times \omega^2$, and the metric $d((i, j, k), (\ell, m, n))$ defined by

- $d((i, j, k), (\ell, m, n)) = 2$ if $j \neq m$ or $k \neq n$ and
- $d((0, j, k), (1, j, n)) = 2^{-j}$.

Finally let P be a binary $\{0, 1\}$ -valued predicate with one variable in each sort such that $P((i, j, k), (\ell, m, n)) = 0$ if and only if $j = m$ and $k = n$, and let Q be a unary predicate on the first sort such that $Q((i, j, k)) = 2^{-j}$. (Q is actually definable already, but it's easier to just put it in by hand.)

Let T be the theory of this structure. If we look at the space of 1-types (in either sort) there is a unique type corresponding to an element which is distance 2 away from everything else in the sort. Restricted to this type, P defines a function (to the corresponding type in the other sort).

P cannot be extended to a definable function on the entire sort, however. To see this, note that for any set of parameters A that such a function, call it f , might be definable over, all but a small number of pairs bc satisfying $0 < d(b, c) \leq 1$ will have $b \equiv_A c$, and likewise for such pairs in the other sort. This implies that f cannot map such elements

¹This answers a question raised in [BY10a] immediately after Lemma 1.23.

in the first sort to such elements in the second sort, and so for each $k < \omega$, the range of f on the (definable) set of pairs bc satisfying $d(b, c) = 2^{-k}$ must be bounded and therefore compact and therefore finite.

Any type of an element satisfying $Q = 0$ also satisfies the following open formula

$$U(x) = \forall y (Q(x) = Q(y) \wedge d(x, y) = 2 \rightarrow d(f(x), f(y)) = 2).$$

Since this is an open formula, there must be some $\varepsilon > 0$ such that $Q(x) \leq \varepsilon \rightarrow U(x)$, but this contradicts the fact that for any $\delta > 0$, the image of f on elements satisfying $Q(x) = \delta$ is finite.

To see that we cannot do this even when we only require that f be defined on some definable set containing $\llbracket Q = 0 \rrbracket$, assume that D is a definable set over the parameters A such that we can extend P to a definable function f on all of D . Let p be a non-forking extension of the \emptyset -type axiomatized by $\{Q(x) = 0\}$, i.e. the type $p(x)$ which says that $Q(x) = 0$ and x is distance 2 from everything in A in the first sort and everything related by P to something in A in the second sort. This type is metrically isolated, so in order for D to contain it, it must contain an open neighborhood of it, but this implies that for sufficiently large $k < \omega$, the non-forking extensions of the types axiomatized by $\{Q(x) = 2^{-k}\}$ are also contained in D , and then we get the same contradiction as before. \square

Counterexample C.1.3. *A closed subset of a type space on which the metric is not open.*

Proof. Let $S_1(\text{PS})$ be the polarized square. Consider the closed set $F = \{0\} \times [0, 1] \cup [0, 1] \times \{0\}$ (oriented so that the interval $\{0\} \times [0, 1]$ is metrically compact). The metric

on this set fails to be open, because for any $\varepsilon > 0$ with $\varepsilon < 1$, $(\{0\} \times (0, 1])^{<\varepsilon}$ fails to be open. \square

Counterexample C.1.4. *Two sorts with $d_{dGH} = 0$, but no definable bijection.*

Proof. Let X and Y be two bounded, non-compact metric spaces with $d_{GH}(X, Y) = 0$, but with X and Y not homeomorphic topologically. Construct a two-sorted structure whose two sorts are X and Y and include predicate symbols for every function $\rho(x:X, y:Y)$ which is the restriction of a pseudo-metric on $X \sqcup Y$ extending d^X and d^Y . Clearly we have that $d_{dGH}(X, Y) = 0$ (thinking of X and Y as sorts), but there can be no definable bijection between X and Y because there is no bi-continuous bijection between X and Y . \square

Note though, that in any such theory, in sufficiently saturated models, the two sorts will be isometric. This raises the following question, which probably has a negative answer.

Question C.1.5. *Is there a theory T with two sorts with $d_{dGH} = 0$ and between which there are no definable bijections, even with parameters?*

Counterexample C.1.6. *A theory T for which \aleph_0 -saturation implies \aleph_1 -saturation.*

Proof. Take the theory of the sort of ω -tuples in any given theory. \square

Counterexample C.1.7 (Fig. 13). *Dictionary type space with a $[0, 1]$ -formula φ such that for every $r \in [0, 1]$, $\llbracket \varphi = r \rrbracket$ is non-empty but fails to be definable.*

Proof. Let \mathcal{L} be a language with two 1-Lipschitz unary predicate symbols, P and Q , and let \mathfrak{M} be an \mathcal{L} -structure whose universe is $\bigsqcup_{p \text{ prime}} \{1, \dots, p-1\}$. Let the elements




Figure 13: $\llbracket Q = r \rrbracket$ is not definable for any $r \in [0, 1]$

be represented by ordered pairs $\langle p, i \rangle$, where p is a prime and $i < p$ is some positive natural number. For any $\langle p, i \rangle$, let $P(\langle p, i \rangle) = \frac{1}{n}$, where p is the n th prime, and let $Q(\langle p, i \rangle) = \frac{i}{p}$. For any $\langle p, i \rangle$ and $\langle q, k \rangle$, if $p \neq q$, then let $d(\langle p, i \rangle, \langle q, k \rangle) = 1$, otherwise let $d(\langle p, i \rangle, \langle q, k \rangle) = \frac{|i-k|}{p}$.

If we let T be the theory of \mathfrak{M} , then it is not hard to check that $S_1(T)$ is the topological space depicted in Figure 13, which can be thought of as a subset of $[0, 1]^2$ consisting of those points of the form $\langle P(\langle p, i \rangle), Q(\langle p, i \rangle) \rangle$, for $\langle p, i \rangle \in \mathfrak{M}$, together with the points of the form $\langle 0, r \rangle$ for $r \in [0, 1]$. It is also not hard to check that the d -metric on this type space is $d(\langle x, y \rangle, \langle z, w \rangle) = 1$ if $x \neq z$ and $d(\langle x, y \rangle, \langle z, w \rangle) = |y - w|$ otherwise (where $\langle x, y \rangle$ and $\langle z, w \rangle$ are elements of the described subspace of $[0, 1]^2$).

For any axis-aligned closed rectangle in \mathbb{R}^2 with corners with irrational coordinates, the intersection with $S_1(T) \subset [0, 1]^2$ is definable, so $S_1(T)$ is dictionaric. (To see this, note that if the rectangle does not intersect $\{0\} \times [0, 1]$, then the resulting subset of $S_1(T)$ is clopen, and in the other case it is not too hard to check that the resulting set is definable.)




Figure 14: Connected \mathbb{R}^2 -embeddable dictionary type space with no d -atomic points

$\llbracket Q = r \rrbracket$ is clearly non-empty for any $r \in [0, 1]$. To see that it is not definable, note that for any $r \in [0, 1]$, $\llbracket Q = r \rrbracket$ contains the type corresponding to $\langle 0, r \rangle$ and at most finitely many other types (this is why we needed to index by the primes). For sufficiently small $\varepsilon > 0$, this implies that $\llbracket Q = r \rrbracket^{<\varepsilon}$ has, as a relatively clopen subset, an open ε -ball of the type corresponding to $\langle 0, r \rangle$, which means that $\llbracket Q = r \rrbracket^{<\varepsilon}$ is not open, and thus $\llbracket Q = r \rrbracket$ is not definable. \square

It is possible to prove that Counterexample C.1.7 is ω -stable (and, moreover, interpretable in the discrete theory of an infinite set).

Counterexample C.1.8 (Fig. 14). *A \mathbb{R}^2 -embeddable, topologically 1-dimension, connected, dictionary type space containing no d -atomic points.*

Proof. Let \mathcal{L} be a language with a countable sequence of 1-Lipschitz $[0, 1]$ -valued predicate symbols $\{P_i\}_{i < \omega}$. Let \mathfrak{M} be a structure whose underlying set is $\mathbb{R}_{\geq 0} \times 2^\omega$, with the metric $d(\langle r, \alpha \rangle, \langle s, \beta \rangle) = 1$ if $\alpha \neq \beta$ and $d(\langle r, \alpha \rangle, \langle s, \alpha \rangle) = |r - s| \downarrow 1$ and with $P_i^{\mathfrak{M}}(\langle r, \alpha \rangle) = \alpha(i)(1 \div r)$. Let T be $\text{Th}(\mathfrak{M})$.

It is not hard to show that $S_1(T)$ is homeomorphic to $[0, 1] \times 2^\omega$ with the set $\{0\} \times 2^\omega$ collapsed to a point. This can be embedded in \mathbb{R}^2 (as seen in Figure 14). It is also easy to see that $S_1(T)$ has a basis of open sets whose boundaries are each homeomorphic to Cantor space, so since $S_1(T)$ is a compact metrizable space, it has topological dimension no more than 1. Since $S_1(T)$ has more than one point and is not totally disconnected, it cannot have topological dimension -1 or 0 , so $\dim S_1(T) = 1$.

If we write the points of $S_1(T)$ as $\langle r, \alpha \rangle$ for $r \in \mathbb{R}_{\geq 0}$ and $\alpha \in 2^\omega$ (corresponding to the type of the same element in \mathfrak{M}) with a single extra point ∞ , the d -metric on $S_1(T)$ is given by

- $d(\langle r, \alpha \rangle, \langle s, \beta \rangle) = 1$, if $\alpha \neq \beta$,
- $d(\langle r, \alpha \rangle, \langle s, \alpha \rangle) = |r - s| \downarrow 1$, and
- $d(\langle r, \alpha \rangle, \infty) = 1$.

This implies that the sets

1. $\{\infty\} \cup [r, \infty) \times 2^\alpha$, for any $r \in \mathbb{R}_{\geq 0}$, and
2. $[r, s] \times Q$, for any r and s with $0 \leq r \leq s$ and any clopen $Q \subseteq 2^\alpha$,

are all definable, so $S_1(T)$ has a basis of definable neighborhoods and is therefore dictionaric.

Finally, sufficiently small balls around ∞ are just $\{\infty\}$, so it is not d -atomic, and for any other $\langle r, \alpha \rangle \in S_1(T)$, sufficiently small balls around it only contain points of the form $\langle s, \alpha \rangle$ and so cannot be d -atomic either. Thus, no points in $S_1(T)$ are d -atomic. \square




Figure 15: Dictionary type space with a closed subspace that fails to be dictionary

Counterexample C.1.9 (Fig. 15). *A closed subset of a dictionary type space that is not itself dictionary.*

Proof. Let \mathcal{L} be a language with two 1-Lipschitz unary predicate symbols, P and Q , and let \mathfrak{M} be a structure whose universe is $\omega \times [0, 1]$, with $P(\langle n, r \rangle) = \frac{1}{1+n}$ and $Q(\langle n, r \rangle) = r$ and with $d(\langle n, r \rangle, \langle k, s \rangle) = 1$ if $n \neq k$ and $|r - s|$ otherwise. Let T be the theory of this structure.

It is not hard to show that $S_1(T)$ is homeomorphic to the set $\{0\} \times [0, 1] \cup \{\frac{1}{1+n} : n < \omega\} \times [0, 1]$ (with $p \mapsto \langle P(p), Q(p) \rangle$ as the homeomorphism), as a subset of $[0, 1]^2$ and that the d -metric on this type space is $d((x, y), (z, w)) = 1$ if $x \neq z$ and $|y - w|$ otherwise.

The intersection of any closed axis-aligned rectangle in \mathbb{R}^2 with $S_1(T) \subset [0, 1]^2$ is definable unless the resulting set is of the form $\{0\} \times I$ for some interval $I \subseteq [0, 1]$. Thus $S_1(T)$ is dictionary. (Another way to prove this is that this theory is ω -stable.)

Let $\{r_n\}_{n < \omega}$ be your favorite dense sequence in $[0, 1]$, and consider the subset of $S_1(T) \subset [0, 1]^2$ given by $\{0\} \times [0, 1] \cup \{(n, a_n) : n < \omega\}$. Call this set F . This is a closed




Figure 16: Infinite type space with precisely 3 definable sets.

subset, which is not dictionaric as a topometric space. To see that it is not dictionaric, consider the point $(0, \frac{1}{2})$, and let U be a small open neighborhood of it. Fix a closed set D such that $p \in D \subseteq U$. We will show that D is not definable. Find the largest $s \in [0, 1]$ such that $(0, s) \in D$ (which exists as D is closed). Find a sequence $\{t_i\}_{i < \omega}$ such that each $t_i > s$ and $t_i \rightarrow s$ as $i \rightarrow \infty$. By construction it is possible to find a sequence $p_{i < \omega}$ of types in $F \setminus \{0\} \times [0, 1]$ such that for each $i < \omega$, $p_i \notin F$ and the \mathbb{R}^2 distance between $(0, t_i)$ and p_i is less than 2^{-i} . This sequence ‘sneaks up’ on D by construction (i.e. topologically limits to D but maintains a uniform metric separation), so D is not definable. \square

Counterexample C.1.10 (Fig. 16). *An infinite type space with precisely 3 definable sets.*

Proof. Let \mathcal{L} be the empty signature. Let \mathfrak{M} be a metric space whose underlying set is $[0, 1]$ with a metric given by $d(0, r) = r$ and $d(r, s) = r + s$. Let T be the theory of this structure.

It is not hard to show that $S_1(T)$ is $[0, 1]$ and that the d -metric on $S_1(T)$ agrees with the metric on \mathfrak{M} . Clearly \emptyset and $[0, 1]$ are both definable sets. Open balls around 0 are of the form $[0, r)$ for $r \in [0, 1]$ (or of the form $[0, 1]$), and so 0 is d -atomic and $\{0\}$ is a definable set.




Figure 17: An infinite type space with precisely 4 definable sets, and a two sided definable endpoint that is not d -atomic.

Let F be a closed proper subset of $[0, 1]$ that contains some $r > 0$. This implies that the boundary of F is non-empty and contains a point other than 0. Let $s > 0$ be on the boundary of F , and consider $F^{<s/2}$. s has a neighborhood consisting solely of points which are distance more than $\frac{s}{2}$ from any other point in $S_1(T)$, therefore $F^{<s/2}$ cannot contain a neighborhood of s and F fails to be definable. \square

Counterexample C.1.11 (Fig. 17). *A two sided definable endpoint in a type space homeomorphic to $[0, 1]$ that is not d -atomic.*

Proof. Let \mathcal{L} be a signature with two 1-Lipschitz $[0, 1]$ -valued unary predicate symbols, P and Q , and let \mathfrak{M} be an \mathcal{L} -structure whose underlying set is $\{\langle 0 \rangle\} \cup (0, 1] \times \{0, 1\}$ with a metric given by $d(\langle r, i \rangle, \langle s, k \rangle) = 1$ if $r \neq s$ and $d(\langle r, 0 \rangle, \langle r, 1 \rangle) = r$ and in which we set $P(\langle r, i \rangle) = r$ and $Q(\langle r, i \rangle) = \frac{1+(-1)^i r}{2}$. Let T be the theory of this structure.

It's not hard to show that $S_1(T)$ is homeomorphic to $[0, 1] \times \{0, 1\}$ with the points $\langle 0, 0 \rangle$ and $\langle 0, 1 \rangle$ identified (and so is homeomorphic to $[0, 1]$) and that the d metric on

$S_1(T)$ agrees with the metric on \mathfrak{M} in the obvious way.

Using an argument similar to the one in Counterexample C.1.10, one can show that the only definable subsets of $S_1(T)$ are \emptyset , $S_1(T)$, $\{\langle 0, 0 \rangle\} \cup \{\langle 0, 0 \rangle\} \times [0, 1]$, and $\{\langle 0, 0 \rangle\} \cup \{\langle 0, 0 \rangle\} \times [0, 1]$, so in particular $\langle 0, 0 \rangle$ is not d -atomic but is a two sided definable endpoint.

□

Counterexample C.1.12. *A non-dictionaric type space $S_1(T)$ in which any two distinct types can be separated by disjoint definable sets.*

Proof. The polarized square (Definition 2.3.34) nearly has this property, but fails for pairs of types that are either both at the top or both at the bottom of the square (Corollary 2.3.44). This can be fixed by modifying the construction of the defining model by making the underlying set either $[0, 1] \times S^1$ or $[0, 1]^3$ and so making the model into either a disjoint collection of copies of the circle S^1 or the square $[0, 1]^2$ (with their standard metrics). If we then replace Q with Q_0 and Q_1 , predicates giving coordinates for either S^1 or $[0, 1]^2$, then a similar argument to that in Corollary 2.3.38 gives the required property.

□

C.2 Minimality and Categoricity

Counterexample C.2.1. *A strongly minimal set D with a definable set $E \subseteq D$ that is neither compact nor co-pre-compact.*

Proof. Let \mathfrak{M} be a structure whose universe is $\omega \times S^1$, where $S^1 \subseteq \mathbb{C}$ is the unit circle with the standard Euclidean metric. Let the distance between any points in distinct circles be 1. D is the entire structure, and let E be $\{(n, e^{2\pi ki/(n+1)}) : k \leq n\}$, i.e. on

the n th circle E consists of the $(n + 1)$ th roots of unity. As a subset of the type space $S_1(\mathfrak{M})$, E is $E(\mathfrak{M})$ together with the unique non-algebraic type. $E \subseteq S_1(\mathfrak{M})$ is clearly closed. To see that it is a definable set, pick $\varepsilon > 0$, and consider $E^{<\varepsilon}$. For any $n > \frac{4\pi}{\varepsilon}$, $E^{<\varepsilon}$ contains all of the circle $\{n\} \times S^1$. There are only finitely many $n \leq \frac{4\pi}{\varepsilon}$, and on each of these $E^{<\varepsilon}$ is an open set since the logic topology agrees with the metric topology on each individual circle in \mathfrak{M} . Therefore $E^{<\varepsilon}$ is an open set, so E is definable. \square

Another example is $(-\infty, 0] \cup \{\ln(1 + n) : n < \omega\}$, which is an \mathbb{R} -definable subset of \mathbb{R} (which is shown to be strongly minimal in Example 5.3.1).

Counterexample C.2.2. *An ω -stable theory with a strongly minimal type over \emptyset but no \emptyset -definable strongly minimal set.*

Proof. For each i let S_i be the sphere in an infinite dimensional Hilbert space with radius 2^{-i} . Let \mathfrak{M} be a metric space whose universe is $\bigsqcup_{i < \omega} S_i$, and let the distance between any points in distinct S_i be 1. Let $T = \text{Th}(\mathfrak{M})$.

The type space $S_1(\emptyset)$ is homeomorphic to $\omega + 1$ with the order topology. The unique non-isolated type is strongly minimal. If we let $\{a_i\}_{i < \omega}$ be a sequence of points such that $a_i \in S_i \subseteq \mathfrak{M}$, then the set $\{a_i\}_{i < \omega}$ is definable, by essentially the same argument as in Counterexample C.2.1. \square

Counterexample C.2.3. *A non-dictionaric superstable theory with a strongly minimal type over \emptyset but no approximately strongly minimal sets over \emptyset .*

Proof. Let $\mathcal{L} = \{P_0, P_1\}$ be a language with two unary 1-Lipschitz $[0, 1]$ -valued predicates, and let \mathfrak{M} be an \mathcal{L} -structure in whose universe is $\omega \times [0, 1]$ and whose metric is given by $d((n, x), (m, y)) = 1$ if $n \neq m$ and $d((n, x), (n, y)) = 2^{-n} + d(x, y)$ if $x \neq y$. Let $P_0((n, x)) = 2^{-n}$ and $P_1((n, x)) = x$. Finally let $T = \text{Th}(\mathfrak{M})$.

The type space $S_1(\emptyset)$ is homeomorphic $(\omega + 1) \times [0, 1]$.

Note that if a definable set has non-empty intersection with one of the sets of types of the form $\{n\} \times [0, 1]$ for $n < \omega$, then it must contain all of it, because this set is metrically isolated from the rest of the type space and is topologically connected, but uniformly metrically discrete. So to show that none of the types in $\{\omega\} \times [0, 1]$ are pointed to by an approximately strongly minimal set, all we need to do is show that if a definable set contains one such type then it must contain some type in $\{n\} \times [0, 1]$ for some $n < \omega$. This follows immediately because if $p \in \{\omega\} \times [0, 1]$, then it is the limit of types in $\{n\} \times [0, 1]$ for $n < \omega$ that are uniformly metrically separated. So if F is a closed set whose intersection with $\{n\} \times [0, 1]$ for $n < \omega$ is empty and whose intersection with $\{\omega\} \times [0, 1]$ is precisely p , then $p \notin \text{int}F^{<\varepsilon}$ for any $0 < \varepsilon < 1$, and F is not definable.

Thus if D is a definable set containing some $p \in \{\omega\} \times [0, 1]$, then D must contain all of $\{n\} \times [0, 1]$ for arbitrarily large $n < \omega$. And so since D is closed it must contain all of $\{\omega\} \times [0, 1]$. So if we let φ be a real formula such that $\varphi(p) = 0$, then for any $\varepsilon > 0$, $D \cap \llbracket \varphi \leq \varepsilon \rrbracket$ contains some $q \in \{\omega\} \times [0, 1]$ with $q \neq p$. \square

Counterexample C.2.4.

(i) If \mathfrak{M} is a model of any discrete theory with a Vaughtian pair, then the structure \mathfrak{M}^ω with the metric $d(\alpha, \beta) = 2^{-i}$ where i is the smallest index such that $\alpha(i) \neq \beta(i)$ and with the function $f(\alpha)(i) = \alpha(i + 1)$, has a definable Vaughtian pair but no open Vaughtian pairs. If $\text{Th}(\mathfrak{M})$ is ω -stable, but is not inseparably categorical then this is an example of an ω -stable theory with no open Vaughtian pairs that is not inseparably categorical.

(ii) If \mathfrak{N} is a discrete metric space with a single unary $[0, 1]$ -valued predicate whose

values are dense in $[0, 1]$, then $\text{Th}(\mathfrak{N})$ has no definable Vaughtian pairs (because every definable set is either finite or cofinite in every model) but does have an open Vaughtian pair. (Note that this is the same as the example in Proposition 2.3.32.)

(iii) The structure \mathfrak{N}^ω —where \mathfrak{N} is the structure in part (ii) with the truncation map $f(\alpha)(i) = \alpha(i + 1)$ and with a $[0, 1]$ -valued unary predicate that is the predicate from part (ii) evaluated on $\alpha(0)$ —is an example of a structure whose theory has no definable Vaughtian pairs and no open Vaughtian pairs but which does have an open-in-definable Vaughtian pair.

Proof. (i) is clear. (ii) is verified in Example 2.3.32. For (iii), $\text{Th}(\mathfrak{N}^\omega)$ has no open Vaughtian pairs, by the same argument as in (i). This is an imaginary sort of \mathfrak{N} , so in particular if D is a definable subset of \mathfrak{N}^ω and for some $\sigma \in \mathfrak{N}^{<\omega}$, we look at the ball $B_\sigma = \{\alpha \in \mathfrak{N}^{<\omega} : \sigma \prec \alpha\}$ and the function $f : B_\sigma \rightarrow \mathfrak{N}$ which maps $f(\alpha) = \alpha(|\sigma|)$ (i.e. the first element of α not in σ), then $D \cap B_\sigma$ is definable (since B_σ is logically clopen) and the image $f(D \cap B_\sigma)$ is a definable subset of \mathfrak{N} . This implies that any definable subset of \mathfrak{N}^ω is actually definable in the reduct where the $[0, 1]$ -valued predicate is removed. This reduct is \aleph_1 -categorical, so it has no definable Vaughtian pairs. Therefore $\text{Th}(\mathfrak{N}^\omega)$ has no definable Vaughtian pairs. However if we fix some element $b \in \mathfrak{N}$, and note that if $D = \{\alpha : (\forall i > 0)\alpha(i) = b\}$, then D is a definable set, since $d(x, D) = \frac{1}{2}d(f(x), b^\omega)$. We then get an open-in-definable Vaughtian pair by considering the same open set as in part (ii), even though this theory does not have any definable Vaughtian pairs or any open Vaughtian pairs. \square

Note that there is a strictly stable discrete theory with no Vaughtian pairs but which does have an imaginary Vaughtian pair, although the same cannot happen in a discrete

superstable theory [Bou89].

Counterexample C.2.5. *An inseparably categorical ultrametric theory with no strongly minimal set.*

Proof. Let $\mathfrak{M} = c_0(\omega, \mathbb{Z}_p)$ be the unit ball of the p -adic Banach space $c_0(\kappa, \mathbb{Q}_p)$, i.e. \mathfrak{M} consists of elements a of \mathbb{Z}_p^ω satisfying $a(i) \rightarrow 0$ as $i \rightarrow \omega$. The metric on \mathfrak{M} is given by $d(a, b) = \sup d^{\mathbb{Z}_p}(a(i), b(i))$. The language consists only of $+$ (note that since \mathbb{Z} is dense in \mathbb{Z}_p , we don't actually need to have explicit scalar multiplication functions).

(*) Clearly we have that the binary relation $[d(x, y) \leq p^{-1}]$ is an equivalence relation. The imaginary obtained by quotienting by this equivalence relation is clearly the infinite dimensional vector field over the finite field \mathbb{F}_p . Furthermore models of $\text{Th}(\mathfrak{M})$ are prime over this imaginary. To see this assume that $\mathfrak{A} \prec \mathfrak{B}$ is a proper elementary pair of models of this theory. Let $b \in \mathfrak{B} \setminus \mathfrak{A}$. Let a be an element of \mathfrak{A} such that $d(b, a) = d(b, \mathfrak{A})$ (this always exists because the set of possible distances is reverse well-ordered). Then we have that $d(b - a, 0) = d(b, a)$ and $b - a \notin \mathfrak{A}$. If this is 1, then we are done, since b is necessarily in its own equivalence class in the imaginary that is not contained in \mathfrak{A} , otherwise $d(b - a, 0) = p^{-n}$ for some $n > 0$. The theory knows that if $d(c, 0) = p^{-n}$, then there is a unique element e satisfying $p^n e = c$. Let $p^n c = b - a$. We have that c must be an element of \mathfrak{B} satisfying $d(c, \mathfrak{A}) = 1$. Therefore the imaginary must be bigger in \mathfrak{B} than it is in \mathfrak{A} .

So since models of $\text{Th}(\mathfrak{M})$ are prime over a strongly minimal imaginary we have that $\text{Th}(\mathfrak{M})$ is inseparably categorical.

To see that $\text{Th}(\mathfrak{M})$ has no strongly minimal sets, note that it is enough to show that $S_1(\mathfrak{M})$ has no pre-minimal types, since \mathfrak{M} is approximately \aleph_0 -saturated. If $p \in S_1(\mathfrak{M})$

is some non-algebraic type, then the argument in paragraph (*) gives an \mathfrak{M} -definable bijection between the smallest ball centered on some element of \mathfrak{M} that contains p and the entire structure. Since the unique non-algebraic type q satisfying $d(q, \mathfrak{M}) = 1$ is not minimal, this implies that p is not minimal. Since we could do this for any non-algebraic p , this implies that no type is minimal, and so no type is strongly minimal. \square

Counterexample C.2.6. *A totally categorical theory T with a strongly minimal set such that for any strongly minimal set $D(x, \bar{a})$ in the home sort of the unique separable model \mathfrak{M} , for any $n \leq \omega$, there is a $\bar{b} \equiv \bar{a}$ such that $\dim(D(\mathfrak{M}, \bar{b})) = n$.*

Proof. The construction of this counterexample is very similar to the construction of T_ω in Theorem C.2.9, so we will only sketch the differences. Make the following two changes to the construction:

- Define the score of an element $\alpha \in \mathfrak{W}$ by

$$s(\alpha) = \min \left\{ g(A) : A \in G, \sum \alpha(A) = 0 \right\}.$$

- Remove the predicates C_v from the language.

A similar analysis gives that the theory of this structure has a \emptyset -definable strongly minimal imaginary I whose dimension in the prime model is infinite. Furthermore types over the prime model \mathfrak{M} in the home sort are strongly minimal if and only if they correspond to elements α where $\alpha(0), \alpha(2), \alpha(4), \dots$ enumerates a linearly independent set and only one $\alpha(2i + 1)$ is not realized in $I(\mathfrak{M})$.

If for example $\alpha(0)$ is not realized in $I(\mathfrak{M})$, then we can choose $\beta(2)\beta(4) \dots \equiv \alpha(2)\alpha(4) \dots$ such that the subspace of $I(\mathfrak{M})$ spanned by $\beta(2)\beta(4) \dots$ has arbitrary

codimension $\leq \omega$, giving the required property for the theory of this structure. \square

Counterexample C.2.7. *An inseparably categorical theory T with a strongly minimal imaginary that is an affine space but for which for any set of parameters A in the home sort, $\text{acl}(A) = \overline{A}$. In particular, this implies that T has no orthonormalizable set in its home sort.*

Proof. The construction of this counterexample is very similar to the construction in Counterexample C.2.10. The only difference is that we lift the requirement that n and m have the same parity in the relation $P_{n,m}$. This ties the two dimensions together. Modding out by the clopen equivalence relation $(d(x, y) < 1) \equiv (d(x, y) \leq \frac{3}{4})$ gives the required strongly minimal imaginary. Just as in Counterexample C.2.10 the theory is ω -categorical. Let \mathfrak{M} be its unique separable model. It is enough to establish the claim regarding acl in \mathfrak{M} .

Let $A \subseteq \mathfrak{M}$ be some set of parameters, and suppose that $b \notin \overline{A}$. We want to show that $b \notin \text{acl}(A)$. Since $b \notin \overline{A}$, for some $n < \omega$, for every $a \in A$, $a(n) \neq b(n)$. This implies that for any $v \in V$ we can find an automorphism σ of \mathfrak{M} fixing A pointwise such that $\sigma b(n+1) = v$, so $\text{tp}(b/A)$ is not algebraic and $b \notin \text{acl}(A)$, as required.

For the ‘in particular’ statement, if D is any non-algebraic \mathfrak{M} -definable set, the same analysis as in Counterexample C.2.10 shows that D must contain more than one non-algebraic type. Let p and q be non-algebraic types in D . If a is a realization of p , then we have that $\text{acl}(\mathfrak{M}a) = \overline{\mathfrak{M}a} = \mathfrak{M}a$, which cannot contain any realizations of q . Therefore D is not orthonormalizable. \square

Counterexample C.2.8. *A type that is d -atomic-in- $S_1(\mathfrak{C}) \setminus \mathfrak{C}^{<\varepsilon}$, but which is not ε -peripheral.*

Proof. Consider the structure \mathfrak{M} in the empty signature with universe ω^2 and the metric $d((a, b), (c, e)) = 1$ if $b \neq e$ and $d((a, b), (c, b)) = 1 - 2^{-b-1}$ if $a \neq c$. The theory of \mathfrak{M} is inter-definable with a discrete theory with a sequence of unary predicates $\{P_i\}_{i < \omega}$ such that the P_i are pairwise disjoint and infinite.

The type over \mathfrak{M} axiomatized by $\forall y d(x, y) = 0 \vee d(x, y) = 1$ is minimal (i.e. has a unique non-algebraic extension over every set of parameters). Let p be its global non-algebraic extension. $S_1(\mathfrak{C}) \setminus \mathfrak{C}^{<1} = \{p\}$, so p is d -atomic-in- $S_1(\mathfrak{C}) \setminus \mathfrak{C}^{<1}$. Now we just need to show that p is not 1-peripheral.

Let D be any \mathfrak{C} -definable set containing p . For each $n < \omega$, the type over \mathfrak{M} axiomatized by $(\forall y d(x, y) = 0 \vee d(x, y) = 1 \vee d(x, y) = 2^{-n-1}) \wedge \exists y d(x, y) = 2^{-n-1}$ is minimal. Let q_n be its global non-algebraic extension. We have that $q_n \rightarrow p$ topologically but for each $n < \omega$, for any global type $r \neq q_n$, $d(r, q_n) \geq 2^{-n-1}$, therefore D must contain q_n for any sufficiently large n . We have that $d_{\text{inf}}(q_n, \mathfrak{C}) \rightarrow 1 = d_{\text{inf}}(p, \mathfrak{C})$ as $n \rightarrow \infty$, but $d(p, q_n) = 1$, so D fails to be a 1-peripheral set pointing to p .

Since we can do this for any definable set containing p , it is not a 1-peripheral type. □

Of course this theory is ω -stable and so by Proposition 4.5.9 it has ε -peripheral types, specifically the strongly minimal types q_n for any $n < \omega$.

C.2.1 A Theory with Strongly Minimal Sets, but Only over Models of Dimension $\geq n$

Theorem C.2.9. *For any $n \leq \omega$ there is an inseparably categorical theory T_n with a \emptyset -definable strongly minimal imaginary I such that*

- T_n has models with $\dim(I(\mathfrak{M})) = k$ for each $k \leq \omega$ but
- T_n has a strongly minimal set over \mathfrak{M} in the home sort if and only if $\dim(I(\mathfrak{M})) \geq n$.

Proof. Fix $n \leq \omega$. Let V be the countable vector space over \mathbb{F}_2 . Let $\mathfrak{W} = V^\omega$ have the standard string ultrametric (i.e. $d(\alpha, \beta) = 2^{-\ell}$ where ℓ is the length of the longest common initial segment of α and β). Let $f : V \rightarrow \omega$ be a fixed bijection. Let $G \subset \mathcal{P}(\{0, 2, 4, \dots, 2n\})$ be the set of all non-empty subsets A of $\{0, 2, 4, \dots, 2n\}$, if $n < \omega$, and $G \subset \mathcal{P}_{\text{fin}}(\{0, 2, 4, \dots\})$ be the set of all non-empty finite subsets of $\{0, 2, 4, \dots\}$, if $n = \omega$. Let $g : G \rightarrow \omega$ be a fixed injection. Let $\langle -, - \rangle : \omega \times \omega \rightarrow \omega$ be a fixed pairing function.

For any $\alpha \in \mathfrak{W}$ we assign a score

$$s(\alpha) = \min \left\{ \left\langle f \left(\sum \alpha(A) \right), g(A) \right\rangle : A \in G \right\}.$$

Now let $D \subset \mathfrak{W}$ be the set of all α such that $\alpha(2i) = 0$ for all $i > n$ and $\alpha(2i + 1) = 0$ for all $i \leq s(\alpha)$.

Claim: D is a \mathfrak{W} -definable set.

Proof of claim: For each $k < \omega$ let $D_k \subseteq \mathfrak{W}^\omega$ be the set of all α satisfying:

- For any i with $i > n$ and $2i < k$, the formula $\alpha(2i) = 0$.
- For any $i < \omega$ with $2i + 1 < k$, the formula

$$\alpha(2i + 1) \neq 0 \rightarrow \bigvee_{\langle v, A \rangle \leq i} \sum (\alpha(A)) = v.$$

In the language of \mathfrak{W} we can say $\alpha(j) = v$ for any $j < \omega$ and $v \in V$ with a $\{0, 1\}$ -valued formula. Likewise note that there are only finitely many pairs v, A with $\langle v, A \rangle \leq i$, so the disjunction in the second family of formulas is first-order. So we have that each D_k is definable by a $\{0, 1\}$ -valued formula.

It's clear that $D_k \supseteq D_{k+1}$ and $D(\mathfrak{W}) = \bigcap_{k < \omega} D_k(\mathfrak{W})$, so we may regard D as the \mathfrak{W} -zeroset $\bigcap_{k < \omega} D_k$. Fix $\ell < \omega$, and find $m < \omega$ such that for any $v \in V$ with $f(v) \leq \ell$ and any $A \in G$ with $\max A \leq \ell$ and $\langle v, A \rangle \leq m$. Fix $\alpha \in D_{2m+1}$. Now find β such that $\beta \upharpoonright \ell = \alpha \upharpoonright \ell$ and $\beta(i) = 0$ for all $i \geq \ell$, so in particular $d(\alpha, \beta) \leq 2^{-\ell}$. Now we have that $\beta \in D(\mathfrak{W})$, so we know that $d(\alpha, D(\mathfrak{W})) \leq 2^{-\ell}$ for all $\alpha \in D_{2m+1}$.

Since we can do this for arbitrarily large $\ell < \omega$ we can define a distance predicate for D and D is a definable set. \square_{claim}

Now consider the following set of $\{0, 1\}$ -valued definable predicates:

- For each $v \in V$, $C_v(\alpha) = 0$ if and only if $\alpha(0) = v$.
- $P(\alpha, \beta, \gamma) = 0$ if and only if $\alpha(0) = \beta(0) + \gamma(0)$.
- For any even $k < \omega$, $Q_k(\alpha, \beta) = 0$ if and only if $\alpha(0) = \beta(k)$.
- For any odd $k < \omega$, $P_k(\alpha, \beta, \gamma) = 0$ if and only if $d(\beta, \gamma) \leq 2^{-k}$ and $\alpha(0) = \beta(k) + \gamma(k)$.

Now let $\mathfrak{A} = \langle D(\mathfrak{W}), \{C_v\}_{v \in V}, P, \{Q_{2k}\}_{k < \omega}, \{R_{2k+1}\}_{k < \omega} \rangle$ and $T_n = \text{Th}(\mathfrak{A})$. As a reduct of $\text{Th}(\mathfrak{W})$, it's clear that T_n is ω -stable. Furthermore since $\text{Th}(\mathfrak{W})$ is inseparably categorical we can easily construct an \aleph_1 -saturated model of T_n . Let $V' \supset V$ be the \mathbb{F}_2 vector space with dimension \aleph_1 , let $\mathfrak{W}' = (V')^\omega$, and let $\mathfrak{B} \succ \mathfrak{A}$ be the corresponding

model of T_n . It's clear that \mathfrak{W}' is \aleph_1 -saturated over \mathfrak{W} , so \mathfrak{B} is \aleph_1 -saturated as a model of T_n .

We can extend the definition of $s(\alpha)$ to cover $\alpha \in \mathfrak{W}'$ if we allow for $s(\alpha) = \infty = \min \emptyset$. Then we can check, by the definitions of D_k , that $D(\mathfrak{W}')$ is precisely the set of $\alpha \in \mathfrak{W}'$ such that $\alpha(2i) = 0$ for all $i > n$ and $\alpha(2i + 1) = 0$ for all $i \leq s(\alpha)$ (where $i \leq \infty$ for all $i < \omega$).

Claim: T_n is inseparably categorical with a \emptyset -definable strongly minimal imaginary.

Proof of claim: If we let E denote the $\{0, 1\}$ -valued equivalence relation given by $E(x, y) = 2(d(x, y) \div \frac{1}{2})$, then it's clear than the quotient H/E of the home sort H by E is a \emptyset -definable strongly minimal imaginary which is equivalent to a \mathbb{F}_2 vector space with constants for elements of the prime model. Now we just need to show that T_n has no Vaughtian pairs over H/E . Suppose that $\mathfrak{M} \prec \mathfrak{N}$ are models of T_n such that $H(\mathfrak{M})/E = H(\mathfrak{N})/E$. Let α be in $\mathfrak{N} \setminus \mathfrak{M}$, and assume that $d(\alpha, \mathfrak{M}) = 2^{-k}$. Find $\beta \in \mathfrak{M}$ such that $d(\alpha, \beta) = 2^{-k}$ (this exists because the distance set is reverse well-ordered).

We may assume that \mathfrak{M} and \mathfrak{N} are separated and furthermore that $\mathfrak{N} \prec \mathfrak{B}$, so we can regard the elements of \mathfrak{M} and \mathfrak{N} as elements of $(V')^\omega$. In particular we have that for all $\alpha \in \mathfrak{N}$, $\alpha(2i) = 0$ for all $i > n$ and $\alpha(2i + 1) = 0$ for all $i \leq s(\alpha)$.

Now we can see that k cannot be even. If $k = 2m$ for some m , then we can find $\gamma \in \mathfrak{M}$ such that $d(\alpha, \gamma) \leq 2^{-k-1}$ by finding the element of $H(\mathfrak{M})/E$ corresponding to $\alpha(2k)$ and replacing β with γ satisfying $\gamma(i) = \beta(i)$ for all $i \neq 2k$ and $\gamma(2k) = \alpha(2k)$.

So assume that $k = 2m + 1$ for some m . There is some element c of $H(\mathfrak{N})/E$ such that for any σ with $\sigma(0) = c$, we have $\mathfrak{N} \models P_{2m+1}(\sigma, \beta, \alpha)$. Therefore we can find such a σ in \mathfrak{M} . But now there must be some α' with $\mathfrak{M} \models P_{2m+1}(\sigma, \beta, \alpha')$, which implies that $d(\alpha, \alpha') \leq 2^{-k-1}$, which is also a contradiction.

Therefore there are no Vaughtian pairs over H/E and T_n is inseparably categorical.

□_{claim}

Now it's clear that the models of T_n are uniquely determined by $\dim(H/E)$ and that every dimension ≥ 0 is possible. So for each $k \leq \omega$ let $V_k \subset V'$ be a vector space such that $\dim(V_k/V) = k$, and let $\mathfrak{M}_k \prec \mathfrak{B}$ be the corresponding model of T_n . We need to characterize the types in $S_1(\mathfrak{M}_k)$. Note that every type in $S_1(\mathfrak{M}_k)$ is realized in \mathfrak{B} , since it is \aleph_1 -saturated.

Fix $k \leq \omega$. For each $\tau \in V_k^{<\omega}$, let $\zeta_\tau \in \mathfrak{M}_k$ be a fixed element satisfying $\zeta_\tau \supset \tau$. For any $\alpha \in \mathfrak{W}'$, assign an 'index' $h(\alpha) = \langle \eta(\alpha), X(\alpha) \rangle$ by these specifications:

- Let $\eta(\alpha) \in V_k^{\leq \omega}$ be the longest initial segment of α such that $\eta(\alpha)(i) \in V_k$ for every $i < |\eta(\alpha)|$.
- If $|\eta(\alpha)|$ is infinite then $X(\alpha) = \emptyset$.
- If $|\eta(\alpha)|$ is finite and even, then $X(\alpha) \subseteq \mathcal{P}_{\text{fin}}(\omega) \times V_k$ is the set of pairs $\langle A, v \rangle$ such that $\sum \alpha(2A) = v$.
- If $|\eta(\alpha)|$ is finite and odd, then $X(\alpha) \subseteq \mathcal{P}_{\text{fin}}(\omega \cup \{-1\}) \times V_k$ is the set of pairs $\langle A, v \rangle$ such that $\sum \alpha(2A) = v$, where we set $\alpha(-2) = \zeta_{\alpha \upharpoonright |\eta(\alpha)|}(|\eta(\alpha)|) + \alpha(|\eta(\alpha)|)$. (Recall that $\alpha \upharpoonright m$ is the sequence $\alpha(0), \alpha(1), \dots, \alpha(m-1)$, with length m .)

With the following two claims we will show that $h(\alpha)$ exactly captures $\text{tp}(\alpha/\mathfrak{M}_k)$.

Claim 1: For any $\alpha, \beta \in \mathfrak{B}$, if $h(\alpha) \neq h(\beta)$ then $\alpha \not\equiv_{\mathfrak{M}_k} \beta$.

Claim 2: For any $\alpha, \beta \in \mathfrak{B}$, if $h(\alpha) = h(\beta)$ then for any $\varepsilon > 0$ there is an automorphism σ of \mathfrak{B} , fixing \mathfrak{M}_k , such that $d(\sigma(\alpha), \beta) \leq \varepsilon$.

Proof of claim 1: If either of $\eta(\alpha)$ or $\eta(\beta)$ is infinitely long, then the corresponding element is an element of \mathfrak{M}_k , so then α and β clearly have different types over \mathfrak{M}_k . So assume that both $\eta(\alpha)$ and $\eta(\beta)$ are finite.

If $\eta(\alpha) \neq \eta(\beta)$, then there are elements of \mathfrak{M}_k with different distances to α and β , implying that they have different types over \mathfrak{M}_k . So assume that $\eta(\alpha) = \eta(\beta)$.

If $\eta(\alpha) = \eta(\beta)$, then $X(\alpha) \neq X(\beta)$, and so this clearly gives an \mathfrak{M}_k -formula satisfied by α and not satisfied by β , so we have that $\alpha \not\equiv_{\mathfrak{M}_k} \beta$.

So we have that $h(\alpha) \neq h(\beta) \Rightarrow \alpha \not\equiv_{\mathfrak{M}_k} \beta$, as required. $\square_{\text{claim 1}}$

Proof of claim 2: Assume that $h(\alpha) = h(\beta)$. If $\eta(\alpha) = \eta(\beta)$ is infinitely long, then $\alpha = \beta \in \mathfrak{M}_k$ and there is nothing to prove, so assume that $\eta(\alpha) = \eta(\beta)$ is finitely long.

First we will prove that there is an automorphism σ_0 of \mathfrak{B} fixing \mathfrak{M}_k such that $\sigma_0(\alpha)(2i) = \beta(2i)$ for every $i < \omega$ and $\sigma_0(\alpha)(|\eta(\alpha)|) = \beta(|\eta(\alpha)|)$ if $|\eta(\alpha)|$ is odd.

If $|\eta(\alpha)| = |\eta(\beta)|$ is finite and even, then $X(\alpha) = X(\beta)$ is precisely the statement that

$$\alpha(0)\alpha(2)\alpha(4) \cdots \equiv_{V_k} \beta(0)\beta(2)\beta(4) \cdots$$

in the structure V' . So we easily get an automorphism of V' taking $\alpha(0)\alpha(2)\alpha(4) \cdots$ to $\beta(0)\beta(2)\beta(4) \cdots$ fixing V_k . This extends to an automorphism of all of \mathfrak{W}' fixing V_k^ω which then induces an automorphism of \mathfrak{B} fixing \mathfrak{M}_k with the required property.

If $|\eta(\alpha)| = |\eta(\beta)|$ is finite and odd, then $X(\alpha) = X(\beta)$ is precisely the statement that

$$\begin{aligned} & (\zeta_{\alpha|\eta(\alpha)}(|\eta(\alpha)|) + \alpha(|\eta(\alpha)|))\alpha(0)\alpha(2)\alpha(4) \cdots \\ & \equiv_{V_k} (\zeta_{\alpha|\eta(\alpha)}(|\eta(\alpha)|) + \beta(|\eta(\alpha)|))\beta(0)\beta(2)\beta(4) \cdots \end{aligned}$$

in the structure V' . So we easily get an automorphism of V' taking

$$(\zeta_{\alpha \upharpoonright |\eta(\alpha)|}(|\eta(\alpha)|) + \alpha(|\eta(\alpha)|)), \alpha(0), \alpha(2), \alpha(4), \dots$$

to

$$(\zeta_{\alpha \upharpoonright |\eta(\alpha)|}(|\eta(\alpha)|) + \beta(|\eta(\alpha)|)), \beta(0), \beta(2), \beta(4), \dots$$

and fixing V_k . This extends to an automorphism of all of \mathfrak{W}' fixing V_k^ω which then induces an automorphism of \mathfrak{B} fixing \mathfrak{M}_k with the required property.

Now we need to argue that $h(\sigma_0(\alpha)) = h(\beta)$. If $|\eta(\alpha)| = |\eta(\beta)|$ is even, then there is nothing to prove. If $|\eta(\alpha)| = |\eta(\beta)|$ is odd, then the only thing to worry about is that we might have moved $\zeta_{\alpha \upharpoonright |\eta(\alpha)|}(|\eta(\alpha)|)$, but this is determined by $\alpha \upharpoonright |\eta(\alpha)| = \beta \upharpoonright |\eta(\beta)|$ which is unchanged and equal to $\sigma_0(\alpha) \upharpoonright |\eta(\sigma_0(\alpha))|$. So we have $h(\sigma_0(\alpha)) = h(\beta)$.

So now let $\gamma \in (V')^{\leq \omega}$ be the longest common initial segment of $\sigma_0(\alpha)$ and β , so that in particular $d(\sigma_0(\alpha), \beta) = 2^{-|\gamma|}$. If γ is infinitely long then we're done, so assume that γ is finitely long. It must be the case that $|\gamma|$ is odd, since we ensured that $\sigma_0(\alpha)(2i) = \beta(2i)$ for every i . If $|\eta(\beta)|$ is odd, then it must be the case that $|\gamma| > |\eta(\beta)|$, since we ensured that $\sigma_0(\alpha)(|\eta(\beta)|) = \beta(|\eta(\beta)|)$. In any case we always have $|\gamma| > |\eta(\beta)| \geq 0$, so in particular we always have that the last element of γ , $\gamma(|\gamma| - 1)$, is not in V_k .

Now define a map $\sigma_1 : \mathfrak{B} \rightarrow \mathfrak{B}$ by these specifications:

- $\sigma_1(\chi) = \chi$ if $d(\chi, \beta) \geq 2^{-|\gamma|+1}$.
- If $d(\chi, \beta) \leq 2^{-|\gamma|}$, so in particular $\chi(|\gamma| - 1) = \beta(|\gamma| - 1)$, then $\sigma_1(\chi)(j) = \chi(j)$, if $j \neq |\gamma|$, and $\sigma_1(\chi)(|\gamma|) = \chi(|\gamma|) + \beta(|\gamma|) + \sigma_0(\alpha(|\gamma|))$.

By checking the definition of D we can see that σ_1 is a bijection on \mathfrak{B} . It's also clearly an isometric map. By checking the predicates in the language of \mathfrak{B} we can see that σ_1 is an automorphism of \mathfrak{B} . Furthermore, clearly $\sigma_1(\alpha)(2i + 1) = \beta(2i + 1)$, so we have that $d(\sigma_1(\sigma_0(\alpha)), \beta) < d(\alpha, \beta)$, as required.

Now the only thing left to verify is that σ_1 fixes \mathfrak{M}_k . If χ is moved by σ_1 , then $d(\gamma, \alpha) \leq 2^{-|\gamma|}$, so in particular $\chi(|\gamma| - 1) = \beta(|\gamma| - 1) \notin V_k$, so $\chi \notin \mathfrak{M}_k$. Therefore everything moved by σ_1 is not in \mathfrak{M}_k and we have that σ_1 fixes \mathfrak{M}_k .

So by iterating the construction of σ_1 we can get the required automorphisms bringing α arbitrarily close to β . $\square_{\text{claim 2}}$

So we see that $\alpha \equiv_{\mathfrak{M}_k} \beta$ if and only if $h(\alpha) = h(\beta)$.

Now we need to determine the values of $h(\alpha)$ that correspond to a strongly minimal type.

Claim: If $\alpha(2i + 1) \neq 0$ for some $i < \omega$, then $\text{tp}(\alpha/\mathfrak{M}_k)$ is not strongly minimal.

Proof of claim: If $\alpha \in \mathfrak{M}_k$, then $\text{tp}(\alpha/\mathfrak{M}_k)$ is clearly not strongly minimal, so assume that $\alpha \notin \mathfrak{M}_k$. Since $\alpha \notin \mathfrak{M}_k$, we have that $\eta(\alpha)$ has finite length. Since $\alpha(2i + 1) \neq 0$ for some $i < \omega$, we must have $s(\alpha) < \infty$ (this is determined by $\text{tp}(\alpha/\mathfrak{A})$ so *a fortiori* it is determined by $\text{tp}(\alpha/\mathfrak{M}_k)$). This implies that if β has $\alpha(i) = \beta(i)$ for all $i \leq |\eta(\alpha)| \uparrow (2s(\alpha) + 1)$ and $\alpha(2i) = \beta(2i)$ for all $i < \omega$, then $\alpha \equiv_{\mathfrak{M}_k} \beta$, implying that α/\mathfrak{M}_k has many non-algebraic global extensions and is not strongly minimal. \square_{claim}

As a corollary of this we get that if $\alpha(0), \alpha(2), \dots, \alpha(2n)$ (where ' $\alpha(0), \alpha(2), \dots, \alpha(2\omega)$ ' is understood to mean $\alpha(0), \alpha(2), \dots$) are not linearly independent over V , then $\text{tp}(\alpha/\mathfrak{M}_k)$ is not strongly minimal. Conversely we have that if $\alpha(0), \alpha(2), \dots, \alpha(2n)$ are linearly independent over V , then $\alpha(2i + 1) = 0$ necessarily for every $i < \omega$, by the definition of D .

Claim: If $\alpha(0), \alpha(2), \dots, \alpha(2n)$ are linearly independent over V and

$$\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) > 1,$$

then $\text{tp}(\alpha/\mathfrak{M}_k)$ is not strongly minimal.

Proof of claim: If $\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) > 1$ then $\text{tp}(\alpha/\mathfrak{M}_k)$ has more than one global non-algebraic extension and so is not strongly minimal. \square_{claim}

Clearly if $\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) = 0$, then we have that $\text{tp}(\alpha/\mathfrak{M}_k)$ is atomic and therefore not strongly minimal. So the only way for $\text{tp}(\alpha/\mathfrak{M}_k)$ to be strongly minimal is if $\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) = 1$.

As it turns out this is a precise characterization.

Claim: $\text{tp}(\alpha/\mathfrak{M}_k)$ is strongly minimal if and only if

$$\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) = 1.$$

Proof of claim: We have already shown the \Rightarrow direction, so we just need to show the converse, but this is easy since $\text{tp}(\alpha/\mathfrak{M}_k)$ is clearly non-algebraic but also has a unique non-algebraic extension over any parameter set. \square_{claim}

Now finally we see that

$$\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/\mathfrak{M}_k) = 1$$

with

$$\dim(\{\alpha(0), \alpha(2), \dots, \alpha(2n)\}/V) = n + 1$$

is possible if and only if $k \geq n$, and so T_n has the required property. \square

C.2.2 A Counterexample to a Direct Translation of the Baldwin-Lachlan Characterization

One might hope that somehow the condition that a theory be ω -stable and have no Vaughtian pairs might be strong enough to ensure that a theory is inseparably categorical, but it is not so. The full Baldwin-Lachlan characterization fails in continuous logic, even after strengthening the no Vaughtian pairs condition to no locatable Vaughtian pairs.

Counterexample C.2.10. *A countable ω -stable theory with no locatable Vaughtian pairs that is not inseparably categorical.*

Proof. Let V be a countable vector space over a finite field \mathbb{F}_p . Let \mathfrak{M} be the structure whose universe is V^ω with the standard string ultrametric. For each $n, m < \omega$ with $n \equiv m \pmod{2}$, let $P_{n,m}$ be a $\{0, 1\}$ -valued quaternary relation such that $P_{n,m}(a, b, c, e) = 0$ if and only if $d(a, b) \leq 2^{-n+1}$, $d(c, e) \leq 2^{-m+1}$, and $a(n) - b(n) = c(m) - e(m)$. For any fixed $n, m < \omega$ we have that if $d(a_0, a_1), d(b_0, b_1) < 2^{-n-1}$ and $d(c_0, c_1), d(e_0, e_1) < 2^{-m-1}$, then $P_{n,m}(a_0, b_0, c_0, e_0) = P_{n,m}(a_1, b_1, c_1, e_1)$, so each $P_{n,m}$ is uniformly continuous.

We will show that T is \aleph_0 -categorical. To see that $T = \text{Th}(\mathfrak{M})$ is ω -stable and \aleph_0 -categorical, note that it is a reduct of an imaginary of $(V, +)$, which is ω -stable and \aleph_0 -categorical. It is also a reduct of Counterexample C.2.5.

To see that T is not inseparably categorical, notice that if W is the unique elementary extension of V of cardinality \aleph_1 , then $(V \times W)^\omega$, $(W \times V)^\omega$, and $(W \times W)^\omega$ are the universes of three non-isomorphic models of T .

Finally to see that T has no locatable Vaughtian pairs we need to analyze the structure of its type spaces more carefully. It is sufficient to prove this statement considering locatable subsets of $S_1(\mathfrak{M})$, where \mathfrak{M} is the unique separable model, since any countable set of parameters can be found inside \mathfrak{M} .

Claim: Every non-algebraic type $p \in S_1(\mathfrak{M})$ is uniquely determined by the $k < \omega$ for which $d(p, \mathfrak{M}) = 2^{-k}$ and $\alpha \upharpoonright k$ for any $\alpha \in \mathfrak{M}$ with $d(\alpha, p) = 2^{-k}$.

Proof of claim. To see that this is true, let a, b be two elements of some $\mathfrak{N} \succ \mathfrak{M}$ with $d(a, \mathfrak{M}) = d(b, \mathfrak{M}) = 2^{-k}$, and suppose that there are $\alpha, \beta \in \mathfrak{M}$ with $d(a, \alpha) = 2^{-k}$ and $d(b, \beta) = 2^{-k}$ and $\alpha \upharpoonright k = \beta \upharpoonright k$. The statement $\alpha \upharpoonright k = \beta \upharpoonright k$ is the same as saying that $d(\alpha, \beta) \leq 2^{-k}$, so since T is ultrametric we have that $d(b, \alpha) \leq 2^{-k} \uparrow 2^{-k} = 2^{-k}$, so $d(b, \alpha) = 2^{-k}$. Therefore also $d(a, b) \leq 2^{-k}$. \square_{claim}

\mathfrak{N} can be written in the form $(U \times W)^\omega$ for some \mathbb{F}_p -vector spaces $U, W \succeq V$.

Assume that $d(a, b) = 2^{-k}$. Either $a(k), b(k) \in U \setminus V$ or $a(k), b(k) \in W \setminus V$ (depending on whether k is even or odd), and $a(k) \neq b(k)$, since $d(a, b) = 2^{-k}$, so we can find an explicit automorphism of U or W fixing V and taking $a(k)$ to $b(k)$. This extends to an automorphism f of all of \mathfrak{N} fixing \mathfrak{M} . So by replacing a with $f(a)$, we may assume that $d(a, b) < 2^{-k}$.

Assume that $d(a, b) < 2^{-k}$. This implies that $a(k) = b(k) \notin V$. Let $\ell > k$ be the first such that $a(\ell) \neq b(\ell)$. Let $g : \mathfrak{N} \rightarrow \mathfrak{N}$ defined by the following: $g(c) = c$ if $c \upharpoonright \ell \neq a \upharpoonright \ell$, otherwise $g(c)(m) = c(m)$ for $m \neq \ell$ and $g(c)(\ell) = c(\ell) + b(\ell) - a(\ell)$. One can check that this is an automorphism of \mathfrak{N} with the property that $d(g(a), b) < d(a, b)$. Note that g fixes \mathfrak{M} .

By iterating, for any $k < \omega$ we can find an automorphism h of \mathfrak{N} , fixing \mathfrak{M} , such that $d(h(a), b) \leq 2^{-k}$. Therefore $a \equiv_{\mathfrak{M}} b$, as required.

Let each non-algebraic type in $S_1(\mathfrak{M})$ be denoted by an element of $V^{<\omega}$.

Note that each non-algebraic type is metrically isolated. To see this consider $\sigma_0, \sigma_1 \in V^{<\omega}$ with $\sigma_0 \neq \sigma_1$. Let τ be the longest common initial segment of σ_0 and σ_1 . We have that $d(\sigma_0, \sigma_1) = 2^{-|\tau|}$. In particular this implies that other types have distance at least $2^{-|\sigma|}$ to the type associated to some $\sigma \in V^{<\omega}$.

Now also note that for any $\sigma \in V^{<\omega}$, the sequence of types $\{\sigma \frown v\}_{v \in V}$ limits to σ , because the limiting type must be in the $(\leq 2^{-|\sigma|})$ -ball whose center starts with σ , but it must have distance greater than $2^{-|\sigma|-1}$ from \mathfrak{M} , so σ is the unique type that it can be.

Let $L \subseteq S_1(\mathfrak{M})$ be a locatable set containing a non-algebraic type σ . The claim is that for some $v \in V$, L must contain $\sigma \frown v$ as well. To see this, note that $L^{<2^{-|\sigma|-2}}$ must be a neighborhood of σ , so it must contain infinitely many types of the form $\sigma \frown v$. Let $\sigma \frown v$ be in $L^{<2^{-|\sigma|-2}}$, and let $\tau \in L$ be such that $d(\sigma \frown v, \tau) < 2^{-|\sigma|-2}$. But since σ is $(\geq 2^{-|\sigma|-1})$ -metrically isolated, this implies that $\sigma \frown v = \tau$.

Now note that in any proper elementary extension of $\mathfrak{M} = (V \times V)^\omega$, one of the two copies of V must grow, but this implies that one of the types σ or $\sigma \frown v$ must be realized in the extension, so there cannot be a Vaughtian pair over L . \square

Of course the theory of this structure has an imaginary Vaughtian pair.

As mentioned before, in discrete logic a superstable theory has no Vaughtian pairs if and only if it has no imaginary Vaughtian pairs (although the same is not true of strictly stable theories [Bou89]). This example shows that the same does not even hold for ω -stable theories in continuous logic.

The construction used in Theorem C.2.9 and Counterexample C.2.10 clearly rely

very heavily on the ability to associate an entire copy of a strongly minimal set to a single element of another strongly minimal set. This seems like something that can only be done with discrete structures, which raises the following question.

Question C.2.11. *Does there exist a countable ω -stable theory T with no Vaughtian pairs that is not inseparably categorical and does not interpret a discrete strongly minimal set? A strongly minimal set at all? In particular is there something like Counterexample C.2.10 whose ‘underlying pregeometries’ are not discrete?*

C.3 Banach and Hilbert Structures

Here we present some relevant counterexamples and in particular we resolve (in the negative) the question of Shelah and Usvyatsov presented at the end of Section 5 of [SU19], in which they ask whether or not the span of a Morley sequence in a minimal wide type is always a type-definable set.

C.3.1 No Infinitary Ramsey-Dvoretzky-Milman Phenomena in General

Unfortunately some elements of the analogy between the Ramsey-Dvoretzky-Milman phenomenon and discrete Ramsey theory do not work. In particular, there is no extension of Dvoretzky’s theorem, and therefore Fact 5.4.3, to $k \geq \omega$, even for a fixed $\varepsilon > 0$. Recall that a linear map $T : X \rightarrow Y$ between Banach spaces is an *isomorphism* if it is a continuous bijection. This is enough to imply that T is invertible and that both T and T^{-1} are Lipschitz. An analog of Dvoretzky’s theorem for $k \geq \omega$ would imply that

every sufficiently large Banach space has an infinite dimensional subspace isomorphic to Hilbert space, which is known to be false. Here we will see a specific example of this.

The following is a well known result in Banach space theory (for a proof see the comment after Proposition 2.a.2 in [LT96]).

Fact C.3.1. *For any distinct $X, Y \in \{\ell_p : 1 \leq p < \infty\} \cup \{c_0\}$, no subspace of X is isomorphic to Y .*

Note that, whereas Corollary 5.4.10 says that every Banach theory is consistent with the partial type of an indiscernible subspace, the following corollary says that this type can sometimes be omitted in arbitrarily large models (contrast this with the fact that the existence of an Erdős cardinal implies that you can find indiscernible sequences in any sufficiently large structure in a countable language [Kan03, Thm. 9.3]).

Corollary C.3.2. *For $p \in [1, \infty) \setminus \{2\}$, there are arbitrarily large models of $\text{Th}(\ell_p)$ that do not contain any infinite dimensional subspaces isomorphic to a Hilbert space.*

Proof. Fix $p \in [1, \infty) \setminus \{2\}$ and $\kappa \geq \aleph_0$. Let $\ell_p(\kappa)$ be the Banach space of functions $f : \kappa \rightarrow \mathbb{R}$ such that $\sum_{i < \kappa} |f(i)|^p < \infty$. Note that $\ell_p(\kappa) \equiv \ell_p$.² Pick a subspace $V \subseteq \ell_p(\kappa)$. If V is isomorphic to a Hilbert space, then any separable $V_0 \subseteq V$ will also be isomorphic to a Hilbert space. There exists a countable set $A \subseteq \kappa$ such that $V_0 \subseteq \ell_p(A) \subseteq \ell_p(\kappa)$. By Fact C.3.1, V_0 is not isomorphic to a Hilbert space, which is a contradiction. Thus no such V can exist. \square

²To see this, we can find an elementary sub-structure of $\ell_p(\kappa)$ that is isomorphic to ℓ_p : Let \mathfrak{L}_0 be a separable elementary sub-structure of $\ell_p(\kappa)$. For each $i < \omega$, given \mathfrak{L}_i , let B_i be the set of all $f \in \ell_p(\kappa)$ that are the indicator function of a singleton $\{i\}$ for some i in the support of some element of \mathfrak{L}_i . B_i is countable. Let \mathfrak{L}_{i+1} be a separable elementary sub-structure of $\ell_p(\kappa)$ containing $\mathfrak{L}_i \cup B_i$. $\overline{\bigcup_{i < \omega} \mathfrak{L}_{i+1}}$ is equal to the span of $\bigcup_{i < \omega} B_i$ and so is a separable elementary sub-structure of $\ell_p(\kappa)$ isomorphic to ℓ_p .

Even assuming we start with a Hilbert space we do not get an analog of the infinitary pigeonhole principle (i.e. a generalization of Fact 5.4.3). The discussion by Hájek and Novotný in [HN18, after Thm. 1] of a result of Maurey [Mau95] implies that there is a Hilbert theory T with a unary predicate P such that for some $\varepsilon > 0$ there are arbitrarily large models \mathfrak{M} of T such that for any infinite dimensional subspace $V \subseteq \mathfrak{M}$ there are unit vectors $a, b \in V$ with $|P^{\mathfrak{M}}(a) - P^{\mathfrak{M}}(b)| \geq \varepsilon$.

Stability of a theory often has the effect of making Ramsey phenomena more prevalent in its models, so there is a natural question as to whether anything similar will happen here. Recall that a function $f : S(X) \rightarrow \mathbb{R}$ on the unit sphere $S(X)$ of a Banach space X is *oscillation stable* if for every infinite dimensional subspace $Y \subseteq X$ and every $\varepsilon > 0$ there is an infinite dimensional subspace $Z \subseteq Y$ such that for any $a, b \in S(Z)$, $|f(a) - f(b)| \leq \varepsilon$.

Question C.3.3. *Does (model theoretic) stability imply oscillation stability? That is to say, if T is a stable Banach theory, is every unary formula oscillation stable on models of T ?*

C.3.2 (Type-)Definability of Indiscernible Subspaces and Complex Banach Structures

A central question in the study of inseparably categorical Banach space theories is the degree of definability of the ‘minimal Hilbert space’ that controls a given inseparable model of the theory. Results of Henson and Raynaud in [HR16] imply that in general the Hilbert space may not be definable. In [SU19], Shelah and Usvyatsov ask whether or not the Hilbert space can be taken to be type-definable or a zeroset. In Counterexample

C.3.5 we present a simple, but hopefully clarifying, example showing that this is slightly too much to ask.

It is somewhat uncomfortable that even in complex Hilbert structures we are only thinking about *real* indiscernible subspaces rather than *complex* indiscernible subspaces. One problem is that Ramsey-Dvoretzky-Milman phenomena only deal with real subspaces in general. The other problem is that Definition 5.4.4 is incompatible with complex structure:

Proposition C.3.4. *Let T be a complex Banach theory. Let V be an indiscernible subspace in some model of T . For any non-zero $a \in V$ and $\lambda \in \mathbb{C} \setminus \{0\}$, if $\lambda a \in V$, then $\lambda \in \mathbb{R}$.*

Proof. Assume that for some non-zero vector a , both a and ia are in V . We have that $(a, ia) \equiv (ia, a)$, but $(a, ia) \models d(ix, y) = 0$ and $(ia, a) \not\models d(ix, y) = 0$, which contradicts indiscernibility. Therefore we cannot have that both a and ia are in V . The same statement for a and λa with $\lambda \in \mathbb{C} \setminus \mathbb{R}$ follows immediately, since $a, \lambda a \in V \Rightarrow ia \in V$. □

In the case of complex Hilbert space and other Hilbert spaces with a unitary Lie group action, this is the reason that indiscernible subspaces can fail to be type-definable. We will explicitly give the simplest example of this.

Counterexample C.3.5. *Let T be the theory of an infinite dimensional complex Hilbert space, and let \mathfrak{C} be the monster model of T . T is inseparably categorical, but for any partial type Σ over any small set of parameters A , $\Sigma(\mathfrak{C})$ is not an infinite dimensional indiscernible subspace (over \emptyset).*

Verification. T is clearly inseparably categorical by the same reasoning that the theory of real infinite dimensional Hilbert spaces is inseparably categorical (being an infinite dimensional complex Hilbert space is first-order and there is a unique infinite dimensional complex Hilbert space of each infinite density character).

If $\Sigma(\mathfrak{C})$ is not an infinite dimensional subspace of \mathfrak{C} , then we are done, so assume that $\Sigma(\mathfrak{C})$ is an infinite dimensional subspace of \mathfrak{C} . Let \mathfrak{N} be a small model containing A . Since \mathfrak{N} is a subspace of \mathfrak{C} , $\Sigma(\mathfrak{N}) = \Sigma(\mathfrak{C}) \cap \mathfrak{N}$ is a subspace of \mathfrak{N} . Let $v \in \Sigma(\mathfrak{C}) \setminus \Sigma(\mathfrak{N})$. This implies that $v \in \mathfrak{C} \setminus \mathfrak{N}$, so we can write v as $v_{\parallel} + v_{\perp}$, where v_{\parallel} is the orthogonal projection of v onto \mathfrak{N} and v_{\perp} is complex orthogonal to \mathfrak{N} . Necessarily we have that $v_{\perp} \neq 0$. Let \mathfrak{N}^{\perp} be the orthocomplement of \mathfrak{N} in \mathfrak{C} . If we write elements of \mathfrak{C} as (x, y) with $x \in \mathfrak{N}$ and $y \in \mathfrak{N}^{\perp}$, then the maps $(x, y) \mapsto (x, -y)$, $(x, y) \mapsto (x, iy)$, and $(x, y) \mapsto (x, -iy)$ are automorphisms of \mathfrak{C} fixing \mathfrak{N} . Therefore $(v_{\parallel} + v_{\perp}) \equiv_{\mathfrak{N}} (v_{\parallel} - v_{\perp}) \equiv_{\mathfrak{N}} (v_{\parallel} + iv_{\perp}) \equiv_{\mathfrak{N}} (v_{\parallel} - iv_{\perp})$, so we must have that $(v_{\parallel} - v_{\perp}), (v_{\parallel} + iv_{\perp}), (v_{\parallel} - iv_{\perp}) \in \Sigma(\mathfrak{C})$ as well. Since $\Sigma(\mathfrak{C})$ is a subspace, we have that $b_{\perp} \in \Sigma(\mathfrak{C})$ and $ib_{\perp} \in \Sigma(\mathfrak{C})$. Thus by Proposition C.3.4, $\Sigma(\mathfrak{C})$ is not an indiscernible subspace over \emptyset . \square

This example is a special case of this more general construction: If G is a compact Lie group with an irreducible unitary representation on \mathbb{R}^n for some n (i.e. the group action is transitive on the unit sphere), then we can extend this action to ℓ_2 by taking the Hilbert space direct sum of countably many copies of the irreducible unitary representation of G , and we can think of this as a structure by adding function symbols for the elements of G . The theory of this structure will be totally categorical and satisfy the conclusion of Counterexample C.3.5.

Counterexample C.3.5 is analogous to the fact that in many strongly minimal theories the set of generic elements in a model is not itself a basis/Morley sequence. The immediate response would be to ask the question of whether or not the unit sphere of the complex linear span (or more generally the ‘ G -linear span,’ i.e. the linear span of $G \cdot V$) of the indiscernible subspace in a minimal wide type agrees with the set of realizations of that minimal wide type, but this can overshoot:

Counterexample C.3.6. *Consider the structure whose universe is (the unit ball of) $\ell_2 \oplus \ell_2$ (where we are taking ℓ_2 as a real Hilbert space), with a complex action $(x, y) \mapsto (-y, x)$ and orthogonal projections P_0 and P_1 for the sets $\ell_2 \oplus \{0\}$ and $\{0\} \oplus \ell_2$, respectively. Let T be the theory of this structure. This is a totally categorical complex Hilbert structure, but for any complete type p and $\mathfrak{M} \models T$, $p(\mathfrak{M})$ does not contain the unit sphere of a non-trivial complex subspace.*

Verification. T is bi-interpretable with a real Hilbert space, so it is totally categorical. For any complete type p , there are unique values of $\|P_0(x)\|$ and $\|P_1(x)\|$ that are consistent with p , so the set of realizations of p in any model cannot contain $\{\lambda a\}_{\lambda \in U(1)}$ for a , a unit vector, and $U(1) \subset \mathbb{C}$, the set of unit complex numbers. \square

The issue, of course, being that, while we declared by fiat that this is a complex Hilbert structure, the expanded structure does not respect the complex structure.

So, on the one hand, Counterexample C.3.6 shows that in general the unit sphere of the complex span won’t be contained in the minimal wide type. On the other hand, a priori the set of realizations of the minimal wide type could contain more than just the unit sphere of the complex span, such as if we have an $SU(n)$ action. The complex (or G -linear) span of a set is of course part of the algebraic closure of the set in question,

so this suggests a small refinement of the original question of Shelah and Usvyatsov:

Question C.3.7. *If T is an inseparably categorical Banach theory, p is a minimal wide type, and \mathfrak{M} is a model of T which is prime over an indiscernible subspace V in p , does it follow that $p(\mathfrak{M})$ is the unit sphere of a subspace contained in the algebraic closure of V ?*

This would be analogous to the statement that if p is a strongly minimal type in an uncountably categorical discrete theory and \mathfrak{M} is a model prime over a Morley sequence I in p , then $p(\mathfrak{M}) \subseteq \text{acl}(I)$.

C.3.3 Non-minimal Wide Types

The following example shows, unsurprisingly, that Theorem 5.4.11 does not hold for non-minimal wide types.

Counterexample C.3.8. *Let T be the theory of (the unit ball of) the infinite Hilbert space sum $\ell_2 \oplus \ell_2 \oplus \dots$, where we add a predicate D that is the distance to $S^\infty \sqcup S^\infty \sqcup \dots$, where S^∞ is the unit sphere of the corresponding copy of ℓ_2 . This theory is ω -stable. The partial type $\{D = 0\}$ has a unique global non-forking extension p that is wide, but the unit sphere of the linear span of any Morley sequence in p is not contained in $p(\mathfrak{C})$.*

Verification. This follows from the fact that on D the equivalence relation ‘ x and y are contained in a common unit sphere’ is definable by a formula, namely

$$E(x, y) = \inf_{z, w \in D} (d(x, z) \div 1) + (d(z, w) \div 1) + (d(w, y) \div 1),$$

where $a \div b = \max\{a - b, 0\}$. If x, y are in the same sphere, then let S be a great circle

passing through x and y and choose z and w evenly spaced along the shorter path of S . It will always hold that $d(x, z), d(z, w), d(w, y) \leq 1$, so we will have $E(x, y) = 0$. On the other hand, if x and y are in different spheres, then $E(x, y) = \sqrt{2} - 1$.

Therefore a Morley sequence in p is just any sequence of elements of D which are pairwise non- E -equivalent and the unit sphere of the span of any such set is clearly not contained in D . \square

C.3.4 Minimal Wide Types That are Not Strongly Minimal Wide

The following two examples are analogous to the simplest examples of weakly minimal theories in discrete logic, namely the superstable theory of a countable sequence of independent unary predicates, each infinite and co-infinite, and the ω -stable theory of a countable sequence of disjoint infinite and co-infinite unary predicates.

Counterexample C.3.9 (Stable Theory with No Strongly Minimal Wide Types). *Let \mathfrak{A} be (the unit ball of) the complex Hilbert space $L_2[0, 1]$ (i.e. square integrable functions on $[0, 1]$ with the standard Lebesgue measure) together with the linear operator X defined by $(Xf)(x) = xf(x)$.³ Let T be the theory of \mathfrak{A} .*

The theory T is superstable, but not ω -stable, and has no strongly minimal wide types.

Verification. $X^{\mathfrak{A}}$ is clearly a self-adjoint bounded operator whose spectrum is $[0, 1]$ (as a subset of \mathbb{C}). The self-adjunction of X is clearly first-order. Recall that an operator is invertible if and only if it is bounded below and has dense image. In a Hilbert space,

³ \mathfrak{A} can be either a real or a complex Hilbert space, but accounts of the spectral theorem are easier to find for complex Hilbert spaces.

if an operator fails to have dense image then this is witnessed by some unit vector that is orthogonal to the image of the operator (any element of the orthocomplement of the image). For any complex number λ , consider the sentence

$$\varphi_\lambda = \min \left\{ \inf_x \|Xx - \lambda x\| + (1 \div \|x\|), \sum_{i=1}^{\infty} 2^{-i} \inf_x \sup_y 1 \div d(x, s_i(Xy)) \right\}.$$

Semantically, if \mathfrak{M} is a Hilbert structure in which X is a bounded operator, $\mathfrak{M} \models \varphi_\lambda = 0$ if and only if either $X - \lambda I$ fails to be bounded below (the first line) or $X - \lambda I$ fails to have dense image (the second line). So for any fixed λ , we have that $\mathfrak{B} \models \varphi_\lambda = 0$ if and only if λ is in the spectrum of $X^{\mathfrak{B}}$. So if $\mathfrak{B} \models T$, then $X^{\mathfrak{B}}$ must have $[0, 1]$ as its spectrum.

We want to characterize the type space $S_1(C)$ for any set of parameters C in the monster model \mathfrak{C} of T . We may assume without loss of generality that C is a linear subspace that is closed under X (this is all contained in $\text{dcl}(C)$).

Note that to characterize the types in $S_1(C)$, it is enough to characterize types $p(x)$ which entail $\langle x, c \rangle = 0$ for every $c \in C$ as well as $\|x\| = 1$. Call such types *orthonormal*.⁴ This is because every type $p \in S_1(C)$ has a unique $c \in C$ for which $d(p, c)$ is minimal (this, as well as everything else in this paragraph, is true in any Hilbert structure and over any set of parameters which is a subspace). If a is some realization of p , then, assuming p is not realized in C , we can consider the type of the vector $\frac{a-c}{\|a-c\|}$. Let q be the type of this vector. We have that p and q are interdefinable over C . This means that once we have a characterization of the orthonormal types, we get that every type p in $S_1(C)$ that is not realized in C is uniquely determined by a triple (c, α, q)

⁴Note that wide types are necessarily orthonormal, but the converse does not hold.

with $\|c\|^2 + |\alpha|^2 \leq 1$, $\alpha \neq 0$, and q an orthonormal type. Realizations of p precisely correspond to vectors of the form $c + \alpha f$, with f a realization of q . Note that if $O(C)$ is the set of orthonormal types in $S_1(C)$ (which is closed), then this argument establishes that the metric density character of $S_1(C)$ (with regards to the d -metric) is no greater than $\#^{\text{dc}}C + \#^{\text{dc}}O(C)$, where $\#^{\text{dc}}X$ is the density character of X . We will show that $\#^{\text{dc}}O(C) \leq 2^{\aleph_0}$, regardless of the choice of C , which will establish that T is superstable.

Any vector $f \in \mathfrak{C}$ induces a linear map $g \mapsto \langle f, g(X)f \rangle$ on the space of complex polynomials g . By the spectral theorem we know that this map extends uniquely to a map on the space of continuous functions from $[0, 1]$ to \mathbb{C} . Furthermore, we know that this map takes the constant function 1 to 1, has operator norm 1, and is positive semi-definite, so by the Riesz representation theorem there is a unique Borel probability measure μ_f on $[0, 1]$ such that for any polynomial g , $\langle f, g(X)f \rangle = \int g d\mu_f$. Clearly the type of f fixes μ_f . We want to show that for any Borel probability measure (B.p.m.)⁵ ν on $[0, 1]$ there is an orthonormal type p such that $\mu_f = \nu$ for any realization f of p and that an orthonormal type p is determined by μ_f for any f realizing p .

To show that such types exist for any B.p.m. ν on $[0, 1]$, let \mathfrak{B} be a model containing \mathfrak{A} and C , fix a non-principal ultrafilter \mathcal{U} on ω , and consider the ultrapower $\mathfrak{B}^{\mathcal{U}}$.

Construct a sequence of finite sets of intervals $\{I_n\}_{n < \omega}$ (where we require that intervals have positive length) with the following properties.

1. For each $n < \omega$, and each interval of the form $[i2^{-n}, (i+1)2^{-n}]$ for some $i < 2^n$, there is precisely one interval $J \in I_n$ satisfying $J \subseteq [i2^{-n}, (i+1)2^{-n}]$.
2. These are the only intervals in each I_n .

⁵Note that a Borel probability measure on a metric space is automatically regular.

3. For $n \neq m$, $J \in I_n$, and $K \in I_m$, $J \cap K = \emptyset$.

It is not hard to construct such a sequence. Let I_n^i be the interval in I_n contained in $[i2^{-n}, (i+1)2^{-n}]$. For each $n < \omega$, let $f_n(x) = \sum_{i < 2^n} \nu([i2^{-n}, (i+1)2^{-n}]) \chi_{I_n^i}(x)$, where $\chi_J(x)$ is the indicator function of J .

It is clear that for each $n < \omega$, f_n is a non-negative element of $L_1[0, 1]$ with norm 1. Furthermore, for any distinct $n, m < \omega$, $f_n f_m = 0$. It is also not hard to show that for any fixed polynomial $p(x)$, $\int p(x) f_n(x) d\lambda \rightarrow \int p(x) d\nu$ as $n \rightarrow \infty$, where λ is the Lebesgue measure on $[0, 1]$.

Let g be the element of \mathfrak{B}^u corresponding to the sequence $\{\sqrt{f_n}\}_{n < \omega}$. By construction we have that $\langle g, p(X)g \rangle = \int p(x) d\nu$ for every polynomial $p(x)$. This implies that $\mu_g = \nu$. Also note that $\|g\| = 1$.

Now to show that $\text{tp}(g/C)$ is orthonormal, for each $n < \omega$, let $h_n = \sqrt{f_n}$ (thought of as an element of \mathfrak{B}). Let $H = \overline{\text{span}}\{h_n\}_{n < \omega}$. For any $c \in C$, let c_{\parallel} be the orthogonal projection of c onto H . Since the sequence $\langle h_n, c \rangle = \langle h_n, c_{\parallel} \rangle$ is square summable, it limits to 0. Therefore we have that $\langle g, c \rangle = 0$. Since we can do this for any $c \in C$, g realizes an orthonormal type over C . (It actually realizes an orthonormal type over all of \mathfrak{B} . Also note that in this paragraph we really only used the fact that h_n is an orthonormal sequence.)

So we have that there is an orthonormal type for every B.p.m. on $[0, 1]$.

If $p \in S_1(C)$ is an orthonormal type, we'll write μ_p for μ_f for some realization f of p (μ_p doesn't depend on the choice of f). We now need to show that for orthonormal types, p is uniquely determined by μ_p .

Fix orthonormal $p, q \in S_1(C)$ satisfying $\mu_p = \mu_q$. Let \mathfrak{B}_0 be an elementary extension of \mathfrak{B} realizing both p and q . Let f and g be these realizations. We are going to construct

an elementary chain starting with \mathfrak{B}_0 .

Recall that there are 2^{\aleph_0} many B.p.m.s on $[0, 1]$. Let $\{\nu_i\}_{i < 2^{\aleph_0}}$ be an enumeration of the B.p.m.s on $[0, 1]$. Let $\kappa = (2^{\aleph_0} \cdot \#^{\text{dc}}\mathfrak{B}_0)^+$. By basic cardinal arithmetic, $\kappa = 2^{\aleph_0} \cdot \kappa$, so we may regard ordinals $i < \kappa$ as an ordered pair (j, k) , with $j < 2^{\aleph_0}$ and $k < \kappa$. Given \mathfrak{B}_i , construct \mathfrak{B}_{i+1} as an elementary extension of \mathfrak{B}_i that satisfies the following conditions.

- \mathfrak{B}_{i+1} has density character at most κ .
- \mathfrak{B}_{i+1} realizes an orthonormal type $p \in S_1(\mathfrak{B}_i)$ satisfying $\mu_p = \nu_j$, where j is the first element of the ordered pair (j, k) corresponding to i .

For $i < \kappa$ a limit ordinal, let $\mathfrak{B}_i = \overline{\bigcup_{j < i} \mathfrak{B}_j}$. Note that since κ is a regular cardinal, $\#^{\text{dc}}\mathfrak{B}_i \leq \kappa$.

Finally let $\mathfrak{B}_\kappa = \bigcup_{i < \kappa} \mathfrak{B}_i$. Since κ is regular and uncountable, \mathfrak{B}_κ is complete. We are going to construct an automorphism of \mathfrak{B}_κ taking f to g . Let $\{b_i\}_{i < \kappa}$ be an enumeration of a dense subset of \mathfrak{B}_κ .

Let $f_0 = f$ and $g_0 = g$. We will proceed with a back-and-forth construction of length κ .

For any set of vectors A , let $\overline{\text{span}}_X A$ be the smallest closed subspace containing A and closed under X . It is clear that the dimension of $\overline{\text{span}}_X A$ is no greater than $\aleph_0 + \#^{\text{dc}} A$. For any $i \leq \kappa$, let $f_{<i}$ represent the sequence $\{f_j\}_{j < i}$ and likewise for $g_{<i}$.

At even stage $i \cdot 2 < \kappa$ with $i > 0$, given $f_{<i}$ and $g_{<i}$, find the first element of $\{b_i\}_{i < \kappa}$ not contained in $\overline{\text{span}}_X(\mathfrak{B}_0 f_{<i})$. Call this element b_k . Find a unit vector f_i , orthogonal to $\overline{\text{span}}_X(\mathfrak{B}_0 f_{<i})$, such that $b_k \in \overline{\text{span}}(\{f_i\} \cup \overline{\text{span}}_X(\mathfrak{B} f_{<i}))$. This is always possible by construction. (Also note that, up to multiplication by a unit norm complex number, f_i

is unique.) The odd stages proceed analogously.

Given the full sequences $f_{<\kappa}$ and $g_{<\kappa}$, by construction we have ensured that $\overline{\text{span}}_X f_{<\kappa} = \overline{\text{span}}_X g_{<\kappa} = \mathfrak{B}_\kappa$. Let σ_0 be the partial function from \mathfrak{B}_κ to \mathfrak{B}_κ mapping f_i to g_i for each $i < \kappa$ as well as h to itself for each $h \in \mathfrak{B}_0$. By orthogonality, this clearly extends to a linear map on $\overline{\text{span}}(\mathfrak{B}_0 f_{<\kappa})$. Furthermore, the theory T says that X is a self-adjoint operator, which implies that for any vector h in a model of T and any $n, m < \omega$, $\langle X^n h, X^m h \rangle = \langle h, X^{n+m} h \rangle = \int x^{n+m} d\mu_h(x)$. So since these inner products only depend on μ_{f_i} , we have that σ_0 extends to a linear map that respects X on all of $\overline{\text{span}}_X(\mathfrak{B}_0 f_{<\kappa})$, where by respecting X we mean that $\sigma(Xh) = X\sigma(h)$. Let σ be this extension. By construction, the domain of σ is all of \mathfrak{B}_κ and the range of σ is also \mathfrak{B}_κ . Since σ is a linear map that respects X , it is an automorphism of \mathfrak{B}_κ . Furthermore, we have that $\sigma(f_0) = g_0$, and so we have that $\text{tp}(f_0/\mathfrak{B}_0) = \text{tp}(g_0/\mathfrak{B}_0)$, and so in particular $\text{tp}(f/C) = \text{tp}(g/C)$, as required.

We have now established that $O(C)$, the set of orthonormal types over C , does not depend on the choice of C , as it always corresponds to the set of B.p.m.s on $[0, 1]$. This set has cardinality 2^{\aleph_0} and so in particular has $\#^{\text{dc}} O(C) \leq 2^{\aleph_0}$. Thus T is superstable.

To show that T is not ω -stable we need to show that there is an $\varepsilon > 0$ and an uncountable set of types in $O(C)$ which are $(>\varepsilon)$ -separated for any set of parameters C . For each $r \in [0, 1]$, let δ_r be the Dirac measure centered at r . If we let p_r be the type such that $\mu_{p_r} = \delta_r$, then we have that realizations of p_r are eigenvectors of X with eigenvalue r . Since X is a self-adjoint operator, eigenvectors with distinct eigenvalues are necessarily orthogonal, therefore we have that $d(p_r, p_s) = \sqrt{2}$ for any distinct $r, s \in [0, 1]$. Therefore no type space $S_1(C)$ has countable metric density character, and T is not ω -stable.

Now, finally, we need to show that T has no strongly minimal wide types. To do this

it is sufficient to show that no $p \in O(C)$ is relatively d -atomic. Given $p \in O(C)$, we will construct a sequence $q_i \in O(C)$ limiting to p topologically but for which there is some $\varepsilon > 0$ such that $d(p, q_i) > \varepsilon$ for all $i < \omega$.

Claim: The logical topology on $O(C)$ agrees with the topology of weak convergence of measure (i.e. the coarsest topology for which $\mu \mapsto \int f d\mu$ is continuous for each continuous function $f : [0, 1] \rightarrow \mathbb{R}$), thinking of the types as their corresponding B.p.m.s.

Proof of claim. It is clear that the logical topology on $O(C)$ is finer than the topology of weak convergence, since for any polynomial $f(x)$, the function $O(C) \rightarrow \mathbb{R}$ defined by $p \mapsto \int f(x) d\mu_p$ is continuous, and these functions generate the topology of weak convergence. Since polynomials are dense in the space of continuous functions on $[0, 1]$ under the uniform norm, we have that $p \mapsto \int f(x) d\mu_p$ is continuous for each continuous $f : [0, 1] \rightarrow \mathbb{R}$. Now the fact that the topologies agree follows from the fact that if τ_0 and τ_1 are two compact Hausdorff topologies on the same set such that $\tau_0 \subseteq \tau_1$, then in fact $\tau_0 = \tau_1$. □_{claim}

So now given $p \in O(C)$ we just need to construct a sequence of B.p.m.s $\{\nu_i\}_{i < \omega}$ converging weakly to μ_p , but which are all singular with regards to it, which will guarantee that $\|\nu_i - \mu_p\|_{\text{tv}} = 2$ for all $i < \omega$, where $\|\cdot\|_{\text{tv}}$ is the total variation norm. Then we will use this to give a lower bound on $d(p, q_i)$, where q_i satisfies $\mu_{q_i} = \nu_i$, for all $i < \omega$ (in particular we will show that any realization of p and any realization of q_i must be orthogonal).

Let A be the set of all $r \in [0, 1]$ such that $\mu_p(\{r\}) > 0$. Since μ_p is a probability measure, A can be at most countable. For each $i < \omega$, let ν_i be a finite sum of Dirac measures, $\sum_{k < 2^i} \mu_p([k2^{-i}, (k+1)2^{-i})) \delta_{r_k^i}$, with each r_k^i chosen so that $r_k^i \in [k2^{-i}, (k+1)2^{-i}) \setminus A$. This is always possible since A is at most countable.

Clearly we have by construction that ν_i limits to μ_p weakly, so now we just need to show that that realizations of p and q_i must actually be orthogonal. Let a be a realization of p and b a realization of q_i . For any $\varepsilon > 0$, find an open set $U \subset [0, 1]$ not containing any r_k^i for any $i < 2^k$ such that $\mu_p(U) > 1 - \varepsilon$. Find a real polynomial $f(x)$ such that for any $x \in U$, $|f(x) - 1| < \varepsilon$; for each $i < 2^n$, $|f(r_k^i)| < \varepsilon$; and for each $x \in [0, 1]$, $-\varepsilon < f(x) < 1 + \varepsilon$.

Consider $\langle a, f(X)b \rangle$. Note that since f is a real polynomial, $f(X)$ is a self-adjoint operator, so we have that $\langle a, f(X)b \rangle = \langle f(X)a, b \rangle$. $|\langle f(X)b, f(X)b \rangle| = |\langle b, f(X)^2b \rangle| = |\int f^2 d\nu_i| < \varepsilon^2$, so $\|f(X)b\| < \varepsilon$. Since a is a unit vector, this implies that $|\langle a, f(X)b \rangle| = |\langle f(X)a, b \rangle| < \varepsilon$ as well. Now consider

$$\begin{aligned} \langle a - f(X)a, a - f(X)a \rangle &= \langle a, a \rangle - \langle a, f(X)a \rangle - \langle f(X)a, a \rangle + \langle f(X)a, f(X)a \rangle \\ &= 1 - 2\langle a, f(X)a \rangle + \langle a, f(X)^2a \rangle \\ &= 1 - 2 \int f d\mu_p + \int f^2 d\mu_p \\ &= 1 + \int f^2 - 2f d\mu_p. \end{aligned}$$

Considering the U and $[0, 1] \setminus U$ parts of the integral, we get that $|\langle a - f(X)a, a - f(X)a \rangle| < (\varepsilon + \varepsilon^2) + \varepsilon(2 + 2\varepsilon + 1 + 2\varepsilon + \varepsilon^2) = 4\varepsilon + 5\varepsilon^2 + \varepsilon^3$.

So now, since b is a unit vector, we have that

$$\begin{aligned} |\langle a, b \rangle| &\leq |\langle a - f(X)a, b \rangle| + |\langle f(X)a, b \rangle| \\ &< 4\varepsilon + 5\varepsilon^2 + \varepsilon^3 + \varepsilon = 5\varepsilon + 5\varepsilon^2 + \varepsilon^3. \end{aligned}$$

Since we can do this for any $\varepsilon > 0$, we have that $\langle a, b \rangle = 0$ for any $a \models p$ and $b \models q_i$, as

required.

Since this is true of each q_i , we have that the sequence $\{q_i\}_{i<\omega}$ does not limit to p in the d -metric and thus p is not relatively d -atomic in $O(C)$.

Since this is true for any $p \in O(C)$ and since all wide types in $S_1(C)$ are contained in $O(C)$, T has no strongly minimal wide types. \square

It is also possible to construct an example of an ω -stable theory with a minimal wide type that is not strongly minimal wide.

Counterexample C.3.10. *Let \mathfrak{A} be (the unit ball of) the infinite Hilbert space sum $\ell_2 \oplus \ell_2 \oplus \dots$ together with a linear operator M defined so that for any vector a confined to the n th summand, $Ma = 2^{-n}a$. In other words, M is a linear functional on a infinite dimensional Hilbert space whose eigenvalues are of the form 2^{-n} and each have an infinite dimensional eigenspace.*

Let T be the theory of \mathfrak{A} . T is ω -stable and has a unique type $p \in S_1(T)$ with the property that any realization a of p has norm 1 and satisfies $Ma = 0$. This type is minimal wide but not strongly minimal wide.

Verification. It is not hard to verify that if \mathfrak{B} is any model of T , then \mathfrak{B} is a Hilbert space that decomposes into a direct sum of eigenspaces $\bigoplus_{\lambda \in \{0, 2^{-0}, 2^{-1}, 2^{-2}, \dots\}} V_\lambda$, where V_λ is the eigenspace of M corresponding to the eigenvalue λ .⁶ It is also not hard to verify that for any such model each space V_λ for $\lambda > 0$ must be infinite dimensional and that models of T are precisely of this form. The approximately ω -saturated separable model of T is the one in which each V_λ has countably infinite dimension.⁷ By looking at the

⁶Note that this relies on the fact that the spectrum of M is countable. The analogous statement is not true for operators such as X in Counterexample C.3.9, as witnessed by the model \mathfrak{A} .

⁷It is actually ω -saturated, not just approximately ω -saturated.

automorphisms of this structure, we can conclude that elements of $S_1(T)$ are in a one-to-one correspondence with square summable functions $f : \{0, 2^{-0}, 2^{-1}, 2^{-2}, \dots\} \rightarrow [0, 1]$ and that the d -metric corresponds to the ℓ_2 metric on this set. We have that the type corresponding to the function g defined by $g(0) = 1$ and $g(2^{-n}) = 0$ for all $n < \omega$ is axiomatized by $Mx = \mathbf{0}$ (i.e. $d(Mx, \mathbf{0}) = 0$). Let $p(x)$ be this type. It is clear that this type is wide.

The space of global types $S_1(\mathfrak{C})$ has a one-to-one correspondence with pairs (c, f) , where $c \in \mathfrak{C}$ and f is a function as before, satisfying $\|c\|_2^2 + \|f\|_2^2 \leq 1$. Furthermore, we have that if q and r correspond to (c, f) and (e, h) , respectively, then $d(q, r) = \sqrt{\|c - e\|_2^2 + \|f - h\|_2^2}$. Furthermore, we have that if q and r correspond to (c, f) and (e, h) , respectively, then $d(q, r) = \sqrt{\|c - e\|_2^2 + \|f - h\|_2^2}$. (This characterization is true of the space of types over any model, not just \mathfrak{C} , which allows us to show that T is ω -stable.)

No type corresponding to a pair (c, f) with $c \neq \mathbf{0}$ can be wide. This implies that the only wide global extension of p is the type corresponding to $(0, g)$, and hence that p is minimal wide. We also have that the types corresponding to pairs of the form $(\mathbf{0}, g)$, with $g(\lambda) = 1$ if $\lambda = 2^{-n}$ (for some particular n) and $g(\lambda) = 0$ otherwise, are wide for any $n < \omega$. Call the corresponding type q_n . For any such type we have that $d(p, q_n) = 1$, but nevertheless the sequence $\{q_n\}_{n < \omega}$ limits to p topologically, therefore p is not d -atomic in the set of wide global types, and is therefore not strongly minimal wide. \square

This example is of course in some sense the Hilbert space analog of the theory of the discrete structure \mathfrak{M} in a language consisting of a countable sequence of unary predicates $\{U_i\}_{i < \omega}$ which partition \mathfrak{M} into infinite sets, with the unique minimal type not contained

in any U_i corresponding to $p(x)$. Whereas in discrete logic it is known that any minimal type in an uncountably categorical theory is actually strongly minimal, the analogous statements for minimal and minimal wide types in continuous logic are unknown.

Question C.3.11. *If T is an inseparably categorical Banach theory, is every minimal wide type strongly minimal wide?*

C.4 Approximate Isomorphism and Approximate Categoricity

The simplest non-trivial example of an approximately ω -categorical theory is easiest to describe in terms of a stratified language (Definition 6.2.12).

Counterexample C.4.1. *Let $\mathcal{L}_0 = \{a_i\}_{i < \omega} \cup \{b_0\}$, and let $\mathcal{L}_{n+1} = \mathcal{L}_n \cup \{b_{n+1}\}$. Consider the \mathcal{L} -structure in which all of the constants are assigned to different elements and there are no other elements. The theory of this structure is approximately ω -categorical. We can describe the distortion system corresponding to this stratified language easily as the one generated by $\Delta_0 = \{d(x, a_i)\}_{i < \omega} \cup \{2^{-j}d(x, b_j)\}_{j < \omega}$, where we're taking the metric to be $\{0, 1\}$ -valued. Let $\Delta = \overline{\Delta_0}$.*

This serves as a counterexample to three things. It is an example of a Δ - ω -categorical theory that is not ω -categorical. It is an example of a theory in a stratified language that is approximately ω -categorical but which does not have an ω -categorical reduct $T \upharpoonright \mathcal{L}_n$ for any n . And finally it is an example of the failure of the converse of Theorem 6.6.13 part (iii). The unique 1-type of a non-realized element is weakly d_Δ -atomic, but not d_Δ -atomic, so $S_1(T)$ cannot be metrically compact with regards to d_Δ .

A more subtle hope would be that we could eliminate the parameters from Theorem 6.6.13. Unfortunately this is impossible in general, but an example is more complicated.

Counterexample C.4.2. *A theory T with distortion system Δ such that for every $n < \omega$, every types $p \in S_n(T)$ is weakly d_Δ -atomic, but such that T is not Δ - ω -categorical.*

Proof. Let E be a binary relation, let $\{f_i\}_{i < \omega}$ be a sequence of unary functions, and let $\{a_{i,j}\}_{i,j < \omega}$ be an array of constant symbols.

Let $\mathcal{L}_0 = \{E\} \cup \{f_i\}_{i < \omega} \cup \{a_{0,0}\}$. For each $k < \omega$, let $\mathcal{L}_{k+1} = \mathcal{L}_k \cup \{a_{i,j}\}_{i,j \leq k+1}$. Finally let $\mathcal{L} = \bigcup_{k < \omega} \mathcal{L}_k$.

Let \mathfrak{M} be an \mathcal{L} structure with universe $\omega \times (\omega + \omega)$. Set $(i, j)E^{\mathfrak{M}}(k, \ell)$ if and only if $i = k$. Let $f_k((i, j)) = (i, k)$ and $a_{i,j} = (i, \omega + j)$. Finally let $T = \text{Th}(\mathfrak{M})$, and let Δ be a distortion system equivalent to the stratified language \mathcal{L} .

To see that T is not Δ - ω -categorical, let \mathfrak{N} be an elementary extension of \mathfrak{M} that realizes a new E -equivalence class, but only the outputs of the functions f_i in that class. \mathfrak{M} and \mathfrak{N} are not Δ -approximately isomorphic since there is nothing to correlate the constants $a_{i,j}$ in \mathfrak{M} with in the new equivalence class in \mathfrak{N} .

To see that for every $n < \omega$ and every $p \in S_n(T)$, p is weakly d_Δ -atomic, let \mathfrak{D} be the countable elementary extension of \mathfrak{M} that realizes infinitely many new E -equivalence classes and realizes infinitely many elements not equal to any $f_i(c)$ or $a_{i,j}$ in every E -equivalence class. This is the unique countable ω -saturated model of this theory. We would like to show that the type of any finite tuple of elements of this model is weakly d_Δ -atomic. It is enough to consider tuples of the following form. Fix $N < \omega$. Consider the tuple whose elements are

- the constants $a_{i,j}$ for $i, j \leq N$,

- N elements E equivalent to $a_{i,0}$ for each $i \leq N$ but not equal to any $f_i(c)$ for any c ,
- N elements not equal to any $f_i(c)$ from N distinct E equivalence classes that contain no elements of \mathfrak{M} (N^2 elements in total), and
- the image of all of those elements under f_i for $i \leq N$.

Let the type of this tuple be p_N . Every finitary \emptyset -type is the type of some sub-tuple of p_N for some N . The restriction of a weakly d_Δ -atomic type to some sub-tuple is still d_Δ -atomic (since projection maps are open and 1-Lipschitz with regards to d_Δ). For any k , the type $p_N \upharpoonright \mathcal{L}_k$ is realized by a tuple in \mathfrak{M} , and therefore in any model of T , since \mathfrak{M} is the prime model. Therefore by Proposition 6.6.12, p_N is weakly d_Δ -atomic. \square

C.4.1 An Irregular Distortion System

Here we will give some explicit examples of the pathological behavior of a_Δ and the separation of $r_\infty^{\Delta,\Omega}$ and ρ_Δ for irregular Δ .

Definition C.4.3. Let \mathcal{L} be the single-sorted language with a $[0, 1]$ -valued metric and a single $[0, 1]$ -valued 1-Lipschitz unary predicate U .

Let $\text{GH}_0 = \{\frac{1}{2}d(x, y)\}$, and let $\text{GH} = \overline{\text{GH}_0}$.

Let $\text{IU}_0 = \text{GH}_0 \cup \{nU(x)\}_{n < \omega}$, and let $\text{IU} = \overline{\text{IU}_0}$. \triangleleft

Clearly GH corresponds to the ordinary Gromov-Hausdorff distance, ignoring U . The notion of approximate isomorphism induced by IU is strange.

Proposition C.4.4. Let $\mathfrak{M}, \mathfrak{N}$ be \mathcal{L} -structures, and let $R \subseteq \mathfrak{M} \times \mathfrak{N}$.

$\text{dis}_{\text{IU}}(R) \leq \varepsilon < \infty$ if and only if $\text{dis}_{\text{GH}}(R) \leq \varepsilon$ and for all $(a, b) \in R$, $U^{\mathfrak{M}}(a) = U^{\mathfrak{N}}(b)$.

So R needs to be a correlation that rigidly obeys U but which is as loose as the Gromov-Hausdorff distance for d .

Definition C.4.5. For any $D \subseteq [0, 1]$ and $\varepsilon \in [0, 1]$, let $\mathfrak{J}(D, \varepsilon)$ be the \mathcal{L} -pre-structure whose universe is D with $U^{\mathfrak{J}(D, \varepsilon)}(x) = x$ for $x \in D$ and $d^{\mathfrak{J}(D, \varepsilon)}(x, y) = |x - y| \uparrow \varepsilon$ for $x, y \in D$ with $x \neq y$. \triangleleft

Proposition C.4.6. $(\text{Mod}(T, \leq \omega), \rho_{\text{IU}})$ is not complete as a metric space, where T is the empty theory in the language \mathcal{L} .

Proof. Let $X = \{2^{-i}\}_{i < \omega}$, and consider the sequence of structures $\{\mathfrak{J}(X, 2^{-k})\}_{k < \omega}$. These converge in ρ_{IU} to the pre-structure $\mathfrak{J}(X, 0)$, but the corresponding completion is $\mathfrak{J}(X \cup \{0\}, 0)$ and $\rho_{\Delta}(\mathfrak{J}(X, 0), \mathfrak{J}(X \cup \{0\}, 0)) = \infty$. \square

Proposition C.4.7. Let D_0 and D_1 be countable, disjoint, dense subsets of $[0, 1]$, then

$$\rho_{\text{IU}}(\mathfrak{J}(D_0, \varepsilon), \mathfrak{J}([0, 1], 0)) = \infty,$$

$$\rho_{\text{IU}}(\mathfrak{J}(D_1, \varepsilon), \mathfrak{J}([0, 1], 0)) = \infty, \text{ and}$$

$$\rho_{\text{IU}}(\mathfrak{J}(D_0, \varepsilon), \mathfrak{J}(D_1, \varepsilon)) = \infty,$$

but

$$a_{\text{IU}}(\mathfrak{J}(D_0, \varepsilon), \mathfrak{J}([0, 1], 0)) = \frac{1}{2}\varepsilon,$$

$$a_{\text{IU}}(\mathfrak{J}(D_1, \varepsilon), \mathfrak{J}([0, 1], 0)) = \frac{1}{2}\varepsilon, \text{ and}$$

$$a_{\text{IU}}(\mathfrak{J}(D_0, \varepsilon), \mathfrak{J}(D_1, \varepsilon)) = \infty.$$

Furthermore $\mathfrak{I}(D_0, \varepsilon)$, $\mathfrak{I}(D_1, \varepsilon)$, and $\mathfrak{I}([0, 1], 0)$ are all metrically complete and separable.

Proof. For ρ_Δ there simply are no correlations between these structures that satisfy the requirement on U given in Proposition C.4.4. Any almost correlation between $\mathfrak{I}(D_0, \varepsilon)$ and $\mathfrak{I}(D_1, \varepsilon)$ is a correlation anyways, so the same holds for almost correlations.

For $\mathfrak{I}(D_i, \varepsilon)$ and $\mathfrak{I}([0, 1], \varepsilon)$, let $R_i = \{(x, x) : x \in D_i\}$, then we have that $\text{dis}_\Delta(R) \leq \frac{1}{2}\varepsilon$. Finally R_i is clearly an almost correlation so we have that $a_\Delta(\mathfrak{I}(D_i, \varepsilon), \mathfrak{I}([0, 1], 0)) \leq \frac{1}{2}\varepsilon$. To see that they are actually equal to $\frac{1}{2}\varepsilon$, note that these are the *only* almost correlations between these structures with finite IU-distortion. \square

We know that for any given weak modulus Ω something different must happen with $r_\infty^{\text{IU}, \Omega}$ on these structures, because $r_\infty^{\text{IU}, \Omega}$ is a pseudo-metric, so it cannot be equal to a_{IU} , and therefore also cannot be equal to ρ_{IU} . In particular we must have $r_\infty^{\text{IU}, \Omega}(\mathfrak{I}(D_0, \varepsilon), \mathfrak{I}(D_1, \varepsilon)) \leq \varepsilon$.

Finally to get an example of $\rho_\Delta(\mathfrak{M}, \mathfrak{N}) = \infty$ yet $a_\Delta(\mathfrak{M}, \mathfrak{N}) = 0$ for \mathfrak{M} and \mathfrak{N} complete structures, fix a countable dense $D \subseteq [0, 1]$, and let

$$\mathfrak{M} = \bigsqcup_{i < \omega} \mathfrak{I}(D, 2^{-i}) \text{ and}$$

$$\mathfrak{N} = \mathfrak{I}([0, 1], 0) \sqcup \mathfrak{M},$$

where in a disjoint union the distances between things in different structures are always 1. It's easy to check that \mathfrak{M} and \mathfrak{N} are separable \mathcal{L} -structures.

Proposition C.4.8. $\rho_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) = \infty$ but $a_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) = 0$.

Proof. $\rho_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) = \infty$ is clear because there are no correlations between \mathfrak{M} and \mathfrak{N} which can correlate the elements of $[0, 1] \setminus D$ in \mathfrak{N} to anything in \mathfrak{M} while having finite distortion.

To see that $a_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) = 0$, for each $k < \omega$, let $R_k \subseteq \mathfrak{M} \times \mathfrak{N}$ be the almost correlation that relates the copies of $\mathfrak{I}(D, 2^{-i})$ for $i < k$ isomorphically, which relates $\mathfrak{I}(D, 2^{-k}) \subset \mathfrak{M}$ to the subset of $\mathfrak{I}([0, 1], 0) \subset \mathfrak{N}$ whose points are those in D , and which relates $\mathfrak{I}(D, 2^{-i-1}) \subset \mathfrak{M}$ to $\mathfrak{I}(D, 2^{-i}) \subset \mathfrak{N}$ for $i > k$. Then we have that $\text{dis}_{\text{IU}}(R_k) = \text{dis}_{\text{IU}_0}(R_k) \leq 2^{-k-1}$, so $a_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) \leq 2^{-k-1}$ for each $k < \omega$ and $a_{\text{IU}}(\mathfrak{M}, \mathfrak{N}) = 0$. \square

This all raises the question of whether or not $a_{\Delta} = 0$ is an equivalence relation for irregular Δ . In general we have that for any Ω , that $a_{\Delta}(\mathfrak{M}, \mathfrak{N}) = 0$ implies $r_{\infty}^{\Delta, \Omega}(\mathfrak{M}, \mathfrak{N}) = 0$, which is an equivalence relation.

Question C.4.9. *Is the relation $a_{\Delta}(\mathfrak{M}, \mathfrak{N}) = 0$ transitive for irregular Δ ?*

It would be very surprising if the answer were yes.

Appendix D

More on Definable Sets

D.1 Almost Clopen Sets

Clopen formulas and definable formulas are both generalizations of definable sets to the context of continuous logic, but they aren't the only ones.

Definition D.1.1. In a topometric space (X, τ, d) , a set $A \subseteq X$ is *almost clopen* if ∂A (i.e. the topological boundary, $\text{cl } A \setminus \text{int } A$) is metrically nowhere metrically dense. A is *open almost clopen*, or *o.a.c.*, if it is open and almost clopen. A is *closed almost clopen*, or *c.a.c.*, if it is closed and almost clopen. \triangleleft

An easy observation is that if d is uniformly discrete, then a set is almost clopen if and only if it is clopen (because the only metrically nowhere metrically dense set is the empty set). Almost clopen sets do not seem to be as useful in continuous logic as definable sets. They do have one thing over definable sets, however: They are closed under intersections.

Proposition D.1.2. *For any topometric space (X, τ, d) , if A and B are almost clopen, then $A \cup B$, $A \cap B$, $X \setminus A$, $\text{cl } A$, $\text{int } A$, and ∂A are all almost clopen.*

Proof. This follows from the facts that $\partial(A \cup B) \subseteq \partial A \cup \partial B$, $\partial(A \cap B) \subseteq \partial A \cup \partial B$, $\partial(X \setminus A) = \partial A$, $\partial \text{cl } A \subseteq \partial A$, $\partial \text{int } A \subseteq \partial A$, and $\partial \partial A = \partial A$, as well as the facts that

topologically closed sets are metrically closed, a boundary is always topologically closed, and the union of two closed nowhere dense sets in a metric space is nowhere dense. \square

On the other hand, the ‘finite ε ’ version of almost clopen isn’t as well behaved as the ‘finite ε ’ version of definable.

Definition D.1.3. An *almost clopen partition* of X is a triple (U, L, V) of pairwise disjoint subsets of X such that U and V are open, L is closed and metrically nowhere metrically dense, and $X = U \cup L \cup V$.

Open sets U and V form an *almost clopen partition of X* if $(U, X \setminus (U \cup V), V)$ is an almost clopen partition of X . \triangleleft

Proposition D.1.4. (U, L, V) is an almost clopen partition of (X, τ, d) if and only if it is a partition with U, V open, L closed, and for every $\varepsilon > 0$, $L \subseteq U^{<\varepsilon} \cup V^{<\varepsilon}$.

Proof. The \Rightarrow direction follows from the fact that the assumed property implies that L is metrically nowhere metrically dense.

For the \Leftarrow direction, assuming that L is metrically nowhere metrically dense, we have that $U \cup V$ is metrically dense, so $(U \cup V)^{<\varepsilon} = U^{<\varepsilon} \cup V^{<\varepsilon} = X \supseteq L$ for every $\varepsilon > 0$. \square

Almost clopen formulas are characterized by a kind of relative quantification.

Proposition D.1.5. For any partial type $\Sigma(\bar{x})$, open formulas $U(\bar{v}, \bar{x})$ and $V(\bar{v}, \bar{x})$ form an almost clopen partition of $(S_{\bar{v}\bar{x}}(\Sigma), d_{/\bar{x}})$ if and only if for every open formula $W(\bar{v}, \bar{w}, \bar{x})$, $(\exists \bar{v}U(\bar{v}, \bar{x}) \wedge W(\bar{v}, \bar{w}, \bar{x})) \vee (\exists \bar{v}V(\bar{v}, \bar{x}) \wedge W(\bar{v}, \bar{w}, \bar{x}))$ is logically equivalent to $\exists \bar{v}W(\bar{v}, \bar{w}, \bar{x})$.

Proof. This follows from representing $W(\bar{v}, \bar{w}, \bar{x})$ as $\varphi(\bar{v}, \bar{w}, \bar{x}) > 0$, and considering the uniform continuity of φ . \square

Proposition D.1.6. *A compact topometric space (X, τ, d) has a basis of o.a.c. sets if and only if for every pair of disjoint closed sets $F, G \subseteq X$ and every $\varepsilon > 0$, there exists open $U \supseteq F$ and $V \supseteq G$ with $\text{cl}U \cap \text{cl}V = \emptyset$ such that $X = U^{<\varepsilon} \cup V^{<\varepsilon}$.*

Proof. The \Rightarrow direction is immediate.

For the \Leftarrow direction, we run an iterative construction with $F_0 = F$ and $G_0 = G$.

At stage k , given F_k and G_k , disjoint open sets, we can find $U_k \supseteq F_k$ and $V_k \supseteq G_k$ such that $\text{cl}U_k \cap \text{cl}V_k = \emptyset$ and $X = U_k^{<2^{-k}} \cup V_k^{<2^{-k}}$. It is not hard to see that $\bigcup_{k < \omega} U_k$ and $\bigcup_{k < \omega} V_k$ form an almost clopen partition of X separating F and G . \square

Recall that a topological space is *zero-dimensional* if every open cover has a pairwise disjoint refinement (which a fortiori consists of clopen sets). For compact Hausdorff spaces this is equivalent to having a basis of clopen sets, but it is not equivalent in general.

Proposition D.1.7. *A compact separable topometric space (X, τ, d) has a basis of o.a.c. sets if and only if there is a topologically zero-dimensional set $A \subseteq X$ such that $X \setminus A$ is metrically meager (i.e. a subset of a countable union of metrically nowhere metrically dense sets).*

Proof. For the \Rightarrow direction, let \mathcal{B} be a countable basis of o.a.c. sets for X closed under unions, intersections, and exteriors. Let $A = \bigcap_{U \in \mathcal{B}} U \cup \text{ext} U$. A has a basis of (relatively) clopen sets by construction, and so is zero-dimensional. A is metrically comeager by the Baire category theorem.

For the \Leftarrow direction, by [EKNK78, Prop. 1.2.13], X has a countable base \mathcal{B} such that $A \cap \text{ext } U = \emptyset$ for every $U \in \mathcal{B}$. This is a basis of o.a.c. sets. \square

In general, having a basis of o.a.c. sets does not imply dictionaricness. An easy example is $S_1(\text{PS})$ (Definition 2.3.34). We will see in Corollary D.2.7 that the converse does hold, however, despite the fact that not all definable sets are o.a.c.

D.2 Dictionaric Type Spaces and Generic Separators

Definition D.2.1. Given a topological space X and disjoint closed sets F and G , a *separator between F and G* is an ordered pair (U, V) such that U and V are disjoint open sets, $U \supseteq F$, and $V \supseteq G$.¹ A separator (U, V) is *strict* if $\text{cl } U \cap \text{cl } V = \emptyset$. An *ordered separator in X* is any ordered pair (U, V) of disjoint open sets. The set of all ordered separators in X is written $OS(X)$.

If X is a normal topological space (e.g. a compact Hausdorff space), $P \subseteq OS(X)$, F and G are disjoint closed subsets of X , the *separator game with payoff set P* , written $SG(X, P)$, is a two player game specified by: On the zeroth turn player I plays a strict separator (U_0, V_0) between \emptyset and \emptyset (i.e. a pair of open sets with disjoint closures). After turn i , the other player places a strict separator (U_{i+1}, V_{i+1}) between $\text{cl } U_i$ and $\text{cl } V_i$. After ω turns, player II wins if and only if $(\bigcup_{i < \omega} U_i, \bigcup_{i < \omega} V_i) \in P$.

A set $P \subseteq OS(X)$ is *generic* if player II has a winning strategy for $SG(X, P)$. \triangleleft

Normality is only required so that both players always have a move available to them. We could have also defined $SG(X, P)$ so that the plays are pairs of disjoint closed sets and it wouldn't change very much for normal spaces. We would have to show that

¹Strictly speaking the separator is $X \setminus (U \cup V)$.

the resulting pair of sets being open is generic, however. Note that by symmetry we immediately get that if P is generic, then $P^{-1} = \{(U, V) : (V, U) \in P\}$ is also generic.

Proposition D.2.2. *If X is a second countable normal space (or, equivalently, a separable metrizable space), then*

$$\{(U, V) \in OS(X) : U = \text{int}(X \setminus V)\}$$

is generic.

Proof. Let $\{W_i\}_{i < \omega}$ be an enumeration of a countable basis for X . The winning strategy for player II is as follows: On turn $2i + 1$, given (U_{2i}, V_{2i}) , look at W_i and if $\text{cl } W_i \cap \text{cl } V_{2i} = \emptyset$, then play a strict separator between $\text{cl}(U_{2i} \cup W_i)$ and $\text{cl } V_{2i}$, otherwise just play a strict separator between $\text{cl } U_{2i}$ and $\text{cl } V_{2i}$.

Let (U_{2i+1}, V_{2i+1}) be player II's move on turn $2i + 1$, and let $U_\omega = \bigcup_{i < \omega} U_i$ and $V_\omega = \bigcup_{i < \omega} V_i$. We need to verify that $U_\omega = \text{int}(X \setminus V_\omega)$ and $V_\omega = \text{int}(X \setminus U_\omega)$. Assume that $W_i \subseteq \text{int}(X \setminus V_\omega)$. Since $\text{cl } V_{2i} \subseteq V_\omega$, $\text{cl } W_i \subseteq X \setminus V_\omega$ is disjoint from $\text{cl } V_{2i}$, so by the move player II made on turn $2i + 1$, we have that $W_i \subseteq U_{2i+1} \subseteq U_\omega$. Therefore $U_\omega = \text{int}(X \setminus V_\omega)$, as required. \square

Lemma D.2.3. *If $P \subseteq OS(X)$ is generic, then for any disjoint closed sets $F, G \subseteq X$, there exists open $U \supseteq F$ and $V \supseteq G$ such that $(U, V) \in P$.*

Proof. Let S be a winning strategy for player II in $SG(X, P)$. Since X is normal, we can find a strict separator (U_0, V_0) between F and G . Let this be our first play and then after that play arbitrarily. Then $U = \bigcup_{i < \omega} U_i$ and $V = \bigcup_{i < \omega} V_i$ are the required sets, because S is a winning strategy. \square

As the name might suggest, we have an analog of the Baire category theorem for generic properties of separators.²

Proposition D.2.4. *For any normal topological space X , if $\{P_i\}_{i<\omega}$ is a sequence of generic subsets of $OS(X)$, then $\bigcap_{i<\omega} P_i$ is also generic.*

Proof. For each i , let S_i be a winning strategy for player II for $SG(X, P_i)$. Let $f : \omega \rightarrow \omega$ be a function such that for each $i < \omega$, $f^{-1}(i)$ is infinite. We get a winning strategy for $SG(X, \bigcap_{i<\omega} P_i)$ by playing each S_i against everything else. On turn $2i+1$, if $k_0, \dots, k_{\ell-1}$ are the natural numbers less than i such that $f(k_j) = f(i)$ (i.e. the previous turns on which we used the strategy $S_{f(i)}$), present $S_{f(i)}$ with the following play.

$$\begin{array}{l} \text{I} \quad (U_{2k_0}, V_{2k_0}) \quad \dots \quad (U_{2k_{\ell-1}}, V_{2k_{\ell-1}}) \\ \text{II} \quad \quad \quad (U_{2k_0+1}, V_{2k_0+1}) \quad \dots \quad (U_{2k_{\ell-1}+1}, V_{2k_{\ell-1}+1}) \end{array}$$

Then play $S_{f(i)}$'s response. Since each S_i is a winning strategy, we will have

$$\left(\bigcup_{i<\omega} U_i, \bigcup_{i<\omega} V_i \right) \in P_i$$

for every $i < \omega$. □

Corollary D.2.5. *If X is a second countable normal topological space, then*

$$\{(U, V) \in OS(X) : U, V \text{ regular, } U = \text{ext } V, V = \text{ext } U\}$$

is generic, where $\text{ext } U = \text{int}(X \setminus A)$.

²In fact, Proposition D.2.4 is actually just a proof that $OS(X)$ is a Baire space with the topology generated by sets of the form $W_{F,G} = \{(U, V) \in OS(X) : F \subseteq U, G \subseteq V\}$ for closed F and G .

Proof. By Proposition D.2.2 and D.2.4, we have that the set of ordered separators (U, V) such than $U = \text{ext } V$ and $V = \text{ext } U$ is generic. The exterior of an open set is always regular, since it is the interior of a closed set. \square

The connection of the concept of generic separators to definable sets in type spaces is summarized in the following proposition.

Proposition D.2.6. *For any compact topometric space (X, τ, d) , X is dictionaric if and only if the set $\mathcal{D} = \{(U, V) \in OS(X) : X \setminus V \text{ is definable}\}$ is generic.*

Proof. (\Rightarrow) Assume that X is dictionaric. It follows that for every $\varepsilon > 0$, X has a basis of open sets U satisfying $\text{cl}U \subseteq \text{int } U^{<\varepsilon}$ (for any $p \in D \subseteq U$, find positive $\delta < \varepsilon$ small enough that $D^{\leq\delta} \subseteq U$, then find positive $\gamma < \delta$ small enough that $D^{\leq\gamma} \subseteq \text{int } D^{<\delta}$, then $\text{int } D^{<\gamma}$ is the required set). Note that any finite union of open sets satisfying this condition also satisfies this condition. The winning strategy for player II is as follows: On turn $2i + 1$, given (U_{2i}, V_{2i}) , cover $\text{cl}U_{2i}$ with open sets W satisfying $\text{cl}W \subseteq \text{int } W^{<2^{-i}}$ and $\text{cl}W \cap \text{cl}V_{2i} = \emptyset$. By compactness this has a finite subcover, let U_{2i+1} be its union. By the previous comment we have that $\text{cl}U_{2i+1} \subseteq \text{int } U_{2i+1}^{<2^{-i}}$ and we also have $\text{cl}U_{2i+1} \cap \text{cl}V_{2i} = \emptyset$. Find $V_{2i+1} \supseteq \text{cl}V_{2i} \cup (X \setminus \text{int } U_{2i+1}^{<2^{-i}})$ with $\text{cl}U_{2i+1} \cap \text{cl}V_{2i+1} = \emptyset$. Our move is (U_{2i+1}, V_{2i+1}) .

Now to see that this is a winning strategy, consider $U_\omega = \bigcup_{i < \omega} U_i$ and $V_\omega = \bigcup_{i < \omega} V_i$. By construction, for any $i < \omega$, we have that $U_\omega^{<2^{-i}} \supseteq U_{2i+1}^{<2^{-i}}$ and $\text{int } U_{2i+1}^{<2^{-i}} \cup V_{2i+1} = X$, so we have $\text{int } U_\omega^{<2^{-i}} \cup V_\omega = X$. Therefore $\text{int } U_\omega^{<2^{-i}} \supseteq X \setminus V_\omega$ and so $X \setminus V_\omega \subseteq \text{int}(X \setminus V_\omega)^{<2^{-i}}$ and $X \setminus V_\omega$ is a definable set, and so $(U_\omega, V_\omega) \in \mathcal{D}$.

(\Leftarrow) This follows immediately from Lemma D.2.3 applied to $\{p\}$ and $X \setminus U$ for U and open neighborhood of p . \square

Corollary D.2.7. *If (X, d, τ) is a compact dictionary topometric space, then for any pair of disjoint closed sets F and G there exists disjoint open $U \supseteq F$ and $V \supseteq G$ such that $X \setminus U$ and $X \setminus V$ are both definable. In particular, (X, d, τ) has a basis of o.a.c. sets.*

Proof. \mathcal{D}^{-1} is also generic. To see that we get a basis of o.a.c. sets, note that if U and V are disjoint open sets such that $X \setminus V$ and $X \setminus U$ are both definable, then $(X \setminus (U \cup V)) \subseteq (U \cup V)^{<\varepsilon}$ for every $\varepsilon > 0$, implying that $X \setminus (U \cup V)$ is metrically meager. Finally, $\partial U \subseteq X \setminus (U \cup V)$, so it is metrically meager as well. \square

Note that the converse—having an o.a.c. basis implies dictionaryness—does not hold, specifically the polarized square has an o.a.c. basis but is not dictionary. We can say something similar to Proposition D.2.6 about having a basis of o.a.c. sets, although we need an open metric.

Proposition D.2.8. *A compact topometric space (X, τ, d) with an open metric has an o.a.c. basis if and only if the set $\mathcal{AC} = \{(U, V) \in OS(X) : U \text{ is almost clopen}\}$ is generic.*

Proof. (\Rightarrow) Assume that X has an o.a.c. basis. The winning strategy for player II is as follows: On turn $2i + 1$, given (U_{2i}, V_{2i}) , cover $\text{cl} U_{2i}$ with o.a.c. sets A satisfying $\text{cl} A \cap \text{cl} V_{2i} = \emptyset$. By compactness this has a finite subcover. Let W be its union. We have that W is almost clopen and satisfies $\text{cl} W \cap \text{cl} V_{2i} = \emptyset$. Let $O = \text{int}(X \setminus W)$. Since W is almost clopen we have that $X = W^{<2^{-i}} \cup O^{<2^{-i}}$ for every $\varepsilon > 0$. This implies that X is covered by $\bigcup\{A^{<2^{-i}} : A \text{ open, } \text{cl} A \subseteq W\} \cup \bigcup\{B^{<2^{-i}} : B \text{ open, } \text{cl} B \subseteq O\}$. By compactness, we can find finite lists A_0, \dots, A_n and B_0, \dots, B_m with $\text{cl} A_k \subseteq W$ and $\text{cl} B_k \subseteq O$ such that

- $\text{cl} U_{2i} \subseteq \bigcup_{k \leq n} A_k$,
- $\text{cl} V_{2i} \subseteq \bigcup_{k \leq m} B_k$, and
- $X = \bigcup_{k \leq n} A_k^{<2^{-i}} \cup \bigcup_{k \leq m} B_k^{<2^{-i}}$.

Let $U_{2i+1} = \bigcup_{k \leq n} A_k$ and $V_{2i+1} = \bigcup_{k \leq m} B_k$. By construction, $\text{cl} U_{2i+1} \cap \text{cl} V_{2i+1} = \emptyset$, $\text{cl} U_{2i} \subseteq U_{2i+1}$, and $\text{cl} V_{2i} \subseteq V_{2i+1}$. So play (U_{2i+1}, V_{2i+1}) .

Let $U_\omega = \bigcup_{i < \omega} U_i$ and $V_\omega = \bigcup_{i < \omega} V_i$. By construction, we have that $U_\omega^{<2^{-i}} \cup V_\omega^{<2^{-i}} = X$ for every $i < \omega$, therefore U_ω is almost clopen and $(U_\omega, V_\omega) \in \mathcal{AC}$.

(\Leftarrow) This follows immediately from Lemma D.2.3. □

Proposition D.2.9. *Let X be a normal topological space, and let $Y \subseteq X$ be a closed subspace. Suppose $P \subseteq OS(Y)$ is generic, then*

$$P^\dagger = \{(U, V) \in OS(X) : (U \cap Y, V \cap Y) \in P\}$$

is generic as well.

Proof. Let S be a winning strategy for $SG(Y, P)$ for player II. The winning strategy for player II for $SG(X, P^\dagger)$ is as follows: Let (U_0, V_0) be player I's first move. Let $W_0 = U_0 \cap Y$ and $O_0 = V_0 \cap Y$.

On turn $2i+1$, given (U_{2i}, V_{2i}) played by player I, let $W_{2i} = U_{2i} \cap Y$ and $O_{2i} = V_{2i} \cap Y$ and present S with the following play.

$$\begin{array}{llll} \text{I} & (W_0, O_0) & \dots & (W_{2i}, O_{2i}) \\ \text{II} & & (W_1, O_1) & \dots \end{array}$$

Let (W_{2i+1}, O_{2i+1}) be its response. Since Y is a closed subspace of a normal space it is normal, so we can find a continuous function $f : Y \rightarrow [0, 1]$ such that $\text{cl} W_{2i+1} \subseteq f^{-1}(0)$

and $\text{cl } O_{2i+1} \subseteq f^{-1}(1)$. Now if we let $g : Y \cup \text{cl } U_{2i} \cup V_{2i} \rightarrow [0, 1]$ be defined by $g(x) = f(x)$ if $x \in Y$, $g(x) = 0$ if $x \in \text{cl } U_{2i}$, and $g(x) = 1$ if $x \in \text{cl } V_{2i}$, then this is a well defined function that is continuous. Since X is normal, by the Tietze extension theorem we can find a continuous function $h : X \rightarrow [0, 1]$ extending g to all of X . Now finally we actually play (U_{2i+1}, V_{2i+1}) with $U_{2i+1} = \{x \in X : h(x) < \frac{1}{3}\}$ and $V_{2i+1} = \{x \in X : h(x) > \frac{2}{3}\}$. Note that these do indeed have disjoint closures and that $W_{2i+1} \subseteq U_{2i+1}$ and $O_{2i+1} \subseteq V_{2i+1}$.

Finally, since S is a winning strategy for $SG(Y, P)$, we must have that $(W_\omega, O_\omega) = (\bigcup_{i < \omega} W_i, \bigcup_{i < \omega} O_i) \in P$, so since $W_\omega = Y \cap \bigcup_{i < \omega} U_i$ and $O_\omega = Y \cap \bigcup_{i < \omega} V_i$, we have that $(\bigcup_{i < \omega} U_i, \bigcup_{i < \omega} V_i) \in P^\dagger$. \square

Corollary D.2.10. *If (X, τ, d) is a compact dictionary topometric space with open metric, $\{D_i\}_{i < \omega}$ is a countable collection of definable subsets of X , and $F, U \subseteq X$ are closed and open respectively, with $F \subseteq U$, then there exists a definable set $E \subseteq X$ such that $F \subseteq E \subseteq U$ and $D_i \cap E$ is definable for every $i < \omega$.*

Proof. By Proposition 2.4.3, D_i is dictionary for every $i < \omega$, so the set

$$\{(U, V) \in OS(D_i) : D_i \setminus V \text{ is definable}\}$$

is generic for each $i < \omega$ by Proposition D.2.6. Therefore by Propositions D.2.9 and D.2.4 we have that

$$\{(U, V) \in OS(X) : (\forall i < \omega) D_i \setminus V \text{ is definable}\}$$

is generic, so the required result follows by Lemma D.2.3. \square

Corollary D.2.11. *If (X, τ, d) is a compact dictionary topometric space with open metric and (X, τ) has a basis \mathcal{B} of open sets of cardinality $\leq \aleph_1$, then there is a collection \mathcal{D} of definable subsets of X that is closed under finite unions and intersections, which satisfies $|\mathcal{D}| \leq |\mathcal{B}| + \aleph_0$, and which contains a basis of neighborhoods for X .*

Proof. If \mathcal{B} is finite then X is finite and we can just let \mathcal{D} be the collection of all subsets of X .

Assume that \mathcal{B} is infinite, and let $\kappa = |\mathcal{B}|$. Let $\{(U_i, V_i)\}_{i < \kappa}$ be an enumeration of all pairs $U_i, V_i \in \mathcal{B}$ such that $\text{cl} U_i \subseteq V_i$.

Let $\mathcal{D}_0 = \{\emptyset, X\}$. Clearly \mathcal{D}_0 is an at most countable collection of definable subsets of X closed under finite unions and intersections. For each limit ordinal $\beta < \kappa$, let $\mathcal{D}_\beta = \bigcup_{\alpha < \beta} \mathcal{D}_\alpha$. If each \mathcal{D}_α is an at most countable collection of definable sets closed under finite unions and intersections, then \mathcal{D}_β is as well. For each ordinal $\alpha < \kappa$, given \mathcal{D}_α , an at most countable collection of definable sets closed under finite unions and intersections, use Corollary D.2.10 to find a definable set $E \subseteq X$ such that $\text{cl} U_\alpha \subseteq E \subseteq V_\alpha$ and such that $D \cap E$ is definable for every $D \in \mathcal{D}_\alpha$. Let $\mathcal{D}_{\alpha+1}$ be the collection of all finite positive Boolean combinations of elements of $\mathcal{D}_\alpha \cup \{E\}$. By construction and by passing to conjunctive normal form, every element of this set is a finite union of sets of the form $D \cap E$ with $D \in \mathcal{D}_\alpha$ and is therefore a definable set. Furthermore, $\mathcal{D}_{\alpha+1}$ is clearly at most countable, so we have that $\mathcal{D}_{\alpha+1}$ is an at most countable collection of definable sets closed under finite unions and intersections.

Now finally let $\mathcal{D} = \bigcup_{\alpha < \kappa} \mathcal{D}_\alpha$. Clearly we have that \mathcal{D} is a collection of definable sets closed under finite unions and intersections. To see that \mathcal{D} contains a basis of neighborhoods, note that for any $x \in X$ and $U \ni x$ with $V \in \mathcal{B}$, since X is a compact Hausdorff space we can find $U' \in \mathcal{B}$ with $x \in U'$ and $\text{cl} U' \subseteq V$. At some point in the

construction we added a definable set D to \mathcal{D} with $U \subseteq D \subseteq V$. Therefore D is a neighborhood of p contained in V . Since we can do this for any p and V , we have that \mathcal{D} contains a basis of neighborhoods. \square

It is well known that Martin's axiom implies that any union of fewer than continuum many meager sets in a metric space is still meager. This suggests the question of whether or not something analogous can be done in this context.

Question D.2.12. *Can Corollary D.2.11 be extended to spaces with weight 2^{\aleph_0} under Martin's axiom? Without Martin's axiom? Can it be extended further than 2^{\aleph_0} ?*

Definition D.2.13. For any sets X, Y , subset $A \subseteq X$, and function $f : X \rightarrow Y$, the *co-image* of A under f , written $f_{co}(A)$, is the set $Y \setminus f(X \setminus A)$. \triangleleft

Lemma D.2.14. *Let X and Y be compact Hausdorff spaces, and let $f : X \rightarrow Y$ be a continuous map.*

- (i) *For any set $A \subseteq X$, $\text{cl}_Y f(A) = f(\text{cl}_X A)$.*
- (ii) *If f is an open map, then $\text{cl}_Y f_{co}(A) \subseteq f_{co}(\text{cl}_X A)$.*
- (iii) *If f is an open map and (U, V) is a strict separator in X , then*

$$(\text{int}_Y f(\text{cl}_X U), \text{int}_Y f_{co}(\text{cl}_X V))$$

is a strict separator in Y . Furthermore, $f(U) \subseteq \text{int}_Y f(\text{cl}_X U)$ and $f_{co}(V) \subseteq \text{int}_Y f_{co}(\text{cl}_X V)$.

- (iv) *If f is an open surjection, (U, V) is a strict separator in X , and (W, O) is a strict separator between $f(\text{cl}_X U)$ and $f_{co}(\text{cl}_X V)$ in Y , then there exists open $U', V' \subseteq X$*

such that (U', V') is a strict separator between $\text{cl}_X U$ and $\text{cl}_X V$ and such that $(f(U'), f_{co}(V'))$ is a strict separator between $\text{cl}_Y W$ and $\text{cl}_Y O$.

Proof. (i) For any $y \in \text{cl}_Y f(A)$, find a net $\{y_i\}_{i \in I}$ of elements of $f(A)$ converging to y . For each y_i , find $x_i \in A$ such that $f(x_i) = y_i$. By compactness, the net $\{x_i\}_{i \in I}$ has a convergent sub-net, converging to $x \in \text{cl}_X A$. By continuity, $f(x) = y$. Therefore, for any $\text{cl}_Y f(A) \subseteq f(\text{cl}_X A)$. Since $f(\text{cl}_X A)$ is a closed set we have $\text{cl}_Y f(A) \subseteq f(\text{cl}_X A)$, so the sets are equal.

(ii) $X \setminus \text{cl}_X A$ is open, so $f(X \setminus \text{cl}_X A)$ is open, and then clearly $\text{cl}_Y f_{co}(A) \subseteq Y \setminus f(X \setminus \text{cl}_X A) = f_{co}(\text{cl}_X A)$.

(iii) Suppose $y \in f(\text{cl}_X U) \cap f_{co}(\text{cl}_X V)$, then this implies that for some $x \in \text{cl}_X U$, $f(x) = y$ and $f^{-1}(y) \subseteq \text{cl}_X V$, but this implies that $x \in \text{cl}_X V$ as well, contradicting that (U, V) is a strict separator.

For the ‘furthermore,’ statement, since image of co-image are both monotonic, we have $f(U) \subseteq f(\text{cl}_X U)$ and $f_{co}(U) \subseteq f_{co}(\text{cl}_X V)$ and the statement follows immediately.

(iv) Since $\text{cl}_Y W \cap (f_{co}(\text{cl}_X V) \cup \text{cl}_Y O) = \emptyset$, for each $y \in \text{cl}_Y W$, there exists an $x_y \in X$ such that $f(x_y) = y$ (since f is a surjection) and $x_y \notin \text{cl}_X V \cup f^{-1}(\text{cl}_Y O)$. For each such y , let $U_y \subseteq X$ be an open neighborhood of x_y such that $\text{cl}_X U_y \cap (\text{cl}_X V \cup f^{-1}(\text{cl}_Y O)) = \emptyset$. Since f is open, for each $y \in \text{cl}_Y W$, $f(U_y)$ is an open neighborhood of y . By compactness there is a finite set $Y_0 \subseteq \text{cl}_Y W$ such that $\bigcup_{y \in Y_0} f(U_y) \supseteq \text{cl}_Y W$.

Note that since $f(\text{cl}_X U) \cap \text{cl}_Y O = \emptyset$, we also have $\text{cl}_X U \cap f^{-1}(\text{cl}_Y O) = \emptyset$. Because of this we can find U' , an open neighborhood of $\text{cl}_X U \cup \bigcup_{y \in Y_0} \text{cl}_X U_y$ such that $\text{cl}_X U' \cap (\text{cl}_X V \cup f^{-1}(\text{cl}_Y O)) = \emptyset$. Then find V' , an open neighborhood of $\text{cl}_X V \cup f^{-1}(\text{cl}_Y O)$ such that $\text{cl}_X U' \cap \text{cl}_X V' = \emptyset$.

By construction (U', V') is a strict separator between $\text{cl}_X U$ and $\text{cl}_X V$. We also

have $\text{cl}_Y W \subseteq f(U')$ and $\text{cl}_Y O \subseteq f_{co}(V')$, as required. The full result follows from part (iii). \square

Proposition D.2.15. *If X and Y are compact Hausdorff spaces and $f : X \rightarrow Y$ is an open continuous surjection, then for any generic set $P \subseteq OS(Y)$, the set*

$$P^\clubsuit = \{(U, V) \in OS(X) : (f(U), f_{co}(V)) \in P\}$$

is generic.

Proof. Let S be a winning strategy for $SG(Y, P)$ for player II. The winning strategy for player II for $SG(X, P^\clubsuit)$ is as follows: Let (U_0, V_0) be player I's first move. Let $W_0 = f(U_0)$, and let $O_0 = f_{co}(V_0)$. Note that these are both open sets.

On turn $2i + 1$, given (U_{2i}, V_{2i}) played by player I, let (W_{2i}, O_{2i}) be a strict separator between $f(\text{cl}_X U_{2i})$ and $f_{co}(\text{cl}_X V_{2i})$ (which exists by Lemma D.2.14), and present S with the following play.

$$\begin{array}{l} \text{I} \quad (W_0, O_0) \qquad \dots \quad (W_{2i}, O_{2i}) \\ \text{II} \qquad \qquad (W_1, O_1) \quad \dots \end{array}$$

Note that this is a legal play. And let (W_{2i+1}, O_{2i+1}) be its response. Use Lemma D.2.14 part (iv) to find (U_{2i+1}, V_{2i+1}) , a strict separator between $\text{cl}_X U_{2i}$ and $\text{cl}_X V_{2i}$ such that $(f(U_{2i+1}), f_{co}(V_{2i+1}))$ is a strict separator between $\text{cl}_Y W_{2i+1}$ and $\text{cl}_Y O_{2i+1}$.

Let $U_\omega = \bigcup_{i < \omega} U_i$ and $V_\omega = \bigcup_{i < \omega} V_i$.

Claim: $f(U_\omega) = \bigcup_{i < \omega} W_i$ and $f_{co}(V_\omega) = \bigcup_{i < \omega} O_i$.

Proof of claim. Let y be an element of $f(U_\omega)$. This implies that $y \in f(U_{2i})$ for some $i < \omega$. Therefore $y \in W_{2i+1} \subseteq \bigcup_{k < \omega} W_k$.

Let y be an element of $\bigcup_{k < \omega} W_k$, then $y \in W_{2i+1}$ for some $i < \omega$, so $y \in f(U_{2i+1}) \subseteq f(U_\omega)$.

Therefore $f(U_\omega) = \bigcup_{i < \omega} W_i$.

Now let y be an element of $f_{co}(V_\omega)$. This implies that $f^{-1}(y) \subseteq V_\omega$. By compactness, there is an $i < \omega$ such that $f^{-1}(y) \subseteq V_{2i}$. Therefore $y \in f_{co}(V_{2i})$ and we have $y \in O_{2i} \subseteq \bigcup_{k < \omega} O_k$ as well.

Let y be an element of $\bigcup_{k < \omega} O_k$, then $y \in O_{2i}$ for some $i < \omega$ and $f^{-1}(y) \subseteq f^{-1}(\text{cl}_Y O_{2i}) \subseteq V_{2i+1}$, so $y \in f_{co}(V_{2i+1}) \subseteq f_{co}(V_\omega)$.

Therefore $f_{co}(V_\omega) = \bigcup_{k < \omega} O_k$. □_{claim}

Since S is a winning strategy we have that $(\bigcup_{i < \omega} W_i, \bigcup_{i < \omega} O_i) \in P$, therefore $(U_\omega, V_\omega) \in P^\clubsuit$, and we have that P^\clubsuit is a generic. □

Proposition D.2.16. *Let (X, τ, d) be a dictionary compact topometric space with an open metric, and let $f : X \rightarrow X$ be an involutive topometric automorphism (i.e. $f(f(x)) = x$). The set*

$$\{(U, V) \in OS(X) : \text{cl}U \cap \text{cl}f(U) \text{ is definable}\}$$

is generic.

Proof. Let P be the set in the statement of the proposition. Let (U_0, V_0) be player I's first move. If there is a legal move (U_1, V_1) such that $U_1 \cap f(U_1)$ is non-empty, make it, otherwise play arbitrarily for the entire game. Note that if no such move exists, we will necessarily have that $\text{cl}U \cap \text{cl}f(U)$ is empty, and therefore definable.

If $U_1 \cap f(U_1)$ is non-empty, on turn $2k + 1$ for $k > 0$, given (U_{2k}, V_{2k}) , we have that $\text{cl}U_{2k} \cap \text{cl}f(U_{2k}) \cap \text{cl}V_{2k} \cup \text{cl}f(V_{2k}) = \emptyset$. Find a definable set D_{2k} such that

$\text{cl } U_{2k} \cap \text{cl } f(U_{2k}) \subseteq D_{2k}$ and $D_{2k} \cap (\text{cl } V_{2k} \cup \text{cl } f(V_{2k})) = \emptyset$. Let $E_{2k} = D_{2k} \cup f(D_{2k})$. By symmetry, we have $\text{cl } U_{2k} \cap \text{cl } f(U_{2k}) \subseteq E_{2k}$ and $E_{2k} \cap (\text{cl } V_{2k} \cup \text{cl } f(V_{2k})) = \emptyset$. For each $x \in \text{cl } V \cup \text{cl } f(V) \cup (X \setminus E_{2k}^{<2^{-k}})$, since $x \notin \text{cl } U \cap \text{cl } f(U)$, either $x \notin \text{cl } U$ or $x \notin \text{cl } f(U)$. If $x \notin \text{cl } U$, find an open neighborhood $W_x \ni x$ such that $\text{cl } W_x \cap (E_{2k} \cup \text{cl } U) = \emptyset$, otherwise find an open neighborhood $W_x \ni x$ such that $\text{cl } W_x \cap (E_{2k} \cup \text{cl } f(U)) = \emptyset$. By compactness, there is a finite $X_0 \subseteq \text{cl } V \cup \text{cl } f(V) \cup (X \setminus E_{2k}^{<2^{-k}})$ such that $\bigcup_{x \in X_0} W_x \supseteq \text{cl } V \cup \text{cl } f(V) \cup (X \setminus E_{2k}^{<2^{-k}})$. For each $x \in X_0$, let $O_x = W_x$ if $\text{cl } W_x \cap \text{cl } U = \emptyset$, and let $O_x = f(W_x)$ otherwise. Let $A_{2k} = \bigcup_{x \in X_0} O_x$. By construction we have that $E_{2k} \cap (A_{2k} \cup f(A_{2k})) = \emptyset$ and $E_{2k}^{<2^{-k}} \cup A_{2k} \cup f(A_{2k}) = X$. By compactness, we can find $V_{2k+1} \supseteq \text{cl } V_{2k} \cup \bigcup_{x \in X_0} \text{cl } O_x$ such that $E_{2k} \cap (\text{cl } V_{2k+1} \cup \text{cl } f(V_{2k+1})) = \emptyset$. Finally find U_{2k+1} such that $U_{2k+1} \supseteq \text{cl } U_{2k} \cup E_{2k}$ and $\text{cl } U_{2k+1} \cap \text{cl } V_{2k+1} = \emptyset$.

Note that by construction, for each positive $k < \omega$ we have that

$$(\text{cl } U_{2k+1} \cap \text{cl } f(U_{2k+1}))^{<2^{-k}} \cup \text{cl } V_{2k+1} \cup \text{cl } f(V_{2k+1}) = X.$$

Let $U_\omega = \bigcup_{k < \omega} U_k$ and $V_\omega = \bigcup_{k < \omega} V_k$. Fix $\varepsilon > 0$, and find k such that $2^{-k} < \varepsilon$. For any $x \in \text{cl } U_\omega \cap \text{cl } f(U_\omega)$ we have that $x \notin \text{cl } V_{2k+1} \cup \text{cl } f(V_{2k+1})$ and so

$$x \in (\text{cl } U_{2k+1} \cap \text{cl } f(U_{2k+1}))^{<2^{-k}} \subseteq (U_\omega \cap f(U_\omega))^{<2^{-k}} \subseteq \text{int}(\text{cl } U_\omega \cap \text{cl } f(U_\omega))^{<\varepsilon},$$

and therefore $\text{cl } U_\omega \cap \text{cl } f(U_\omega) \subseteq \text{int}(\text{cl } U_\omega \cap \text{cl } f(U_\omega))^{<\varepsilon}$ for every $\varepsilon > 0$, and $\text{cl } U_\omega \cap \text{cl } f(U_\omega)$ is definable. \square

More generally what you get is this:

Proposition D.2.17. *Let (X, τ, d) be a dictionary compact topometric space with an*

open metric, and let $f : X \rightarrow X$ be a topometric automorphism such that for some finite n , $f^n(x) = x$. The set

$$\left\{ (U, V) \in OS(X) : \bigcap_{\ell < n} \text{cl } f^\ell(U) \text{ is definable} \right\}$$

is generic.

Definition D.2.18. Given a topological space X and an open cover \mathcal{U} of X , the *ply* of \mathcal{U} is $\inf_{x \in X} |\{U \in \mathcal{U} : x \in U\}|$. Another open cover \mathcal{V} is a *refinement* of \mathcal{U} if for every $V \in \mathcal{V}$ there is $U \in \mathcal{U}$ such that $V \subseteq U$. The *Lebesgue covering dimension* of X , written $\dim(X)$, is the smallest finite n such that every open cover \mathcal{U} of X has a refinement with ply less than or equal to $n + 1$. If no such n exists, $\dim(X) = \infty$.

The *large inductive dimension* of X , written $\text{Ind}(X)$, is defined inductively. $\text{Ind}(\emptyset) = -1$. If X is non-empty, then $\text{Ind}(X)$ is the smallest $n < \omega$ such that for any pair of disjoint closed sets $F, G \subseteq X$, there exists an open set $U \supseteq F$ with $\text{cl } U \cap G = \emptyset$ and $\text{Ind}(\partial U) < n$. If no such n exists, $\text{Ind}(X) = \infty$. \triangleleft

Fact D.2.19 (Urysohn's theorem). *If X is a second countable normal space, then $\dim(X) = \text{Ind}(X)$.*

Note that by Urysohn's metrization theorem, a topological space is second countable and normal if and only if it is separable and metrizable.

Proposition D.2.20. *If X is a second countable compact Hausdorff space, then for any $n < \omega$, $\dim(X) \leq n$ if and only if*

$$\{(U, V) \in OS(X) : \dim(X \setminus (U \cup V)) < n\}$$

is generic.

Moreover, for $n = 0$ we can drop the requirement that X be second countable.

Proof. Assume that $\dim(X) \leq n$. Since X is a compact Hausdorff space, it is normal. Let d be a metric inducing the topology on X . We will describe a winning strategy for player II. On turn $2i + 1$, given (U_{2i}, V_{2i}) , since $\text{Ind}(X) = n$, we can find open $A \supseteq \text{cl } U_{2i}$ such that $\text{cl } A \cap \text{cl } V_{2i} = \emptyset$ and such that $\text{Ind}(\partial A) < n$. By Fact D.2.19, we have that $\dim(\partial A) < n$. Find $\varepsilon > 0$ with $\varepsilon < 2^{-i-1}$ small enough that $(\partial A)^{\leq \varepsilon}$ is disjoint from $\text{cl } U_{2i}$ and $\text{cl } V_{2i}$. Cover ∂A with open ε -balls, find a refinement \mathcal{W}_i^* with $\text{ply} \leq n$. Find a finite sub-cover \mathcal{W}_i . Note that \mathcal{W}_i 's ply is still $\leq n$ and every $W \in \mathcal{W}_i$ has $\text{diam } W \leq 2^{-i}$. Find U_{2i+1} and V_{2i+1} with $\text{cl } U_{2i} \subseteq U_{2i+1}$, $\text{cl } V_{2i} \subseteq V_{2i+1}$, $\text{cl } U_{2i+1} \cap \text{cl } V_{2i+1} = \emptyset$, and $U_{2i+1} \cup V_{2i+1} \cup \bigcup \mathcal{W}_i = X$.

Now let $U_\omega = \bigcup_{i < \omega} U_i$ and $V_\omega = \bigcup_{i < \omega} V_i$. We want to show that $F = X \setminus (U_\omega \cup V_\omega)$ has $\dim(F) < n$. Note that for every $i < \omega$, \mathcal{W}_i is an open cover of F . Let \mathcal{O} be an open cover of F (specifically let the elements of \mathcal{O} be open subsets of X). Since (X, d) is a compact metric space, \mathcal{O} has a Lebesgue number $\varepsilon > 0$ (see Fact A.1.4). Find $i < \omega$ such that $2^{-i+1} < \varepsilon$. Then by the definition of Lebesgue number, we have that \mathcal{W}_i is a refinement of \mathcal{O} and that by construction \mathcal{W}_i has $\text{ply} \leq n$. Therefore every open cover of F has a refinement with $\text{ply} \leq n$ and we have that $\dim(F) < n$, as required.

The other direction follows from Lemma D.2.3.

If X is a zero-dimensional compact Hausdorff space, then player II can win on turn 1 by playing disjoint clopen sets A and B such that $\text{cl } U_0 \subseteq A$, $\text{cl } V_0 \subseteq B$, and $A \cup B = X$. The other direction again follows from Lemma D.2.3. \square

D.3 Definable Sets in Type Spaces Homeomorphic to $[0, 1]$

A detailed characterization of dictionary type spaces in general seems very difficult. Some additional traction can be gained by making assumptions about the topology of the relevant type space. In this section we will examine the simplest non-trivial case, that of type spaces homeomorphic to $[0, 1]$.

Definition D.3.1. Let d be a topometric on $[0, 1]$. A point $p \in [0, 1]$ is called a *left hand definable endpoint* if for some (or equivalently every) $q \in \mathbb{R}$ with $q > p$, the set $[p, q] \cap [0, 1]$ is locatable. p is called a *right hand definable endpoint* if for some $q \in \mathbb{R}$ with $q < p$, the set $(q, p] \cap [0, 1]$ is locatable. p is called a *two-sided definable endpoint* if it is both a left hand definable endpoint and a right hand definable endpoint. \triangleleft

Obviously a d -atomic point is a two sided definable endpoint. The converse does not hold, however (Counterexample C.1.11). Also, we obviously have that $[a, b] \subseteq [0, 1]$ is definable if and only if either $a = 0$ or a is a left hand definable endpoint and either $b = 1$ or b is a right hand definable endpoint. At the moment the existence of left hand definable endpoints that are not right hand definable (or vice versa) is open, although it would be fairly surprising if such a thing did not exist.

Question D.3.2. *Does there exist a type space $S_1(T)$ homeomorphic to $[0, 1]$ with some $p \in S_1(T)$ that is left hand definable but not right hand definable?*

Thinking in these terms we get some easy observations.

Proposition D.3.3. *If d is a topometric for $[0, 1]$, and $D \subseteq [0, 1]$ is definable in $([0, 1], d)$, then $\sup D$ is a right hand definable endpoint and $\inf D$ is a left hand definable endpoint.*

Furthermore, the convex closure of D is definable.

Proof. First note that the convex closure of D is a closed set that is the union of a definable set and an open set, and is therefore definable. The other statements follow immediately from this. \square

Proposition D.3.4. *Let d be an open topometric on $[0, 1]$. The following are equivalent.*

- (i) $([0, 1], d)$ is dictionaric.
- (ii) For every $p \in [0, 1]$ and $(a, b) \ni p$ there is definable $[c, e] \subseteq (a, b)$ with $p \in [c, e]$.
- (iii) Left hand definable endpoints are dense in $[0, 1]$.
- (iv) Right hand definable endpoints are dense in $[0, 1]$.
- (v) Two-sided definable endpoints are comeager in $[0, 1]$.

Proof. (i) \Rightarrow (ii) and (ii) \Rightarrow (iii) \wedge (iv) follow immediately from Proposition D.3.3.

(iii) \vee (iv) \Rightarrow (v): To demonstrate comeagerness, we will perform this construction with an opponent choosing open sets.

Assume without loss that left hand definable endpoints are dense in $[0, 1]$. Let the opponent choose a non-empty open set $U_0 \subseteq [0, 1]$. Find $a_0 \in U_0$, a left hand definable endpoint not equal to 0 or 1.

For each $i < \omega$, given a_i , a left hand definable endpoint not equal to 0 or 1, and U_i , an open set with $a_i \in U_i$, let $V_i = [a_i, 1]^{<2^{-i}} \setminus [a_i, 1]$, and let $W_i = (V_i^{<2^{-i}} \setminus [0, a_i]) \cap U_i$.

By construction, $[a_i, 1]$ is a definable set, so V_i is an open set whose supremum is a_i . Furthermore, $V_i^{<2^{-i}} \setminus [0, a_i]$ must be non-empty, because for every $x \in V_i$, there is a $y \in [a_i, 1]$ with $d(x, y) < 2^{-i}$. By openness of the metric, $V_i^{<2^{-i}}$ is an open set, so $V_i^{<2^{-i}} \cap [a_i, 1]$ cannot be just $\{a_i\}$.

Claim: $\inf V_i^{<2^{-i}} \cap [a_i, 1] = a_i$.

Proof of claim. For every $\varepsilon > 0$ with $\varepsilon < 2^{-i}$, we have that $O_\varepsilon = [a_i, 1]^{<\varepsilon} \setminus [a_i, 1]$ is an open set whose supremum is a_i . By construction, $O_\varepsilon^{<\varepsilon}$ must have non-empty intersection with $[a_i, 1]$, so a fortiori $(\text{cl } O_\varepsilon)^{\leq\varepsilon}$ has non-empty intersection with $[a_i, 1]$. Since the metric refines the topology, we have that $[a_i, 1] \cap \bigcap_{0 < \varepsilon < 2^{-i}} (\text{cl } O_\varepsilon)^{\leq\varepsilon} = \{a_i\}$. Therefore, for every neighborhood $N \ni a_i$, for any sufficiently small $\varepsilon > 0$, there are elements of $[0, a_i)$ with distance less than ε to elements of $N \cap [a_i, 1]$, but this implies the same for 2^{-i} , so we have that for every neighborhood $N \ni a_i$, there are elements of $[0, a_i)$ with distance less than ε to elements of $N \cap [a_i, 1]$. Therefore, by the definition of V_i we have that $V_i^{<2^{-i}} \cap [a_i, 1] \cap N$ is non-empty and the claim follows. \square_{claim}

This implies that W_i is non-empty, since U_i is a neighborhood of a_i .

Now find a left hand definable endpoint $b_i \in W_i$ not equal to 0 or 1, and then find a non-empty open set Y_{i+1} such that $\text{cl } Y_{i+1} \subseteq W_i \cap ([b_i, 1]^{<2^{-i}} \setminus [b_i, 1])$ and such that $\sup Y_{i+1} < \inf Y_{i+1} + 2^{-i}$. Play this open set and allow the opponent to play some non-empty open set $U_{i+1} \subseteq Y_{i+1}$. Find a left hand definable endpoint $a_{i+1} \in U_{i+1}$ that is not equal to 0 or 1.

Now consider $\bigcap_{i < \omega} \text{cl}(U_i) = \bigcap_{i < \omega} \text{cl}(Y_i)$. By compactness this is non-empty, and by construction it is a singleton. Let b be its unique element. We want to argue that c is a right hand definable endpoint.

Fix $i < \omega$, and consider $[0, c]^{<2^{-i}}$. By construction $c \in U_{i+1} \subseteq V_i^{<2^{-i}} \subseteq [0, c]^{<2^{-i}}$,

so we have that $c \in \text{int}[0, c]^{<2^{-i}}$. This implies that $[0, c] \subseteq \text{int}[0, c]^{<2^{-i}}$. Since we can do this for arbitrarily large $i < \omega$, we have that $[0, c] \subseteq \text{int}[0, c]^{<\varepsilon}$ for every $\varepsilon > 0$, and therefore $[0, c]$ is definable and hence c is a right hand definable endpoint.

On the other hand, note that by construction, for each $i < \omega$, $[b_{i+1}, 1] \subseteq [b_i, 1]^{<2^{-i}}$. Therefore the sequence $\{[b_i, 1]\}_{i < \omega}$ is a Hausdorff metric Cauchy sequence. By construction, the sequence $\{b_i\}_{i < \omega}$ limits to c , so we must have that the limit of $\{[b_i, 1]\}_{i < \omega}$ is $[c, 1]$. Therefore we have that $[c, 1]$ is definable and hence c is a right hand definable endpoint, and hence c is a two-sided definable endpoint.

Since we can do this with an opponent choosing the open sets U_i , we have that the set of two-sided definable endpoints is comeager in $[0, 1]$. \square

What is unclear at the moment is whether or not a dictionary type space homeomorphic to $[0, 1]$ must actually contain d -atomic types.

Question D.3.5. *Does every dictionary type space homeomorphic to $[0, 1]$ contain a d -atomic type?*

Note that if this is true, then actually the set of d -atomic types will always be dense in such type spaces and therefore also comeager (since metric limits of d -atomic types are d -atomic).

Counterexample C.1.8 shows that this already fails in some topologically 1-dimensional type spaces.

List of Symbols

Metric Spaces

$A^{d<\varepsilon}, A^{<\varepsilon}$	Open ε -fattening of A , 11
$A^{d\leq\varepsilon}, A^{\leq\varepsilon}$	Closed ε -fattening of A , 11
$B_{<\varepsilon}^d(x), B_{<\varepsilon}(x)$	Open ball of radius ε with center x , 10
$B_{\leq\varepsilon}^d(x), B_{\leq\varepsilon}(x)$	Closed ball of radius ε with center x , 10
$\text{diam } X$	Diameter of X , 10
$d_{\text{inf}}(x, A)$	Point-set distance from x to A , 10
$X \sqcup_r Y$	Disjoint union of X and Y with separation r , 11
$\#^{\text{dc}} X$	Density character of X , 11
$\#_{\geq\varepsilon}^{\text{ent}} Y, \#_{>\varepsilon}^{\text{ent}} Y$	Metric entropy of Y , 11
$\#_{<\varepsilon}^{\text{ecov}, X} Y, \#_{\leq\varepsilon}^{\text{ecov}, X} Y$	External covering number of Y in X , 12
$\#_{<\varepsilon}^{\text{icov}} Y, \#_{\leq\varepsilon}^{\text{icov}} Y$	Internal covering number of Y , 12
$\#_{<\varepsilon}^{\text{epac}, X} Y, \#_{\leq\varepsilon}^{\text{epac}, X} Y$	External packing number of Y in X , 12
$\#_{<\varepsilon}^{\text{ipac}} Y, \#_{\leq\varepsilon}^{\text{ipac}} Y$	Internal packing number of Y , 12
$\#_{\leq\varepsilon}^{\text{par}} Y$	Partition number of Y , 12

Signatures, Terms, and Formulas (Single-Sorted)

$[x]_r^s$	$[r, s]$ -clamping function, 26
$a(s)$	Arity of predicate or function symbol s , 17
$(\forall x \in F)U$	Relative strong universal quantification, 88
BM	Banach-Mazur distortion system, 360
$[\varphi]$	Element of prob. alg. corresponding to φ in randomization, 320
$\mathcal{C}(\mathcal{L})$	Set of constant symbols in \mathcal{L} , 17
$\mathcal{CL}(V)$	Set of closed $\mathcal{L}(V)$ -formulas, 29
Δ	Distortion system, 337
d	Metric predicate symbol, 17
$\text{db}(\mathcal{L})$	Diameter bound of \mathcal{L} , 17
$\Delta(\delta)$	Distortion system corresponding to the family of topometrics δ , 346
$D(\Delta, C)$	Parameter distortion system of Δ over the constants C , 374
$d_{\equiv}(f, g)$	Logical distance between f and g , 144
$\text{DEF}_{\bar{v}}\varphi$	Distance predicate quantifier, 106
$\text{def}_{\bar{v}}\varphi$	Real valued distance predicate quantifier, 106
$d_{\equiv}(t, s)$	Logical distance between t and s , 38

$d_H(D(\bar{x}), E(\bar{x}))$	Hausdorff distance between the definable sets D and E , 110
$d_{H,\Sigma}(D(\bar{x}), E(\bar{x}))$	Logical Hausdorff distance between D and E , 110
$\bar{\Delta}$	Distortion system completion of Δ , 334
$d_\Sigma(f, g)$	Logical distance between f and g over Σ , 144
$d_\Sigma(t, s)$	Logical distance between t and s modulo Σ , 38
$\Theta_{<\varepsilon}^{\text{cov}\geq n} xF(x)$	Finite (external) covering number quantifier for closed formulas, 129
$\exists_{\leq\varepsilon}^{\text{cov}\geq n} xL(x)$	Finite (external) covering number quantifier for locatable formulas, 129
$\text{cov}_x^n D(x)$	Real valued finite (external) covering quantifier for definable sets, 129
$\Theta_{\geq\varepsilon}^{\text{ent}\geq n} xF(x)$	Finite metric entropy quantifier for closed formulas, 129
$\exists_{>\varepsilon}^{\text{ent}\geq n} xL(x)$	Finite metric entropy quantifier for locatable formulas, 129
eGHK	Elementary Gromov-Hausdorff-Kadets distortion system, 355
$\text{ent}_x^n D(x)$	Real valued finite metric entropy quantifier for definable sets, 129
ΘxF	Weak existential quantification, 56
$(\Theta x \in G)F$	Relative weak existential quantification, 88
$\mathcal{F}(\mathcal{L})$	Set of function symbols in \mathcal{L} , 17
fGHK	Finitary Gromov-Hausdorff-Kadets distortion system, 355
$f(\bar{x})\uparrow$	Partial function is not defined for input \bar{x} , 138

$\text{FUN}_y \varphi(\bar{x}, y)$	Domain of partial function quantifier, 137
$\text{fun}_y \varphi(\bar{x}, y)$	Real valued domain of partial function quantifier, 137
$\text{fv}(\varphi)$	Free variables of φ , 23
$f(\bar{x})\downarrow$	Partial function is defined for input \bar{x} , 138
$F(\bar{x}, f(\bar{x})\downarrow, \bar{z})$	$f(\bar{x})$ is well defined and satisfies F , 138
$f(\bar{x}) \supseteq g(\bar{x})$	Partial function f extends partial function g , 138
GH	Gromov-Hausdorff distortion system, 333
\mathcal{I}	Interpretation, 220
$\inf_{\bar{v} \in D(\cdot, \bar{x})} \varphi$	Relative infimal quantification, 108
$I(P)$	Codomain interval of predicate symbol P , 17
$I(\varphi)$	Codomain interval of formula φ , 23
\mathcal{L}	A metric signature, 17
Lip	Lipschitz distortion system, 333
\mathcal{L}^{Mor}	Morleyization of \mathcal{L} , 37
$\mathcal{L}(V), \mathcal{L}(\bar{v})$	Set of \mathcal{L} -formulas with free variables in V or \bar{v} , 25
\uparrow	Infix notation for max, 26
\downarrow	Infix notation for min, 26

$x \dot{-} y$	Monus operator, $\max\{x - y, 0\}$, 26
$\ \varphi\ _{\equiv}$	Logical norm of φ , 38
$\text{nalg}(p)$	Degree of non-algebraicity of p , 265
$\ \varphi\ _{\Sigma}$	Logical norm of φ modulo Σ , 38
ω_s	Modulus of uniform continuity of predicate or function symbol s , 17
$O\mathcal{L}(V)$	Set of open $\mathcal{L}(V)$ -formulas, 29
$\omega_{\varphi, \bar{x}}$	Syntactic modulus of uniform continuity of φ with regards to \bar{x} , 36
$\omega_{t, \bar{x}}$	Syntactic modulus of uniform continuity of t with regards to \bar{x} , 36
$\mathcal{P}(\mathcal{L})$	Set of predicate symbols in \mathcal{L} , 17
$\varphi[\frac{\bar{t}}{\bar{x}}]$	φ with simultaneous substitution of \bar{t} for \bar{x} , 24
$\varphi \equiv \psi$	φ and ψ are logically equivalent, 38
$\varphi \equiv_{\Sigma} \psi$	φ and ψ are logically equivalent modulo Σ , 38
$\varphi(V), \varphi(\bar{v})$	Formula φ with relevant variables emphasized, 24
$Qx\varphi$	Formula with variable quantifier, 24
\square	Variable for infix relation between reals, such as \leq or \neq , 29
$R\mathcal{L}(V)$	Set of restricted $\mathcal{L}(V)$ -formulas, 27
$\Sigma \models X$	Σ logically entails X , 30

$a \sqcap b, a \sqcup b$	Boolean operations on the prob. alg. sort of a randomization, 320
$a \sqsubseteq b$	a is contained in b in the probability algebra of a randomization, 320
$\sup_{\bar{v} \in D(\cdot, \bar{x})} \varphi$	Relative supremal quantification, 108
$t(V), t(\bar{v})$	Term t with relevant variables emphasized, 23
\mathcal{V}_b	Variable symbols for binding, 23
v_i	i th variable symbol for binding, 23
Ω	Weak modulus [BDNT17], 366

(Pre-)structures

$\text{acl}_{\mathfrak{M}}(A), \text{acl}(A)$	Algebraic closure of A in \mathfrak{M} , 134
$\bar{a} \cong \bar{b}$	\bar{a} is congruent to \bar{b} (Banach and Hilbert structures), 308
$\text{acor}(\mathfrak{M}, \mathfrak{N})$	Collection of almost correlations between \mathfrak{M} and \mathfrak{N} , 331
$a_{\Delta}(\mathfrak{M}, \mathfrak{N})$	Almost Δ -similarity between \mathfrak{M} and \mathfrak{N} , 332
$\alpha_{\mathfrak{M}, \mathfrak{N}}$	(Δ, Ω) -Scott rank of \mathfrak{M} and \mathfrak{N} , 371
$\bar{b} \equiv_A \bar{c}$	\bar{b} and \bar{c} have the same type over A , 31
$\text{cor}(\mathfrak{M}, \mathfrak{N})$	Set of correlations between \mathfrak{M} and \mathfrak{N} , 331
\mathfrak{C}_T	Monster model of the theory T , 74
$d^{\mathfrak{M}}(a, b)$	Metric on (pre-)structure \mathfrak{M} , 19

$\text{dcl}_{\mathfrak{M}}(A), \text{dcl}(A)$	Definable closure of A in \mathfrak{M} , 141
$d^{\mathfrak{M}}(\bar{a}, \bar{b})$	Metric on tuples, 18
$\text{dis}_{\Delta}(R)$	Distortion of R with respect to Δ , 332
$\text{eldiag}(\mathfrak{M})$	Elementary diagram of \mathfrak{M} , 31
$f : \mathfrak{M} \preceq \mathfrak{N}$	f is an elementary map from \mathfrak{M} to \mathfrak{N} , 31
ι	Variable assignment, 28
$\iota \left[\frac{\bar{a}}{\bar{x}} \right]$	ι with \bar{x} assigned to the values \bar{a} , 28
$\iota_{\mathcal{U}}$	Ultraproduct of $\{\iota_i\}_{i \in I}$ with regards to \mathcal{U} , 52
\mathcal{L}_C	\mathcal{L} expanded by the set of constants or elements C , 22
$\mathfrak{M}, \mathfrak{N}, \mathfrak{A}, \mathfrak{B} \dots$	Metric (pre-)structures, 19
$\mathfrak{M}_{\mathcal{U}}^0$	Pre-ultraproduct of $\{\mathfrak{M}_i\}_{i \in I}$ with regards to \mathcal{U} , 52
$\mathfrak{M}^{0, \mathcal{U}}$	Pre-ultrapower of \mathfrak{M} with regards to \mathcal{U} , 52
\mathfrak{M}_A	\mathfrak{M} expanded by constants for A , 22
$\mathfrak{M} \approx_{\Delta} \mathfrak{N}$	\mathfrak{M} and \mathfrak{N} are Δ -approximately isomorphic, 332
$\mathfrak{M} \approx_{\mathcal{L}} \mathfrak{N}$	\mathfrak{M} and \mathfrak{N} are \mathcal{L} -approximately isomorphic, 364
$\mathfrak{M} \equiv \mathfrak{N}$	\mathfrak{M} and \mathfrak{N} are elementarily equivalent, 31
$\mathfrak{M} \preceq \mathfrak{N}$	\mathfrak{M} is an elementary sub-(pre-)structure of \mathfrak{N} , 31

$\mathfrak{M} \models X$	\mathfrak{M} satisfies X , 30
\mathfrak{M}^{Mor}	Morleyization of \mathfrak{M} , 37
$\mathfrak{M}_{\mathcal{U}}$	Ultraproduct of $\{\mathfrak{M}_i\}_{i \in I}$ with regards to \mathcal{U} , 52
$\mathfrak{M}^{\mathcal{U}}$	Ultrapower of \mathfrak{M} with regards to \mathcal{U} , 52
$\mathfrak{M} \upharpoonright \mathcal{L}$	Reduct of \mathfrak{M} to \mathcal{L} , 20
$n(a)$	Name of a , 22
$\overline{\mathfrak{M}}$	Completion of \mathfrak{M} , 20
$\varphi^{\mathfrak{M}}(\bar{a})$	Evaluation of φ on $\bar{a} \in \mathfrak{M}$, 31
$\varphi^{\mathfrak{M}}(\iota)$	Evaluation of formula φ on ι in \mathfrak{M} , 28
$\rho_{\Delta}(\mathfrak{M}, \mathfrak{N})$	Δ -distance between \mathfrak{M} and \mathfrak{N} , 332
$r_{\alpha}^{\Delta, \Omega}(\mathfrak{M}, \bar{m}; \mathfrak{N}, \bar{n})$	(Δ, Ω) -back-and-forth pseudo-metric, 366
$r\mathfrak{M}$	Reduction of \mathfrak{M} , 20
$t^{\mathfrak{M}}(\iota)$	Evaluation of term t on ι in \mathfrak{M} , 28
$X(\mathfrak{M}, \bar{a})$	Set of tuples \bar{b} in \mathfrak{M} such that $\mathfrak{M} \models X(\bar{b}, \bar{a})$, 30

Theories

DI	Discrete Interval Theory, 114
IHS	Infinite Dimensional Hilbert Space, 306

PS	Polarized Square Theory, 115
T^{eq}	The complete imaginary expansion of T , 191
$\text{Th}(\Delta, \varepsilon)$	Theory of approximate isomorphism for Δ at scale ε , 349
$\text{Th}(\mathfrak{M})$	Theory of \mathfrak{M} , 31
T^{Mor}	Morleyization of T , 37

Type Space and Other Topometric Spaces

$\llbracket X \rrbracket_{V, \Sigma, \mathcal{L}}$	Set of types in $S_V(\Sigma, \mathcal{L})$ satisfying X , 47
$\text{cl}_{X, \tau} Q$	Topological closure of Q in (X, τ) , 82
$\overline{Q}, \overline{Q}^{X, d}, \overline{Q}^d$	Metric closure of Q in (X, d) , 82
$\delta_\Delta^\lambda(p, q), \delta_\Delta(p, q)$	Metric on type space induced by Δ , 344
$d_{\Delta, A}(p, q), d_\Delta(p, q)$	Δ -modified d -metric on type space, 375
$\dim(X)$	Lebesgue covering dimension of X , 564
$d_{\mathcal{L}_{\bar{v}/\bar{x}}, d/\bar{x}}$	Induced metric on $S_{\bar{v}\bar{x}}(\mathcal{L})$ over \bar{x} , 76
∂Q	Topological boundary of Q , 92
$d^{\text{tp}}(\bar{b}, \bar{c}/A)$	Alternative notation for $d(\text{tp}(\bar{b}/A), \text{tp}(\bar{c}/A))$, 77
$\text{ext}_{X, \tau} Q$	Topological exterior of Q in (X, τ) , 82
$f_{\text{co}}(A)$	Co-image of A under f , 559

$SG(X, P)$	Separator game with payoff set P , 551
$\text{Ind}(X)$	Large inductive dimension of X , 564
$\text{int}_{X, \tau} Q$	Topological interior of Q in (X, τ) , 82
$Q^\circ, Q_{X, d}^\circ, Q_d^\circ$	Metric interior of Q in (X, d) , 82
$OS(X)$	Set of ordered separators in X , 551
$\varphi(p)$	Evaluation of the formula φ on the type p , 46
$S_V(A)$	Set of V -types over the parameters A , 47
$S_V(\mathcal{L}), S_V$	Space of $\mathcal{L}(V)$ -types over the empty theory, 47
$S_V(\Sigma, \mathcal{L})$	Space of complete $\mathcal{L}(V)$ -types over Σ , 46
$\text{tp}_{\mathfrak{M}}(\iota), \text{tp}_{\mathfrak{M}}(\bar{a})$	Type of ι or \bar{a} in \mathfrak{M} , 31
$\text{tp}_{\mathfrak{M}}(\iota/A), \text{tp}_{\mathfrak{M}}(\bar{b}/A)$	Type of ι or \bar{b} over A in \mathfrak{M} , 31
(X, τ, d)	Topometric space X with topology τ and metric d , 78

Many-Sorted Signatures

0, 1	The constants 0 and 1 in the sort 2 , 188
2	The imaginary sort 2 , 188
n	Finite parameter sort n or constant in larger finite parameter sort, 202
$\bigsqcup_{n < \omega}^* O_n$	Countable metric disjoint union, 461

$d_{dGH}^T(O_0, O_1)$	Definable Gromov-Hausdorff distance between O_0 and O_1 , 213
$D(O)$	The definable set sort of D in O , 187
d_O	Metric of the sort O , 177
i	The inclusion map $i : D(O) \rightarrow O$, 187
$\mathcal{L}^{eq}, \mathcal{L}^{eq}(T)$	The complete imaginary expansion language of T , 191
$\mathcal{L} \upharpoonright \mathcal{S}$	Reduct of \mathcal{L} to the sorts \mathcal{S} , 178
$\mathfrak{M} \times_{\mathcal{S}} \mathfrak{N}$	Sort-by-sort product of \mathfrak{M} and \mathfrak{N} , 331
O/ρ	The quotient sort of O by the pseudo-metric ρ , 187
π_i	The projection map $\pi_i : \prod_{i < n} O_i \rightarrow O_i$, 186
$\prod_{i < n} O$	The imaginary product sort of the sequence of sorts $\{O_i\}$, 186
$\varphi : O_0 \approx_T O_1$	Metric correspondence between O_0 and O_1 over T , 213
q	The quotient map $q : O \rightarrow O/\rho$, 187
$\text{db}(O)$	Diameter bound of the sort O , 177
$S(f)$	Codomain sort of function or constant symbol f , 178
$S(t)$	Codomain sort of the term t , 178
$\mathcal{S}(\mathcal{L})$	Collection of sorts in \mathcal{L} , 177
$\mathcal{V}_b(O)$	The collection of variable symbols for binding of sort O , 178

$\mathcal{V}_b(\mathcal{S}, k)$	The set of variables of the form \dot{v}_k^O for some $O \in \mathcal{S}$, 181
\dot{v}_i^O	The i th variable symbol for binding of sort O , 178
$v:O$	A variable symbol of sort O , 178

Bibliography

- [AK15] Uri Andrews and H. Jerome Keisler. Separable Models of Randomizations. *The Journal of Symbolic Logic*, 80(4):1149–1181, dec 2015.
- [Bar86] Michael Barr. Models of sketches. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 27(2):93–107, 1986.
- [BB02] Philippe Blanchard and Erwin Bruening. *Mathematical Methods in Physics: Distributions, Hilbert Space Operators, and Variational Methods*. Progress in Mathematical Physics. Birkhäuser Boston, 2002.
- [BDNT17] Itai Ben Yaacov, Michal Doucha, André Nies, and Todor Tsankov. Metric Scott analysis. *Advances in Mathematics*, 318:46 – 87, 2017.
- [Ben13] Itai Ben Yaacov. On theories of random variables. *Israel Journal of Mathematics*, 194(2):957–1012, March 2013.
- [Bou89] Elisabeth Bouscaren. Elementary pairs of models. *Annals of Pure and Applied Logic*, 45(2, Part 1):129 – 137, 1989.
- [BPW00] Steven Buechler, Anand Pillay, and Frank Wagner. Supersimple theories. *J. AMS*, 14, 06 2000.
- [Bue17] Steven Buechler. *Essential Stability Theory*. Perspectives in Logic. Cambridge University Press, 2017.

- [BY05] Itai Ben Yaacov. Uncountable dense categoricity in cats. *Journal of Symbolic Logic*, 70(3):829–860, 2005.
- [BY08a] Itai Ben Yaacov. Continuous first order logic for unbounded metric structures. *Journal of Mathematical Logic*, 8(2):197–223, 2008.
- [BY08b] Itai Ben Yaacov. On perturbations of continuous structures. *Journal of Mathematical Logic*, 08(02):225–249, 2008.
- [BY08c] Itai Ben Yaacov. Topometric spaces and perturbations of metric structures. *Logic and Analysis*, 1(3):235, July 2008.
- [BY10a] Itai Ben Yaacov. Definability of groups in \aleph_0 -stable metric structures. *J. Symbolic Logic*, 75(3):817–840, 09 2010.
- [BY10b] Itai Ben Yaacov. Lipschitz functions on topometric spaces. *Journal of Logic and Analysis*, 10 2010.
- [BY14] Itai Ben Yaacov. The linear isometry group of the Gurarij space is universal. *Proceedings of the American Mathematical Society*, 142(7):2459–2467, 2014.
- [BYB04] Itai Ben Yaacov and Alexander Berenstein. Imaginaries in Hilbert spaces. *Archive for Mathematical Logic*, 43(4):459–466, May 2004.
- [BYBHU08] Itai Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov. *Model theory for metric structures*, volume 2 of *London Mathematical Society Lecture Note Series*, pages 315–427. Cambridge University Press, 2008.

- [BYJK09] Itaï Ben Yaacov and H. Jerome Keisler. Randomizations of models as metric structures. *Confluentes Mathematici*, 01(02):197–223, 2009.
- [BYU07] Itaï Ben Yaacov and Alexander Usvyatsov. On d-finiteness in continuous structures. *Fundamenta Mathematicae*, 194(1):67–88, 0 2007.
- [BYU10] Itaï Ben Yaacov and Alexander Usvyatsov. Continuous first order logic and local stability. *Transactions of the American Mathematical Society*, 362(10):5213–5259, 2010.
- [CDK18] Marek Cúth, Michal Doucha, and Ondřej Kurka. Complexity of distances between metric and Banach spaces. *arXiv e-prints*, 2018.
- [CK90] C.C. Chang and H.J. Keisler. *Model Theory*. ISSN. Elsevier Science, 1990.
- [CT16] Gabriel Conant and Caroline Terry. Model theoretic properties of the Urysohn sphere. *Annals of Pure and Applied Logic*, 167(1):49–72, January 2016.
- [EG12] Clifton Ealy and Isaac Goldbring. Thorn-forking in continuous logic. *The Journal of Symbolic Logic*, 77(1):63–93, mar 2012.
- [EKNK78] R. Engelking, D. Krupka, K. Nagami, and Y. Kodama. *Dimension Theory*. Mathematical Studies. North-Holland Publishing Company, 1978.
- [Esc09] Martín Hötzel Escardó. Intersections of compactly many open sets are open. 2009.
- [FM18] Ilijas Farah and Menachem Magidor. Omitting types in logic of metric structures. *Journal of Mathematical Logic*, 18(02):1850006, 2018.

- [GK19] Isaac Goldbring and H. Jerome Keisler. Continuous sentences preserved under reduced products, 2019.
- [GM83] M. Gromov and V. D. Milman. A topological application of the isoperimetric inequality. *American Journal of Mathematics*, 105(4):843–854, 1983.
- [Gro83] Mikhael Gromov. Filling Riemannian manifolds. *J. Differential Geom.*, 18(1):1–147, 1983.
- [GS15] Isaac Goldbring and Thomas Sinclair. Omitting types in operator systems. *Indiana University Mathematics Journal*, 66, 01 2015.
- [Han] James Hanson. Metric spaces are universal for bi-interpretation with metric structures. Submitted.
- [Han20] James Hanson. Strongly Minimal Sets and Categoricity in Continuous Logic. *arXiv e-prints*, page arXiv:2011.00610, November 2020.
- [HN18] Petr Hájek and Matěj Novotný. Distortion of Lipschitz functions on $c_0(\Gamma)$. *Proceedings of the American Mathematical Society*, 146(5):2173–2180, January 2018.
- [Hod93] Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.
- [HR16] C. Ward Henson and Yves Raynaud. Asymptotically Hilbertian modular Banach spaces: Examples of uncountable categoricity. *Commentationes Mathematicae*, 56, 06 2016.

- [Iov99] José Iovino. Stable models and reflexive banach spaces. *J. Symbolic Logic*, 64(4):1595–1600, 12 1999.
- [Joh48] Fritz John. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948*, pages 187–204. Interscience Publishers, Inc., New York, N. Y., 1948.
- [Kan03] A. Kanamori. *The Higher Infinite: Large Cardinals in Set Theory from Their Beginnings*. Springer monographs in mathematics. Springer, 2003.
- [KM99] H. Keisler and Arnold Miller. Categoricity without equality. *Fundamenta Mathematicae*, 170, 07 1999.
- [KP14] T. Kihara and A. Pauly. Point degree spectra of represented spaces. *ArXiv e-prints*, May 2014.
- [LT96] J. Lindenstrauss and L. Tzafriri. *Classical Banach spaces I and II*. Classics in mathematics. Springer-Verlag, 1996.
- [Mau95] B. Maurey. Symmetric distortion in l_2 . In *Geometric aspects of functional analysis (Israel, 1992–1994)*, volume 77 of *Oper. Theory Adv. Appl.*, pages 143–147. Birkhäuser, Basel, 1995.
- [MHS77] S.A. Morris, N.J. Hitchin, and London Mathematical Society. *Pontryagin Duality and the Structure of Locally Compact Abelian Groups*. Anthropological Studies of Systems of Meaning. Cambridge University Press, 1977.

- [Mil71] V. D. Milman. On a property of functions defined on infinite-dimensional manifolds. *Sov. Math., Dokl.*, 12:1487–1491, 1971.
- [Mor70] Michael D. Morley. The number of countable models. *J. Symb. Log.*, 35(1):14–18, 1970.
- [Noq17] Victoria Noquez. *Vaught’s Two-Cardinal Theorem and Notions of Minimality in Continuous Logic*. PhD thesis, University of Illinois at Chicago, 2017.
- [PP87] Anand Pillay and Bruno Poizat. Pas d’imaginaires dans l’infini! *J. Symbolic Logic*, 52(2):400–403, 06 1987.
- [Rei75] Joachim Reineke. Minimale gruppen. *Mathematical Logic Quarterly*, 21(1):357–359, 1975.
- [Roo78] A. C. M. van Rooij. *Non-Archimedean functional analysis / A. C. M. van Rooij*. M. Dekker New York, 1978.
- [SU11] Saharon Shelah and Alexander Usvyatsov. Model theoretic stability and categoricity for complete metric spaces. *Israel J. Math.*, 182:157–198, 2011.
- [SU19] Saharon Shelah and Alexander Usvyatsov. Minimal stable types in Banach spaces. *Advances in Mathematics*, 355:106738, 2019.
- [Tel10] Hernando Tellez. *Contributions to model theory of metric structures*. PhD thesis, University of Illinois at Urbana–Champaign, 2010.
- [TZ12] Katrin Tent and Martin Ziegler. *A Course in Model Theory*. Lecture Notes in Logic. Cambridge University Press, 2012.

- [Usu19] Toshimichi Usuba. G_δ -topology and compact cardinals. *Fundam. Math.*, 246(1):71–87, 2019.
- [Wei00] Klaus Weihrauch. *Computable Analysis: An Introduction*. Springer-Verlag, Berlin, Heidelberg, 2000.