HIGHER-ORDER REVERSE TOPOLOGY

by

James Hunter

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Mathematics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2008

ACKNOWLEDGMENTS

I would like to thank the following people for their invaluable assistance: Ulrich Kohlenbach, who was kind enough to invite me to speak at his university and to meet with him, and who answered my questions over e-mail and in person; Jeremy Avigad, who was helpful as well in explaining those aspects of his research that related to the subject of this dissertation and in providing suggestions for further reading; my Ph.D. advisor Steffen Lempp; and the other UW-Madison logic faculty, in particular Kenneth Kunen and Arnold Miller, who patiently answered my repeated questions.

I would like to thank my wife for supporting me through five years of graduate school.

DISCARD THIS PAGE

TABLE OF CONTENTS

Page

HIGHER-ORDER REVERSE TOPOLOGY

James Hunter

Under the supervision of Professor Steffen Lempp

At the University of Wisconsin-Madison

Reverse mathematics is the study of the relationships between logical axioms and mathematical theorems. Traditional reverse mathematics studies subsystems of second-order arithmetic, which means that it can examine only theorems expressible in the language of second-order arithmetic. To study a higher-order theorem, one must first find a way to encode the theorem using only first- and second-order objects.

Ulrich Kohlenbach described [Koh05] a higher-order theory RCA_0^ω, in all finite types, that is conservative over the second-order theory RCA_0 used as the base theory in traditional reverse mathematics. Using RCA_0^ω as a base theory, one can perform reverse-mathematics on statements of mathematical analysis that either cannot be expressed in the second-order language, or become trivial when expressed using second-order codes; see [Koh05].

We do some reverse mathematics, focusing on topology, over the base theory RCA_0^ω, examining the consequences of certain higher-order analogues of comprehension axioms. We also show that certain of these axioms are conservative over certain subsystems of second-order arithmetic studied in traditional reverse mathematics.

To examine further the (higher-order) reverse mathematics of topology, we construct a new base theory, extending and conservative over RCA_0^ω. We prove some conservation results for this new theory, in which we add a new base type for atoms, the space of which may have any cardinality. Sets (of atoms) are second-order (over atoms) objects, while families of sets—such as topologies—are third-order objects.

We examine two basic topological theorems and show that, over our new base theory, they are implied by a uniform analogue, over atoms, of Π^1_∞-comprehension and do not imply a uniform analogue, over atoms, of arithmetical comprehension.

Steffen Lempp

ABSTRACT

Reverse mathematics is the study of the relationships between logical axioms and mathematical theorems. Traditional reverse mathematics studies subsystems of second-order arithmetic, which means that it can examine only theorems expressible in the language of second-order arithmetic. To study a higher-order theorem, one must first find a way to encode the theorem using only first- and second-order objects.

Ulrich Kohlenbach described [Koh05] a higher-order theory RCA_0^ω, in all finite types, that is conservative over the second-order theory RCA_0 used as the base theory in traditional reverse mathematics. Using RCA_0^ω as a base theory, one can perform reverse-mathematics on statements of mathematical analysis that either cannot be expressed in the second-order language, or become trivial when expressed using second-order codes; see [Koh05].

We do some reverse mathematics, focusing on topology, over the base theory RCA_0^ω, examining the consequences of certain higher-order analogues of comprehension axioms. We also show that certain of these axioms are conservative over certain subsystems of second-order arithmetic studied in traditional reverse mathematics.

To examine further the (higher-order) reverse mathematics of topology, we construct a new base theory, extending and conservative over RCA_0^ω. We prove some conservation results for this new theory, in which we add a new base type for atoms, the space of which may have any cardinality. Sets (of atoms) are second-order (over atoms) objects, while families of sets—such as topologies—are third-order objects.

We examine two basic topological theorems and show that, over our new base theory, they are implied by a uniform analogue, over atoms, of Π^1_∞-comprehension and do not imply a uniform analogue, over atoms, of arithmetical comprehension.

Chapter 1

Reverse mathematics

1.1 Motivation

Reverse mathematics is the study of the relationships between logical axioms and mathematical theorems: reverse mathematics tries to find the minimal set of axioms required to prove a particular mathematical fact. Mathematical logic allows one to examine implications between axioms and theorems; for example, the relationship between choice axioms and various theorems has been widely examined. (For example, over the base theory ZF, the Axiom of Choice is equivalent to Tychonoff's Theorem, which states that the product of arbitrarily many compact spaces is compact.) However, the ZF axioms that form the basis for all of mathematics are sufficiently concise and general that weakening any of its axioms has not traditionally led to an interesting theory.

As a consequence, when logicians sought to prove equivalences among axioms other than choice they found ZF an unsuitable theory; in particular, what makes ZF so flexible—a language consisting only of one binary relation, "$\in$"—also gives some difficulty when trying to construct weaker subtheories. Instead, logicians studying reverse mathematics considered subsystems of second-order arithmetic; see, for example [Sim99].

The language of second-order arithmetic has two sorts of variables—one ranging over natural numbers, the other ranging over sets of natural numbers—along with the binary relation "$\in$," relating natural numbers to sets; the equality relation "$=$" for natural numbers; and those functions and constants, including "0" and "S," for successor, required for one to describe the natural numbers, $\mathbb{N}$. Subsystems of second-order arithmetic differ from what

Simpson, in his survey [Sim99] of the field, calls "full second-order arithmetic," in that they lack the full comprehension schema. Traditional reverse mathematics is inherently second-order.

One advantage of second-order reverse mathematics is that many results from recursion theory and its offshoots are relevant either immediately or after some effort to transfer statements about, say, computable or arithmetical sets on ω to statements about sets on a possibly non-standard version of ω (called "$\mathbb{N}$") that may satisfy only weak forms of induction. Today, techniques from recursion theory, computable model theory, and proof theory are all applicable to reverse mathematics.

However, the restriction that all statements considered in reverse mathematics be written in the language of second-order arithmetic is unsatisfying—and, in some cases, severe. For example, to express a function f on $\mathbb{R}$ in the language of second-order arithmetic (that is, the language of traditional reverse mathematics) one may use only integers (in $\mathbb{N}$) and real numbers (using the correspondence between subsets of $\mathbb{N}$ and the real numbers). Note that there are only $2^{\aleph_0}$ reals, while there are $(2^{\aleph_0})^{2^{\aleph_0}} = 2^{2^{\aleph_0}}$ functions on reals. At the very least the language restriction prohibits us from considering statements concerning all functions on $\mathbb{R}$.

But the restriction, as Kohlenbach [Koh05] noted, is more severe. The traditional reverse-mathematics workaround is to consider only continuous functions f on $\mathbb{R}$, coded using f's behavior on the countable dense subset $\mathbb{Q} \subset \mathbb{R}$. This means, first, that in reverse mathematics "continuous function" is a single unit: there are no discontinuous functions on $\mathbb{R}$. Worse, as Kohlenbach points out, the way one codes a continuous function on $\mathbb{R}$ in reverse mathematics allows one to read from the code—via a uniform procedure, computably from the code—for every real x and every $\varepsilon > 0$, a δ such that:

$$\forall y(|x-y| < \delta \to |f(x) - f(y)| < \varepsilon).$$

In reverse mathematics, a "continuous function" is not simply a function that satisfies the criteria for continuity, but rather a (code for a) package that, among other things, gives

δ's corresponding to ε's. A "continuous function" f tells you, for every input x, not only the value $f(x)$ but also how much f varies around x.

When considering the relative strength of mathematical statements one must be sensitive to the definitions chosen—in this case, the definition includes additional information, and this is an inherent limitation of trying to express a higher-order statement in second-order arithmetic.

Kohlenbach [Koh05] noted that the higher-order theory E-PRA$_0^\omega$ + QF-AC^{1,0} is conservative over its second-order fragment, which is equivalent to the base theory RCA$_0$ used in reverse mathematics (except with function variables in place of set variables). Kohlenbach named this higher-order base theory RCA$_0^\omega$, and proceeded to do reverse analysis over this base theory. RCA$_0^\omega$ allows for variables in all finite types, starting with the integers (type 0); this allows one to formulate sentences involving, for example, functions on $\mathbb{R}$.

In particular, Kohlenbach noted that various statements involving functions on $\mathbb{R}$ are equivalent (over the base theory RCA$_0^\omega$) to the existence of a discontinuous function on $\mathbb{R}$, which is in turn equivalent to a principle that he called $(\exists^2)$ and I will call $(\mathcal{E}_1)$:

$$(\mathcal{E}_1): \ \exists E_1^{(0\to 0)\to 0}\, \forall x^{0\to 0}\, (\exists n^0((xn) \neq_0 0) \leftrightarrow ((E_1\, x) =_0 1)).$$

The superscripts will be explained below—these are type designations for the variables. The principle $(\mathcal{E}_1)$ simply asserts the existence of a functional, E_1, that given a real number x returns the integer 1 if and only if x is not the infinite string of all 0's. The functional E_1 allows one to determine whether two reals are equal. There are natural higher-order analogues $(\mathcal{E}_n)$ as well (see [AF98], p. 350).

We show, in Theorem 2.5, that the theory RCA$_0^\omega$ $+(\mathcal{E}_1)$ is conservative over and implies the reverse-mathematical theory ACA$_0$; hence, the second-order part of a theory that implies the existence of a discontinuous function on $\mathbb{R}$ is no weaker than ACA$_0$.

1.2 Topology

The requirement that all statements be formulated in the language of second-order arithmetic has restricted the development of reverse topology. To formulate topological statements one usually wants to refer both to points in a topological space and to open subsets of that space. Because traditional, second-order reverse mathematics is restricted to sentences with only second-order variables, one can refer only to spaces of size $\aleph_0$ or $2^{\aleph_0}$.

At the same time, open sets must correspond to second-order codes, which in practice means that second-order reverse topology is restricted to spaces with countable bases. See Carl Mummert's recent work, [Mum05] and [Mum07], which advances the state of the art of traditional reverse topology, using (codes for) maximal and unbounded filters on $\mathbb{N}$.

The traditional work-around, as advanced in Mummert [Mum05], is to consider countably-based spaces of size $2^{\aleph_0}$. Then one can code both points and open sets by sets of integers. Because the predicate that determines whether a given point is in a given set is a third-order object, one then defines a formula for "$x \in U$," where x is a (second-order code for a) point and U is a (second-order code for an) open set. For a different approach, see the recent work by Iraj Kalantari and Lawrence Welch on point-free topology, [KW98] and [KW04].

The restrictions of second-order reverse mathematics are particularly unsatisfying with respect to topology. A basic topological theorem such as, "The product of two compact spaces is compact," cannot be expressed in the language of second-order arithmetic, and so traditional reverse mathematics cannot examine the relative strength of this theorem. The best one could hope for is a variant of the form:

"The product of two compact, countably-based spaces of size $2^{\aleph_0}$, where set membership is defined by a particular fixed formula, is compact,"

which is not quite the same. Many basic topological theorems make no reference to the cardinalities of the space.

Much reverse topology—although not called by that name!—has already been done using ZF as a base theory and examining the topological consequences of various choice principles.

(See, for example, Jech's 1973 book *The Axiom of Choice* [Jec73].) However, there are many basic topology theorems that are provable in ZF and not expressible in traditional, second-order, reverse mathematics. For example, Tychonoff's Theorem, "The product of (arbitrarily many) compact spaces is compact," is equivalent to the Axiom of Choice, over ZF.

However, the restriction of Tychonoff's Theorem to finite products is provable in ZF and not expressible in the language of second-order arithmetic. We show, in Section 3.4.1, that the strength of the statement, "The product of two compact spaces is compact," lies strictly between a higher-order axiom conservative over RCA_0 and a higher-order axiom conservative over ACA_0.

Note that once one leaves second-order codes behind many of the well-established and powerful techniques of traditional reverse mathematics become useless. In this sense second-order reverse mathematics is more powerful than its higher-order variant: if you can express a theorem, in a reasonable way, as a sentence in the language of second-order arithmetic, then you can often determine more precisely the logical strength of that theorem. Higher-order reverse mathematics is currently a new field, and the techniques that we currently know how to use are crude. However, to the extent that certain theorems have no reasonable second-order analogue, there is no alternative to higher-order reverse mathematics.

We adopt and extend Kohlenbach's base theory RCA_0^ω and apply it to examine the higher-order reverse mathematics of certain basic topological statements. The base theory RCA_0^ω, while well-suited to examining functions on $\mathbb{R}$ or the Baire space ${}^{\mathbb{N}}\mathbb{N}$, has certain issues when applied to topology. For example, in RCA_0^ω the cardinality of a topological space is built into the type of the objects chosen as elements of the space. If one uses type-0 elements, the topological space is countable. If one instead uses type-1 elements, the topological space has size $2^{\aleph_0}$, in the sense of the model.

Further, the axiom that asserts the existence of a functional determining whether two arbitrary sets of elements are equal—necessary for defining the family of all sets intersecting a particular set—has strong consequences. In particular, if one uses type-1 elements for a topological space, the corresponding axiom's second-order consequences are Π^1_∞-CA_0: second-order arithmetic with full comprehension.

To address to these two issues we define and consider, in Chapter 3, a new base theory in which the elements of our topological spaces are atoms, rather than integers or reals. We show that two basic topological theorems lie strictly between two natural comprehension schemas.

Chapter 2

Base theory $\mathbf{RCA}_0^\omega$ and related theories

2.1 Definitions

2.1.1 Finite types

The language of traditional reverse mathematics includes variables of only two types: the type of integers ($\mathbb{N}$) and the type of sets of integers (some version of $\mathcal{P}(\mathbb{N})$). For higher-order reverse mathematics we use function variables, instead of set variables—*e.g.*, ${}^{\mathbb{N}}\mathbb{N}$ instead of $\mathcal{P}(\mathbb{N})$—and we have variables of all finite types.

The finite types are defined inductively:

Definition 2.1. *0 is a finite type: the type of $\mathbb{N}$.*

If σ and τ are finite types, then so is $(\sigma \to \tau)$: the type of a function with input σ and output τ.

We also have the following abbreviations for the standard types:

Definition 2.2. *0 is a standard type.*

If n is a standard type, then $n+1$ is the standard type $(n \to 0)$.

The types used in traditional, second-order, reverse mathematics are roughly equivalent to the standard types 0 and 1.

In section 3.1.1 we introduce a second atomic type, α, in addition to the standard atomic type 0, and allow finite types to be built from 0, α, and arrows. Every variable in the theories considered in this paper is of one of these types.

We sometimes denote a variable's or term's type by superscripts. For example, "x^σ" denotes a variable x of type σ; however, to improve readability we frequently drop superscripts where the type is clear from context.

2.1.2 The theory $\mathbf{RCA}_0^\omega$

We follow the definition from [Koh05]. The language of RCA_0^ω includes variables x^σ of all finite types σ, quantifiers $\exists^\sigma$ and $\forall^\sigma$ for all finite types σ, and the type-0 equality relation $=_0$.

The axioms of RCA_0^ω consist of axioms defining the standard (typed) lambda calculus combinators $\Sigma_{\rho,\sigma,\tau}$ and $\Pi_{\sigma,\tau}$ (known to computer scientists as "S" and "K," respectively):

- $(\Sigma_{\rho,\sigma,\tau} x^\rho y^\sigma z^\tau) = xz(yz)$, and
- $(\Pi_{\sigma,\tau} x^\sigma y^\tau) =_\sigma x$,

where "xz" denotes the application of the function(al) represented by x to the parameter represented by z. The ordinary left-to-right precedence rules apply, so "$xz(yz)$" is equivalent to "$(xz)(yz)$." The Σ and Π combinators allow one to define λ-abstraction (see [Tro73], pp. 41–42); one can think of Π as a projection operator, and Σ as an application operator. In this paper we will find it easier to use λ in place of the combinators, outside of this section.

We use "$=_\sigma$," for all types $\sigma \neq 0$, as an abbreviation for extensional equality, defined inductively: $=_0$ is a relation, and if $\sigma = (\beta \to \gamma)$ then:

$$x =_{(\beta\to\gamma)} y \text{ abbreviates } \forall z^\beta (xz =_\gamma yz).$$

Note that we have omitted the subscript from the equality relation in the Σ-combinator rule. Formally, in the Σ-combinator rule, we require that $\rho = (\tau \to (\beta \to \gamma))$ and $\sigma = (\tau \to \beta)$, which means the equality relation is $=_\gamma$. For Π, the equality relation is $=_\sigma$, as shown in the defining equation above.

We also include the basic equality axioms for $=_0$, along with axioms defining the successor function S on $\mathbb{N}$ and the constant $\mathbf{0} \in \mathbb{N}$. We include axioms defining the primitive-recursion operator $\mathcal{R}_0$:

- $(\mathcal{R}_0\, x^0\, y^{(0,0)\to 0}\, \mathbf{0}) =_0 x$, and
- $(\mathcal{R}_0\, x^0\, y^{(0,0)\to 0}\, S(z^0)) =_0 (y(\mathcal{R}_0\, x\, y\, z)z)$.

For simplicity we write "$(\rho,\sigma) \to \tau$" for "$\rho \to (\sigma \to \tau)$." (In other words, we sometimes write the type as if the function had been "curried.") Using $\mathbf{0}$, S, and $\mathcal{R}_0$ one can define all primitive recursive type-1 functions on $\mathbb{N}$.

Finally, we include the schema QF-IA of induction for quantifier-free formulas (with parameters) and the quantifier-free choice schema (also with parameters) QF-AC^{1,0}:

$$\text{QF-AC}^{1,0}\text{: } \forall x^1 \exists y^0 \Phi(x,y) \to \exists F^2 \forall x^1 \Phi(x, Fx),$$

where Φ is a quantifier-free formula. Note that adding QF-AC^{1,0} gives us all recursive (or computable) type-1 functions—not just the primitive recursive functions.

Justifying Kohlenbach's choice of higher-order base theory is the following, which is Proposition 3.1 in [Koh05]:

Proposition 2.3. RCA_0^ω *is conservative over and implies* RCA_0.

(Note, further, that we could also add QF-AC$^{\sigma,0}$, for any type σ, and still have the resulting theory be conservative over RCA_0.)

Formally, the theory RCA_0 refers to set variables, not function variables. In fact Kohlenbach defined a function analogue to RCA_0, which he called RCA_0^2. In this paper, as in [Koh05], we omit the superscript and treat RCA_0^2 and RCA_0 as interchangeable.

2.1.3 Additional axioms yielding stronger theories

Kohlenbach [Koh05] showed that several statements in mathematical analysis are equivalent to the principle $(\mathcal{E}_1)$, defined in section 1.1. One can think of $(\mathcal{E}_1)$ as an axiom asserting the existence of the functional E_1 that determines whether a type-1 object is the infinite string consisting only of 0's:

$$\exists n^0\, ((x^1\, n) \neq_0 0) \iff (E_1 x) =_0 1.$$

From E_1 one can define a type-$((1,1) \to 0)$ equality functional $\mathcal{F}$ such that $(\mathcal{F}xy) =_0 1$ if and only if $x =_1 y$:

$$\mathcal{F} :\equiv (\lambda x^1.(\lambda y^1.(E_1(\lambda n^0.((xn) \dot{-} (yn)) + ((yn) \dot{-} (xn))|))))),$$

where "$\dot{-}$" is defined from $\mathbf{0}$, S, and $\mathcal{R}_0$. It often helps to think of E_1 as the type-1 equality operator.

Of course one can similarly define the type-$(\sigma \to 0)$ "equality operator" $E_{\sigma\to 0}$ for any type σ:

$$(\mathcal{E}_{\sigma\to 0}):\ \exists E_{\sigma\to 0}^{(\sigma\to 0)\to 0}\, \forall x^{\sigma\to 0}\, (\exists y^{\sigma}((xy) \neq_0 0) \leftrightarrow ((E_{\sigma\to 0}\, x) =_0 1)).$$

The two such functionals to which we will refer are E_1 and E_2. Abusing notation somewhat, we write "$\sigma + 1$" for $(\sigma \to 0)$.

Proposition 2.4 (RCA_0^ω). *For every finite type σ:*

1. $(\mathcal{E}_{\sigma+2})$ *(i.e., $(\mathcal{E}_{(\sigma\to 0)\to 0})$) implies* $(\mathcal{E}_1)$.
2. $(\mathcal{E}_{\sigma+3})$ *implies* $(\mathcal{E}_2)$.
3. *For all $n \geq 3$, $(\mathcal{E}_{\sigma+(n+1)})$ implies $(\mathcal{E}_n)$.*

In particular, $(\mathcal{E}_3) \implies (\mathcal{E}_2) \implies (\mathcal{E}_1)$.

Proof. **1.** Define E_1' by the rule:

$$E_1' :\equiv (\lambda x^1.\, (E_{\sigma+2}(\lambda y^{\sigma+1}.\, x(ya^{\sigma})))\,),$$

where a is some fixed type-σ element.

(The idea is that we want to define a type-$(1 \to 0)$ functional that returns 1 if and only if its type-1 input is not the constant-0 function. Any such functional must take a type-1 input—we use x–and return either 0 or 1. We want to use $E_{\sigma+2}$, which returns 1 if and only if its type-$(\sigma+2)$ input is not the constant-0 function, so we build a type-$(\sigma+2) = ((\sigma+1) \to 0)$ functional to which to apply $E_{\sigma+2}$.

(That type-$(\sigma+2)$ functional should depend on x, and in particular we want to use the functional's input—y, here—to generate a type-0 input for x. The result is the functional E'_1, defined above, which returns 1 if and only if $(x(ya)) \neq_0 0$ for some y. To show that E'_1 is, in fact, E_1, we need to show that every type-0 number n^0 is (ya), for some y.)

Note that $(E'_1 x) =_0 1$ if and only if $(E_{\sigma+2}(\lambda y.\, x(ya))) =_0 1$, if and only if there is some $y^{\sigma+1}$ such that $x(ya) \neq_0 0$.

For any $n \in \mathbb{N}$, RCA_0^ω proves the existence of the constant type-$\sigma+1$ functional $(\lambda z^\sigma.n)$. As a consequence, there is some $y^{\sigma+1}$ such that $x(ya) \neq_0 0$ if and only if there is some n^0 such that $(xn) \neq_0 0$, so $E'_1 \equiv E_1$.

2. Define E'_2 by the rule:

$$E'_2 :\equiv (\lambda X^2.\, (E_{\sigma+3}(\lambda y^{\sigma+2}.\, X(\lambda m^0.y(\lambda z^\sigma.m))))\,).$$

That is, E'_2 applies $E_{\sigma+3}$ to the functional that maps $y^{\sigma+2}$ to the result of applying X to the type-1 element $(\lambda m.y(\lambda z.m))$—which, for each m^0, applies y to the constant functional $(\lambda z.m)$. Note that E'_2 returns 1 if and only if $(X(\lambda m.y(\lambda z.m))) \neq_0 0$ for some y.

For every type-1 element x, there is a type-$(\sigma+2)$ element y that "codes" x, defined by:

$$y :\equiv (\lambda w^{\sigma+1}.x(wa^\sigma)),$$

where a is some fixed type-σ element. As in the previous proof, this shows that $E'_2 \equiv E_2$, since:

$$\begin{aligned}
(\lambda y.X(\lambda m.y(\lambda z.m)))\ (\lambda w.x(wa)) &=_0 (X(\lambda m.((\lambda w.x(wa))(\lambda z.m)))) \\
&=_0 (X(\lambda m.x((\lambda z.m)a))) \\
&=_0 (X(\lambda m.(xm))) \\
&=_0 (Xx).
\end{aligned}$$

3. Define E'_n by the rule:

$$E'_n :\equiv (\lambda X^n.\, (E_{\sigma+(n+1)}(\lambda y^{\sigma+n}.X(\lambda f^{n-2}.y(\lambda z^{\sigma+(n-2)}.f(za)))))\,),$$

where a is some fixed type-$(\sigma + (n-3))$ element.

Note that $(E'_n X) =_0 1$ if and only if $(E_{\sigma+(n+1)}(\lambda y.X(\lambda f.y(\lambda z.f(za))))) =_0 1$, if and only if there is some y such that $(X(\lambda f.y(\lambda z.f(za)))) \neq 0$.

If $(E'_n X) =_0 1$ then $(E_n X) =_0 1$; it remains for us to show the converse. Fix type-n functional X such that for some type-$(n-1)$ input x, $(Xx) \neq_0 0$. We want to show that there is some type-$(\sigma + n)$ functional y_0 such that $(\lambda f.y_0(\lambda z.f(za))) =_{n-1} x$. But this is true, since we can define such a y_0, in RCA_0^ω, by the rule:

$$(y_0\, w^{\sigma+(n-1)}) :=_0 x(\lambda m^{n-3}.(w(\lambda v^{\sigma+(n-2)}.m))).$$

Then, for all type-$(n-2)$ elements f:

$$\begin{aligned}(y_0(\lambda z.f(za))) &=_0 x(\lambda m.(\lambda z.f(za))(\lambda v.m)) \\ &=_0 x(\lambda m.f((\lambda v.m)a)) \\ &=_0 x(\lambda m.(fm)) \\ &=_0 (xf).\end{aligned}$$

□

One can similarly introduce schemas $\mathrm{QF\text{-}AC}^{\sigma,\tau}$ of choice axioms for any pair of types σ and τ:

$$(\mathrm{QF\text{-}AC}^{\sigma,\tau})\text{: } [\forall x^\sigma\, \exists y^\tau\, \Phi(x,y)] \rightarrow [\exists F^{\sigma\rightarrow\tau}\, \forall x^\sigma\, \Phi(x,(Fx))],$$

where $\Phi(x,y)$ is quantifier-free. For example, $\mathrm{QF\text{-}AC}^{0,1}$ corresponds to Weak Countable Choice. However, there does not seem to be an advantage in using RCA_0^ω as a base theory instead of ZF in considering the logical consequences of choice principles, and so we have mostly avoided examining these schemas.

2.2 Conservation results for higher-order theories

In [Koh05], Kohlenbach refers to the theory $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ and notes (Theorem 3.2, citing [Fef77]) that it is conservative over the first-order theory of Peano Arithmetic (PA).

It is not difficult to show that, over RCA_0^ω, $(\mathcal{E}_1)$ implies arithmetic comprehension (and hence ACA_0). First, if $\Phi(n^0)$ is a quantifier-free formula with only one free variable of type 0 then it is equivalent to:

$$f(n) =_0 1,$$

where f is some function definable from RCA_0^ω. Second:

$$\exists n^0(f(n) =_0 1) \iff (E_1 f) =_0 1,$$

so quantifiers can be replaced by E_1's, yielding a quantifier-free formula equivalent to the original.

If X is a set of integers defined by:

$$n^0 \in X \iff \Phi(n),$$

where Φ is an arithmetical formula, then we also have:

$$n^0 \in X \iff f(n) =_0 1,$$

so f is a characteristic function for X.

So the second-order part of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ includes ACA_0. Jeremy Avigad and Solomon Feferman, in [AF98], state (Theorem 8.3.4) that $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ is conservative over ACA_0 for Π^1_2 sentences. In fact, this holds for all sentences:

Theorem 2.5. $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ *is conservative over and implies* ACA_0.

One can think of E_1 as a uniform version of arithmetical comprehension.

Proof. Let $\mathcal{M}$ be a (second-order) model of ACA_0, where ACA_0 is formalized using function rather than set variables. We will construct, from $\mathcal{M}$, a new model $\mathcal{N}$ of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ whose second-order part—the collection of all type-0 and type-1 elements—is isomorphic to $\mathcal{M}$. If a second-order sentence Φ is provable from $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ then it holds in all models of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$, so it holds in $\mathcal{N}$. Since Φ involves only second-order objects, it holds in $\mathcal{M}$; since $\mathcal{M}$ was arbitrary, Φ holds in all models of ACA_0 and hence is provable from ACA_0.

Our model $\mathcal{N}$ will be a term model: its elements will be equivalence classes of finite strings, or terms. We start by including constant symbols for all elements of $\mathcal{M}$—we call the type-0 constant symbols "numerals"—and all functionals defined in RCA_0^ω. Following a comment by Avigad and Feferman in [AF98], instead of including a symbol for E_1, we include a symbol for μ, defined by:

$$\mu(x^1) =_0 \begin{cases} n+1, & \text{for } n^0 \text{ least such that } (xn) \neq_0 0, \text{ if such } n \text{ exists,} \\ 0, & \text{otherwise.} \end{cases}$$

Avigad and Feferman noted that the principle (μ), which asserts the existence of μ, is equivalent to $\mathrm{QF\text{-}AC}^{1,0} + (\mathcal{E}_1)$. By including μ instead of E_1 we avoid having to verify that $\mathcal{N}$ satisfies $\mathrm{QF\text{-}AC}^{1,0}$.

The terms from which the model $\mathcal{N}$ is constructed consist of all finite strings built up from variables and the constant symbols listed above, by λ-abstraction and application. We require that applications respect types; so, for example:

$$(\lambda x^1.(x\mathbf{n}))$$

is a valid term, for all $n \in \mathbb{N}$, but:

$$(\lambda x^0.(x\mathbf{n}))$$

is not. We use boldface to distinguish constant symbols: n is an element of the model $\mathcal{M}$'s collection $\mathbb{N}$ of type-0 elements, while $\mathbf{n}$ is the constant symbol, or numeral, corresponding to n.

The equivalence relation $\approx$ on terms has two components: syntactic equivalence and semantic equivalence. The syntactic rule is:

$$((\lambda x^\sigma.t_1[x])t_2) \approx t_1[t_2/x],$$

where by "$t_1[t_2/x]$" we mean the finite string t_1 with all (free) instances of x replaced by the finite string t_2. (Notice that we use "$[x]$" to list the variable free in t_1; we use brackets here instead of parentheses to avoid confusion.) This is the standard reduction relation for the λ-calculus: the term $((\lambda x.t_1[x])t_2)$ reduces to $t_1[t_2/x]$. Since we are using a finite-typed λ-calculus, it is straightforward to show that all sequences of reductions must terminate uniquely. The syntactic reduction rule applies to closed terms of all types.

In addition, for closed type-0 terms we add semantic rules for the intrinsic functionals—μ, $\mathcal{R}_0$, S, and the second-order elements of $\mathcal{M}$—for which constant symbols have been included. (Our terms use λ's instead of the combinators Π and Σ, so we do not add rules for the combinators.) The rules are:

- $(\mathbf{fn}) \approx \mathbf{f}(\mathbf{n})$, where n and f are first- (type-0) and second-order (type-1) elements of $\mathcal{M}$, respectively. The application of the type-1 function constant $\mathbf{f}$ to a type-0 constant is equivalent to the type-0 constant, or numeral, $\mathbf{f}(\mathbf{n})$.
- Similarly, $(S\mathbf{n}) \approx$ the numeral $(\mathbf{n}+\mathbf{1})$.
- $(\mu\, t)$ is equivalent to the numeral $\mathbf{0}$ if and only if for all numerals $\mathbf{p}$, $(t\mathbf{p}) \approx \mathbf{0}$, and is equivalent to $(\mathbf{p}+\mathbf{1})$ for the least numeral $\mathbf{p}$ such that $(t\mathbf{p}) \not\approx 0$, otherwise.
- $(\mathcal{R}_0\, \mathbf{a}\, t\, \mathbf{0}) \approx \mathbf{a}$, while $(\mathcal{R}_0\, \mathbf{a}\, t\, (\mathbf{b}+\mathbf{1}))$ is equivalent to the numeral $\mathbf{p}$ if and only if there is a numeral $\mathbf{q}$ such that $(\mathcal{R}_0\, \mathbf{a}\, t\, \mathbf{b}) \approx \mathbf{q}$ and $(t\, \mathbf{q}\, \mathbf{b}) \approx \mathbf{p}$.

Each of these rules explicitly relates a closed term of the specified form to at most one numeral; however, we will still need to verify that distinct numerals yield distinct equivalence classes. Also, all but the last rule, for $\mathcal{R}_0$, relates its term to at least one numeral. Handling $\mathcal{R}_0$ is a special case: we must relate primitive recursion in $\mathcal{N}$ to primitive recursion in $\mathcal{M}$.

Combining the syntactic and semantic rules, and taking their transitive closure, yields an equivalence relation on closed type-0 terms of $\mathcal{N}$. For terms of $\mathcal{N}$ of other types, we add to the syntactic reduction rule the following extensional rule. First, two terms are equivalent only if they have the same type (this is consistent with the existing syntactic and type-0 semantic rules). Second, if s and t have type $(\sigma \to \tau)$ then $s \approx t$ provided that every closed type-σ term u, $su \approx tu$ (this is also consistent with the existing rules). The equivalence relation $\approx$ on all closed terms is the transitive closure of the union of these three sets of rules.

The elements of $\mathcal{N}$ are the equivalence classes of closed terms—terms with no free variables. The relation $=_0$ is just $\approx$, restricted to type-0 objects.

We first show that every closed type-0 term of $\mathcal{N}$ is equivalent to some numeral. Let t be an arbitrary closed type-0 term. Applying the normal-form theorem from λ-calculus, there is a unique closed term t' such that any sequence of reductions starting from t terminates at t'. Let t' be the "normal form" of t, the unique term obtained from t by syntactic equivalence such that no further reduction is possible. So $t \approx t'$, and so without loss of generality we may assume that t is in normal form.

Every variable occurring in t (not necessarily free in t) must be of type 0. This is because if the type-σ variable y, where $\sigma \neq 0$, occurs in t then y is free in a some subterm of t. Since, by assumption, y is not free in t, it must be captured by a λ—that is, y occurs in t as part of a subterm of the form:

$$(\lambda y.s[y]).$$

Notice that this subterm has type $(\sigma \to \tau)$ for some τ, and that $\sigma \neq 0$. Since t is of type 0, the subterm $(\lambda y.s[y])$ must occur in t as part of an application. (Applying a term to another term lowers the combined term's type.)

None of the constants included in the term model $\mathcal{N}$ can be applied to terms of type $(\sigma \to \tau)$ for $\sigma \neq 0$. That means that if the application is:

$$(s'(\lambda y.s[y])),$$

then s' is itself a λ expression, so we can apply reduction—contradicting our assumption that t is in normal form.

If, instead, $(\lambda y.s[y])$ occurs on the left-hand side of an application:

$$((\lambda y.s[y])s'),$$

then we can again apply reduction—a contradiction. So t contains only variables of type 0.

Let s be an arbitrary (not necessarily closed) type-0 subterm of t, with free type-0 variables $x_1, \ldots, x_k$. Using induction, we will show that for every assignment $\mathbf{m}_1, \ldots, \mathbf{m}_k$ of numerals to $x_1, \ldots, x_k$, the closed term:

$$s[\mathbf{m}_1/x_1, \ldots, \mathbf{m}_k/x_k],$$

is equivalent to some numeral $\mathbf{n}$ from $\mathcal{M}$. In particular, this will show that the closed term t is equivalent to a numeral.

Our base cases are where s is either a variable x or a numeral $\mathbf{m}$; in the first case, $s[\mathbf{m}/x]$ is trivially equivalent to $\mathbf{m}$, and in the second case $s[]$ is also trivially equivalent to $\mathbf{m}$.

Our induction cases correspond to our semantic rules. Suppose that for every $\overline{\mathbf{m}}$ there is some $\mathbf{n}$ such that $s[\overline{\mathbf{m}}/\overline{x}] \approx \mathbf{n}$:

- $(\mathbf{f}\, s[\overline{\mathbf{m}}/\overline{x}])$ is equivalent to the numeral $\mathbf{f}(\mathbf{n})$.
- $(S s[\overline{\mathbf{m}}/\overline{x}])$ is equivalent to the numeral $(\mathbf{n} + \mathbf{1})$.
- $(\mu(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]))$ is equivalent to $\mathbf{0}$ if for all numerals $\mathbf{m}_1$, the corresponding numeral $\mathbf{n} \approx \mathbf{0}$, and otherwise is equivalent to $(\mathbf{m}_1 + \mathbf{1})$ for the least $\mathbf{m}_1$ such that the corresponding numeral $\mathbf{n} \not\approx \mathbf{0}$.

  (Note that every type-1 term is equivalent to a term of the form $(\lambda x.s)$, where s is of type 0. This is because of our extensionality rule.)
- $(\mathcal{R}_0\, \mathbf{a}\, t\, \mathbf{0}) \approx \mathbf{a}$. For $(\mathcal{R}_0\, \mathbf{a}\, t\, (\mathbf{b} + \mathbf{1}))$, note that the type-$(0 \to (0 \to 0))$ term t is equivalent to:

  $$(\lambda x_1.(\lambda x_2.s[\overline{x}])),$$

  for some type-0 term s. So we can apply induction. We need to show that there are numerals $\mathbf{p}$ and $\mathbf{q}$ satisfying:

  $$(\mathcal{R}_0\, \mathbf{a}\, (\lambda x_1.(\lambda x_2.s[x_1, x_2, \mathbf{m}_3/x_3, \ldots, \mathbf{m}_k/x_k]))\, \mathbf{b}) \approx \mathbf{q}, \tag{2.1}$$

  and:

  $$((\lambda x_1.(\lambda x_2.s[x_1, x_2, \mathbf{m}_3/x_3, \ldots, \mathbf{m}_k/x_k]))\, \mathbf{q}\, \mathbf{b}) \approx \mathbf{p}. \tag{2.2}$$

  Using syntactic reduction, equation (2.2) is logically equivalent to:

  $$s[\mathbf{q}/x_1, \mathbf{b}/x_2, \mathbf{m}_3/x_3, \ldots, \mathbf{m}_k/x_k] \approx \mathbf{p},$$

so we can choose $\mathbf{p}$ to be the $\mathbf{n}$ corresponding to $s[\mathbf{q}/x_1, \mathbf{b}/x_2, \mathbf{m}_3/x_3, \ldots, \mathbf{m}_k/x_k]$—assuming we have found a numeral $\mathbf{q}$ that satisfies equation (2.1). This is the special case, and we prove the existence of $\mathbf{q}$ by applying induction in $\mathcal{M}$.

The key point here is that $\mathcal{M}$ satisfies ACA, and hence induction for arithmetical formulas. Simultaneously with the induction already described we will prove that for each subterm s there is an arithmetical formula $\Phi(x_1, \ldots, x_k, y)$ in the language of $\mathcal{M}$ such that:

$$(\mathcal{M} \models \Phi[m_1/x_1, \ldots, m_k/x_k, n/y]) \implies (s[\overline{\mathbf{m}}/\overline{x}] \approx \mathbf{n}). \tag{2.3}$$

and:

$$\mathcal{M} \models \forall \overline{x}\, \exists! y\, \Phi(\overline{x}, y). \tag{2.4}$$

The idea is that the arithmetical formula Φ, built from the term $s[\overline{x}]$, tells us how to find $\mathbf{n}$. Equation (2.4) says that, in $\mathcal{M}$, Φ is the graph of a function $(x_1, \ldots, x_k) \mapsto y$. (Notice that we use parentheses to list variables free in the formula Φ, and brackets to list assignments of elements of $\mathcal{M}$ to variables in Φ.)

If $s \equiv x$ then Φ is "$x = y$"; if s is a numeral $\mathbf{m}$ then Φ is "$m = y$." The induction cases, where by assumption we have a Φ that works for $s[\overline{x}]$, are:

- For $(\mathbf{f}s[\overline{\mathbf{m}}/\overline{x}])$, use formula: $\exists z(\Phi(\overline{x}, z) \wedge f(z) = y)$.
- For $(Ss[\overline{\mathbf{m}}/\overline{x}])$, use formula: $\exists z(\Phi(\overline{x}, z) \wedge (z + 1 = y))$.
- For $(\mu(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]))$, use formula:

$$\begin{aligned}&[y = 0 \wedge \forall z \Phi(z, x_2, \ldots, x_k, 0)] \\ &\qquad \vee [\exists z' \neq 0\ (\Phi(y-1, x_2, \ldots, x_k, z') \wedge \forall z < (y-1)\ (\Phi(z, x_2, \ldots, x_k, 0)))].\end{aligned} \tag{2.5}$$

  (The second half of the disjunction just says that y is least such that $(\lambda x_1.s)(y-1) \neq 0$.)

- For $(\mathcal{R}_0\, \mathbf{a}\, t\, \mathbf{0})$, use formula: $y = a$. For $(\mathcal{R}_0\, \mathbf{a}\, t\, (\mathbf{b} + \mathbf{1}))$, use formula:

$$\begin{aligned}\Psi(z, y) \equiv \exists c\, [|c| = z + 1 \ \wedge\ (c)_0 = a \ \wedge\ (c)_z = y \\ \wedge\ \forall (0 \leq k < z)\, \Phi[(c)_k/x_1, k/x_2, (c)_{k+1}/y]],\end{aligned}$$

where $|c|$ gives the length of the finite string encoded by c, $(c)_k$ gives the kth element of c, starting with index 0, and Φ is the formula corresponding to t. (Recall that we have already noted that t is equivalent to $(\lambda x_1.(\lambda x_2.s))$, for some type-0 term s.)

The formula $\Psi[(b+1)/z]$ just asserts the existence of a type-0 code c, in $\mathcal{M}$, for the primitive recursive computation corresponding to $(\mathcal{R}_0\,\mathbf{a}\,t\,(\mathbf{b}+\mathbf{1}))$. (Note that $\Psi[(b+1)/z, n/y]$ holds for at most one n.)

If $\Psi[(b+1)/z, n/y]$ holds for some n, then $(\mathcal{R}_0\,\mathbf{a}\,t\,(\mathbf{b}+\mathbf{1})) \approx \mathbf{n}$. To see this, note that the code c gives a list, possibly infinite, of equivalences:

$$\begin{aligned}
(\mathcal{R}_0\,\mathbf{a}\,t\,\mathbf{0}) \approx \mathbf{a} &\approx \mathbf{n}_0 \\
(\mathcal{R}_0\,\mathbf{a}\,t\,\mathbf{1}) \approx (t\,\mathbf{n}_0\,\mathbf{0}) &\approx \mathbf{n}_1 \\
(\mathcal{R}_0\,\mathbf{a}\,t\,\mathbf{2}) \approx (t\,\mathbf{n}_1\,\mathbf{1}) &\approx \mathbf{n}_2 \\
&\cdots \\
(\mathcal{R}_0\,\mathbf{a}\,t\,(\mathbf{b}+\mathbf{1})) \approx (t\,\mathbf{n}_b\,\mathbf{b}) &\approx \mathbf{n}_{b+1}.
\end{aligned}$$

Let $\mathbf{n}$ be $\mathbf{n}_{b+1}$. (The reason the list may be infinite is that the model $\mathcal{M}$'s collection $\mathbb{N}$ of all type-0 elements may include non-standard integers. This is why we need to perform induction in $\mathcal{M}$ rather than expanding the term $(\mathcal{R}_0\,\mathbf{a}\,t\,(\mathbf{b}+\mathbf{1}))$ in $\mathbb{N}$—if we try to expand we may end up with an infinitely long string, which is no a valid term.)

Finally, we apply arithmetical induction in $\mathcal{M}$ to prove that $\Psi[(b+1)/z, n/y]$ holds for some $n \in \mathcal{M}$:

$$\mathcal{M} \models \exists! y(\Psi(0, y)),$$

since $y = a$ works, and:

$$\mathcal{M} \models \forall z(\exists! y(\Psi(z, y)) \rightarrow \exists! y(\Psi(z+1, y))),$$

since we can just append the type-0 element of $\mathcal{M}$ corresponding to "$(t\,(\mathbf{c})_{\mathbf{z}}\,\mathbf{z})$" to the code c. Since $\mathcal{M}$ satisfies ACA, it satisfies the schema of arithmetical induction, and hence:

$$\mathcal{M} \models \forall z(\exists! y(\Psi(z, y))),$$

and we are done.

Using the formula Ψ from the last comment, we get that there is a numeral $\mathbf{q}$ satisfying equation (2.1). This completes both induction arguments: every closed type-0 term of $\mathcal{N}$ is equivalent to a numeral.

In fact, every closed type-0 term of $\mathcal{N}$ is equivalent to a unique numeral. Suppose $\mathbf{m} \approx \mathbf{n}$ via some finite sequence $r_0, \dots, r_k$ of closed terms, where r_0 is $\mathbf{m}$, r_k is $\mathbf{n}$, and r_{j+1} results from r_j by a single application of one of the equivalence rules. To each of the closed terms r_j corresponds a formula Φ_j, defined above, that picks out a unique numeral. The formulas Φ_0 and Φ_k are "$m = y$" and "$n = y$," respectively.

We will show that for each j, $\mathcal{M} \models \forall y(\Phi_j(y) \leftrightarrow \Phi_{j+1}(y))$. By induction (in the metatheory), this shows that $m = n$, so $\mathbf{m}$ is $\mathbf{n}$. We have several cases; to avoid repeating "or vice versa," we note that all of the cases are symmetric in r_j and r_{j+1}:

- r_j and r_{j+1} are related by syntactic reduction. Recall that in defining Φ we assumed that its closed type-0 term was written in normal form; this means that Φ_j and Φ_{j+1} are the same.

- r_j is the closed term $(\mathbf{f}\,\mathbf{a})$ and r_{j+1} is the numeral $\mathbf{f}(\mathbf{a})$. Then Φ_j is "$\exists z(a = z \wedge f(z) = y)$" and Φ_{j+1} is "$f(a) = y$"; these two formulas are equivalent in $\mathcal{M}$.

- r_j is $(S\mathbf{a})$ and r_{j+1} is $(\mathbf{n}+\mathbf{1})$. Then Φ_j is "$\exists z(a = z \wedge z + 1 = y)$" and Φ_{j+1} is "$z + 1 = y$"; these two formulas are equivalent in $\mathcal{M}$.

- r_j is (μt), with formula Φ corresponding to t, and r_{j+1} is either $\mathbf{0}$ or $(\mathbf{p}+\mathbf{1})$. Then Φ_j is the formula from equation (2.5) and Φ_{j+1} is either "$0 = y$" or "$p + 1 = y$"—one of these two is equivalent to Φ_j.

- r_j is $(\mathcal{R}_0\,\mathbf{a}\,t\,\mathbf{0})$ and r_{j+1} is $\mathbf{a}$, or r_j is $(\mathcal{R}_0\,\mathbf{a}\,t\,(\mathbf{b}+\mathbf{1}))$ and r_{j+1} is some $\mathbf{p}$. In the first case, Φ_j and Φ_{j+1} are both "$a = y$"; in the second case Φ_j is $\Psi(b+1, y)$ and Φ_{j+1} is "$p = y$." In both cases, Φ_j is equivalent to Φ_{j+1}.

  (In other words, if r_j and r_{j+1} are related by a semantic type-0 rule, then the formula Φ defined above shows that r_j and r_{j+1} refer to the same numeral.)

- r_j and r_{j+1} are related because subterms of r_j and r_{j+1} are extensionally equivalent—then by extensionality, Φ_j and Φ_{j+1} are equivalent.

So every closed type-0 term in $\mathcal{N}$ is equivalent to a unique numeral; this shows that the first-order part of $\mathcal{N}$ is isomorphic to the first-order part of $\mathcal{M}$. As a consequence, the implication in equation (2.3) is actually a logical equivalence:

$$(\mathcal{M} \models \Phi[n_1/x_1, \ldots, n_k/x_k, m/y]) \iff (s[\overline{\mathbf{n}}/\overline{x}] \equiv \mathbf{m}).$$

It follows immediately that every closed type-1 term of $\mathcal{N}$ is equivalent to at most one type-1 constant symbol. (In other words, closed type-1 terms are well-defined functions.) To see that to each closed type-1 term t there is an equivalent type-1 constant symbol, construct the formula Φ corresponding to the type-0 term (tx). Then:

$$\{\langle n, m\rangle \in \mathcal{M} : \Phi[n/x, m/y]\}$$

is (in $\mathcal{M}$) the graph of the closed term t. Since Φ is arithmetical and $\mathcal{M}$ satisfies ACA, $\mathcal{M}$ contains a type-1 function f with exactly this graph. So $t \equiv \mathbf{f}$, and the second-order part of $\mathcal{N}$ is isomorphic to $\mathcal{M}$.

It remains for us to show that $\mathcal{N} \models \mathrm{RCA}_0^\omega + (\mathcal{E}_1)$. For this, we need to check that the combinator axioms are satisfied, that the axioms defining μ, $\mathcal{R}_0$, and S are satisfied, and that QF-IA still holds. That the combinator axioms are satisfied follows from the way we constructed our term model—we built in rules for λ-abstraction and application. The axioms defining μ, $\mathcal{R}_0$, and S are satisfied because of their corresponding equivalence rules and the fact that every closed type-0 term is equivalent to a numeral. (So the equivalence rules apply to all terms in $\mathcal{N}$.)

It remains only for us to check that QF-IA still holds. But this is true because if Φ is a quantifier-free formulas with type-0 variable n free, then without loss of generality Φ is equivalent to $(\mathbf{f}n) =_0 1$, where $\mathbf{f}$ is some type-1 function constant from $\mathcal{M}$. (Note that Φ may have parameters other than n, on which the choice of $\mathbf{f}$ may depend.) So QF-IA for $\mathcal{N}$ follows from QF-IA for $\mathcal{M}$, and $\mathcal{N} \models \mathrm{RCA}_0^\omega + (\mathcal{E}_1)$. □

A similar result, proved by a similar argument, holds for $(\mathcal{E}_2)$:

Corollary 2.6. $\mathrm{RCA}_0^\omega + (\mathcal{E}_2)$ *is conservative over and implies* $\Pi^1_\infty\text{-CA}_0$*.*

Traditional reverse mathematics studies subsystems of second-order arithmetic, where by "second-order arithmetic" one typically means the theory $\Pi^1_\infty\text{-CA}_0$. This corollary points out that the second-order consequences of $\mathrm{RCA}_0^\omega + (\mathcal{E}_2)$ are strictly stronger than the theories studied in traditional reverse mathematics.

Proof. The proof is essentially the same as the proof of Theorem 2.5, except that we also include a constant symbol for E_2. (We include μ as well, to avoid having to address QF-AC^{1,0}.) The rule for E_2 is, of course:

$$(E_2F^2) =_0 1 \iff \exists x^1((Fx) \neq_0 0).$$

The rest of the proof is the same, except that we add an additional equivalence rule, for E_2, and that, because of E_2, a term t may now also contain variables of type 1. The formula Φ corresponding to a type-0 term with free type-0 and type-1 variables is no longer arithmetical but allows type-1 quantification. We use the fact that $\mathcal{M}$ satisfies Π^1_∞-CA to handle $\mathcal{R}_0$ (because the formula Φ used for $\mathcal{R}_0$ may be a Π^1_∞ formula). $\square$

The proof of Theorem 2.5 implies the following corollary:

Corollary 2.7. $\mathrm{RCA}_0^\omega + (\mathcal{E}_1) + \text{QF-AC}^{0,1}$ *is conservative over and implies* $\Sigma^1_1\text{-AC}_0$*.*

Proof. Clearly QF-AC^{0,1} implies Σ^1_1-AC, since the former is simply a higher-order generalization of the latter. The theory $\Sigma^1_1\text{-AC}_0$ also includes ACA, which is implied by $(\mathcal{E}_1)$. (See [Sim99], p. 297 for a definition of $\Sigma^1_1\text{-AC}_0$.)

Let $\mathcal{M}$ be a model of $\Sigma^1_1\text{-AC}_0$; construct a term model $\mathcal{N}$ from $\mathcal{M}$ exactly as in the proof of Theorem 2.5. Then $\mathcal{N}$ satisfies $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$, since $\mathcal{M}$ satisfies ACA_0. It remains only for us to show that $\mathcal{N}$ satisfies QF-AC^{0,1} (a.k.a. Weak Countable Choice).

Let Φ be a quantifier-free formula in the language of $\mathcal{N}$ and suppose for every $n \in \mathbb{N}$ there is a type-1 function f such that $\Phi(n, f)$ holds. We will show that there is a type-$(0 \to 1)$ functional F such that for all n, $\Phi(n, F(n))$ holds.

As in the proof of Theorem 2.5 we note that for any assignment to Φ's other parameters, $\Phi(n, f)$ is equivalent to an arithmetical formula Φ' in the language of $\mathcal{M}$. Since $\mathcal{M}$ satisfies Σ^1_1-AC, there is in $\mathcal{M}$ a type-1 function g such that for all n, $\Phi'(n, g_n)$ holds, where $g_n(x) = g(\langle x, n\rangle)$ using the standard pairing function $\langle \cdot, \cdot \rangle$ on $\mathbb{N}$.

Define F by the rule:

$$F(n)(x) = g_n(x) = g(\langle x, n\rangle),$$

and we are done. □

This last corollary is interesting because while QF-AC^{0,1} is a version of Weak Countable Choice, the second-order axiom Σ^1_1-AC follows from a comprehension axiom. So Σ^1_1-AC holds, for example, in any model of ATR$_0$, while QF-AC^{0,1} is a true choice principle and is not implied by any comprehension axiom.

There is a natural uniform version of ATR, which states that for every type-1 countable well-order W there is a corresponding transfinite-recursion functional $\mathcal{R}_W$ defined by:

$$(\mathcal{R}_W\, s^1\, t^{1\to 1}\, n^0) :\equiv (t\, \langle (\mathcal{R}_W\, s\, t\, m^0) : m <_W n\rangle),$$

where $\langle x_m : m <_W n\rangle$ is the result of encoding (in a computable way) a countable list of type-1 objects as a single type-1 object. In other words, the definition says that $(\mathcal{R}_W\, s\, t\, n)$ is the result of iterating the functional t along the well-order W, through all elements $\leq_W n$.

Let "(Uniform ATR)" be the axiom asserting the existence of $\mathcal{R}_W$ for every well-order W. Then we have:

Corollary 2.8. RCA$^\omega_0$ + (Uniform ATR) *is conservative over and implies* ATR$_0$.

Proof. We classify this theorem as a corollary, because many of the techniques from Theorem 2.5 apply here as well—we just add constant symbols for each $\mathcal{R}_W$ corresponding to a well-order W in $\mathcal{M}$.

The only wrinkle is that the axiom schema ATR is stated to allow only one arithmetical formula at a time to be iterated along a given well-ordering, whereas the functional $\mathcal{R}_W$ can be nested. This does not pose a problem, since it is easy to show in RCA$^\omega_0$ that the products and sums of well-orders are also well-orders. This fact allows us to collapse multiple $\mathcal{R}_W$'s occurring in a term to a single, outermost, $\mathcal{R}_W$, at which point we have a correspondence between a closed term in our term model $\mathcal{N}$ and a formula, in the language of the base model $\mathcal{M}$, that fits the ATR schema—so we can apply ATR. □

Note that RCA_0^ω + (Uniform ATR) does not imply Weak Countable Choice—any model of ZF in which Weak Countable Choice fails satisfies the former but not the latter. So while ATR_0 is strictly stronger than Σ^1_1-AC_0, the corresponding (uniform) higher-order theories are independent.

As the next section shows, RCA_0^ω + ($\mathcal{E}_2$) seems to be the minimal theory required to do much topology on the set of type-1 objects. (Without ($\mathcal{E}_2$), for example, you cannot prove the existence of a finite family of sets.) This theory has strong second-order consequences.

2.3 Some reverse-mathematical results in RCA_0^ω and related theories

2.3.1 Type-$(1 \to 1)$ functions vs. graphs of type-$(1 \to 1)$ functions

In reverse mathematics, using one definition instead of another has certain consequences. In [Koh05], Kohlenbach follows the traditional reverse-mathematical definition of $\mathbb{R}$ as the space of all Cauchy sequences of rational numbers whose nth and $(n+1)$st elements are within 2^{-n} of each other. (See, e.g., [Sim99].) He then defines functions on $\mathbb{R}$ to be type-$(1 \to 1)$ objects that respect the natural equivalence relation $\equiv_\mathbb{R}$ on (type-1 codes for) elements of $\mathbb{R}$.

A function, in this sense, is an object g, implementing a rule, that given input x yields output $g(x)$. Consider the space ${}^\mathbb{N}\mathbb{N}$ of all type-1 elements—that is, $\mathbb{R}$ without the equivalence relation $\equiv_\mathbb{R}$. One might ask what would happen if one defined a function on ${}^\mathbb{N}\mathbb{N}$ to be, instead of a type-$(1 \to 1)$ rule g, a type-$(1 \to 0)$ graph G:

$$G := \{\langle x, y\rangle : x \in {}^\mathbb{N}\mathbb{N}\},$$

such that:

$$\forall x\, \exists! y\, (\langle x, y\rangle \in G).$$

Here we assume that the operator $\langle \cdot, \cdot \rangle$ is a suitable pairing operator for type-1 objects, such as:

$$\langle x, y \rangle :=_1 \left(\lambda n^0. \begin{cases} (xk), & \text{if } n =_0 2k \\ (yk), & \text{if } n =_0 2k+1. \end{cases} \right)$$

So one can define a function to be either a rule or a graph; as Kohlenbach notes in [Koh05], over RCA_0^ω there is no difference between these two ways of defining functions on $\mathbb{N}$, the space of type-0 objects. Given a type-$(0 \to 0)$ graph G, one can define a type-$(0 \to 0)$ rule g, recursively, by:

$$g(m) = \text{the least } n \text{ such that } \langle m, n \rangle \in G.$$

By $\text{QF-AC}^{1,0}$ (in fact, $\text{QF-AC}^{0,0}$ suffices), if G exists then g exists as well. (Here we assume that the operator $\langle \cdot, \cdot \rangle$ is a suitable pairing operator for type-0 objects.)

Conversely, one can define the graph G of a rule g by:

$$\langle m, n \rangle \in G \iff g(m) =_0 n.$$

So graphs and rules on $\mathbb{N}$ are equivalent. However, graphs and rules on ${}^{\mathbb{N}}\mathbb{N}$ are not necessarily equivalent over RCA_0^ω.

First, the existence of any graph G implies $(\mathcal{E}_1)$ over RCA_0^ω, since one can use G to define E_1 by first defining a type-2 functional F:

$$(Fy) :=_0 (G \langle x_0, y \rangle),$$

where x_0 is some fixed type-1 object, and then using F to define E_1:

$$(E_1 z) := (\neg F(\lambda n^0.((zn) + (yn)))),$$

where "$\neg$" is the (type-1) boolean negation operator. This means that RCA_0^ω cannot prove the existence of a graph on ${}^{\mathbb{N}}\mathbb{N}$, whereas RCA_0^ω proves the existence of type-$(1 \to 1)$ rules for certain continuous functions, including the identity:

$$g(x) :=_1 (\lambda n.(xn)).$$

Conversely, $(\mathcal{E}_1)$ proves that every type-$(1 \to 1)$ function rule g has a corresponding type-2 graph G. Using E_1, we can define G by:

$$\langle x, y\rangle \in G \iff g(x) =_1 y,$$

where we use E_1 to define a functional that determines whether two type-1 objects are equal.

Going from a graph to a rule is more complicated. First, any model that does not satisfy $(\mathcal{E}_1)$ contains no graphs; in such models it is vacuously true that every graph has a corresponding rule. So "every graph has a corresponding rule" does not imply $(\mathcal{E}_1)$. Second, there is a model of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$ that contains a graph with no corresponding function; so $(\mathcal{E}_1)$ does not imply "every graph has a corresponding rule."

To see this second fact we use a bit of computability theory. Start with the second-order structure (ω, ARITH) whose second-order part consists of all functions corresponding to arithmetical sets; this structure satisfies ACA_0. (It is, in fact, the minimal ω-model satisfying ACA_0; see [Sim99].) Apply the proof of Theorem 2.5 to get a minimal term model $\mathcal{N}$ of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$. Note that for each type-1 set $X \subseteq \mathbb{N}$ in the model, the model also contains all of X's finite Turing jumps $X^{(n)}$, for $n \in \omega$—this follows because the model satisfies ACA_0 and hence is closed under arithmetical definability. In particular, $\mathcal{N}$ contains $\emptyset, \emptyset', \ldots, \emptyset^{(n)}, \ldots$.

However, $\mathcal{N}$ does not contain the ω-jump $\emptyset^{(\omega)}$, defined by:

$$\langle m, n\rangle \in \emptyset^{(\omega)} \iff m \in \emptyset^{(n)},$$

since $\emptyset^{(\omega)}$ is a type-1 object but not an arithmetical set, and the second-order part of $\mathcal{N}$ consists only of arithmetical sets.

From E_1 one can define the type-2 functional G:

$$(G\,\langle n^0, x^1, y^1\rangle) :=_0 1 \iff y \text{ is the join of the first } n \text{ Turing jumps of } x.$$

Note that one can define, using E_1, the Turing jump operator $X \mapsto X'$ as:

$$X' :=_1 (\lambda e^0.(E_1(\lambda s^0.\Phi(e, X, s)))),$$

where Φ is a primitive-recursive functional that tells whether the eth Turing machine, when given oracle X, halts within s steps on input e. One can then define G from the Turing jump operator by primitive recursion on the code y's index n. So G exists in any model of $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$.

Note that if $\langle x, y\rangle \in G$ and $\langle x, y'\rangle \in G$ then $y =_1 y'$. For G to be a graph, it must also contain, for every x, some $\langle x, y\rangle$; in an arbitrary model, it may not. However, since $\mathcal{N}$ contains exactly the arithmetical sets, G is a graph:

$$\mathcal{N} \models \forall n^0 \, \forall x^1 \, \exists! y^1 (\langle n, x, y\rangle \in G),$$

Now suppose, for a contradiction, that $\mathcal{N}$ contains a type-$(1 \to 1)$ rule g for G. From that rule g one can define the ω-jump:

$$\langle m, n\rangle \in \emptyset^{(\omega)} \iff ((g\, n\, \emptyset)m) =_0 1,$$

a contradiction.

The axiom $(\mathcal{E}_2)$ implies that every graph G on ${}^{\mathbb{N}}\mathbb{N}$ has a corresponding rule g—define g by letting $((gx)\, m)$ be least such that for some type-1 $y \supset \langle ((gx)\, 0), \ldots, ((gx)\, m)\rangle$, we have $\langle x, y\rangle \in G$. The functional $\mathcal{E}_2$ lets us determine whether there is such a y, and thus allows us to make this definition.

At the same time, the choice principle $\text{QF-AC}^{1,1}$ clearly suffices, over $(\mathcal{E}_1)$, to prove that every graph G on ${}^{\mathbb{N}}\mathbb{N}$ has a corresponding rule g, since:

$$\forall x^1 \, \exists y^1 (\langle x, y\rangle \in G).$$

In fact, this y is unique, which is why the existence of g can also be proved from the comprehension principle $(\mathcal{E}_2)$.

The statement "every graph has a corresponding rule" is not equivalent to $\text{QF-AC}^{1,1}$, since $(\mathcal{E}_2)$ does not even imply Weak Countable Choice, $\text{QF-AC}^{0,1}$.

Further, the conjunction of the statements "every graph has a corresponding rule" and "a graph exists" is not equivalent to $(\mathcal{E}_2)$. One can see this by a variation of Proposition 2.17, or by the following proposition:

Proposition 2.9 ($\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$). *The conjunction of the statements "every graph has a corresponding rule" and "a graph exists" is conservative over the second-order theory Σ^1_1-AC_0.*

Notice that we do not say that these statements imply Σ^1_1-AC_0—we say only that they are conservative over Σ^1_1-AC_0, so their second-order consequences are strictly weaker than the second-order consequences of $(\mathcal{E}_2)$.

Note that our proof relies on the fact that the choice of $y =_1 g(x)$ in a graph G is unique; we do not claim that (and have not determined whether) the higher-order schema $\mathrm{QF}\text{-}\mathrm{AC}^{1,1}$ is conservative over Σ^1_1-AC_0.

Proof. Let $\mathcal{M}$ be a (second-order) model of Σ^1_1-AC_0; construct a term model $\mathcal{N}$ from $\mathcal{M}$ as in the proof of Theorem 2.5, while also including, for every (term for a) graph G in $\mathcal{N}$, a corresponding type-$(1 \to 1)$ constant symbol g. Since each term is finite, after adding constant symbols g for rules at ω-many stages, the model $\mathcal{N}$ will have stabilized and will satisfy "every graph has a corresponding rule."

As in the proof of Theorem 2.5, $\mathcal{N}$ also satisfies $\mathrm{RCA}_0^\omega + (\mathcal{E}_1)$, so it contains at least one graph. It remains only for us to show that the type-0 and type-1 parts of $\mathcal{N}$ are isomorphic to the corresponding parts of $\mathcal{M}$. The only difference from the proof of Theorem 2.5 is that a closed type-0 term t may contain one or more constant symbols g corresponding to graphs.

This means that for subterm $s[\overline{x}]$ of t, we must consider an additional case:

$$(g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]),$$

where g is one of the new type-$(1 \to 1)$ functionals we added to $\mathcal{N}$. The equivalence rule is:

$$(g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]) =_1 y,$$

for the unique y such that $\langle x, y\rangle \in G$, the graph corresponding to g. Note that the definition of G may itself refer to a rule h, so we also induct on the rank of the rule g.

Note also that this reduction rule relates type-1 terms, rather than type-0 terms. To get the rule for type-0 terms, we just consider:

$$((g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]))\, s'[\overline{\mathbf{m}}/\overline{x}]),$$

where s' is another type-0 subterm of t.

The key point is that the Σ^1_1-AC schema proves that there is a type-1 sequence $\overline{\mathbf{m}} \mapsto y_{\overline{\mathbf{m}}}$ such that, for all $\overline{\mathbf{m}}$:

$$\langle(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]), y_{\overline{\mathbf{m}}}\rangle \in G.$$

Since G is a graph, this $y_{\overline{\mathbf{m}}}$ is unique, so $y_{\overline{\mathbf{m}}} =_1 (g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]))$. So $\mathcal{M}$ contains a type-1 sequence:

$$\overline{\mathbf{m}} \mapsto y_{\overline{\mathbf{m}}} =_1 (g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k])).$$

The formula Φ for $((g(\lambda x_1.s[\mathbf{m}_2/x_2, \ldots, \mathbf{m}_k/x_k]))\, s'[\overline{\mathbf{m}}/\overline{x}])$, assuming we have formula Ψ' for s', is just:

$$\exists z^0(\Psi'(\overline{x}, z) \wedge (y_{\overline{\mathbf{m}}}\, z) =_0 y).$$

The rest of the proof is as before. $\square$

2.3.2 Set cardinalities

Definition 2.10. *A type-$(\sigma \to 0)$ **set of elements of type-**σ is a type-$(\sigma \to 0)$ functional X such that for all type-σ objects x, $(Xx) =_0 0 \vee (Xx) =_0 1$.*

In other words, we call a characteristic function a set. We sometimes write "$x \in X$" for "$(Xx) =_0 1$."

Definition 2.11. *A type-$(\sigma \to 0)$ set X of type-σ elements is **countable** if and only if there is a type-$(0 \to \sigma)$ enumeration $m \mapsto x_m$ of its contents:*

$$x \in X \iff \exists m\, (x =_\sigma x_m).$$

*A type-$(\sigma \to 0)$ set X of type-σ elements has **cardinality** $\leq \beth_n$ if and only if there is a type-$(n \to \sigma)$ enumeration $f \mapsto x_f$ of its contents:*

$$x \in X \iff \exists f^n\, (x =_\sigma x_f).$$

*A type-$(\sigma \to 0)$ set X of type-σ elements is **finite** if and only if it is countable and there exists an N such that:*

$$\forall k > N\, \exists m \leq N (x_k =_\sigma x_m).$$

In the above definition the enumerations need not be injective.

Proposition 2.12 (RCA_0^ω)**.** *Every non-empty type-$(n \to 0)$ set X of type-n elements has cardinality $\leq \beth_n$.*

Proof. Fix $x_0 \in X$. If $n = 0$ then define the type-$(0 \to 0)$ enumeration $m \mapsto x_m$ of X:

$$x_m :=_0 \begin{cases} m, & \text{if } m \in X \\ x_0, & \text{otherwise.} \end{cases}$$

If $n > 0$ then define the type-$(n \to n)$ enumeration $f \mapsto x_f$:

$$x_f :=_n \left(\lambda y^{n-1}. \begin{cases} (fy), & \text{if } f \in X \\ (x_0 y), & \text{otherwise.} \end{cases} \right).$$

□

Proposition 2.13 (RCA_0^ω)**.** *For all n, $\beth_{n+1} \not\leq \beth_n$.*

Proof. Suppose, for a contradiction, that there is a type-$(n \to (n+1))$ enumeration $f \mapsto x_f$ of all the type-$(n+1)$ elements. We can use diagonalization to define a new type-$(n+1)$ element x:

$$x :=_{n+1} (\lambda f^n.(x_f\, f) + 1).$$

Then for every type-n object f, $x \neq_{n+1} x_f$ since the two disagree on input f. This is a contradiction. □

Proposition 2.14 (RCA_0^ω)**.** $\beth_0 \leq \beth_1 \leq \beth_2$.

Proof. The proof is similar to that of Proposition 2.4. For $\beth_0 \leq \beth_1$, use the type-$(1 \to 0)$ enumeration:

$$f^1 \mapsto (f0).$$

This map is surjective since it maps the constant-k function $(\lambda m^0.k^0)$ to k.

For $\beth_1 \leq \beth_2$, use use the type-$(2 \to 1)$ enumeration:

$$F^2 \mapsto (\lambda m^0.(F(\lambda k^0.m))).$$

This map is surjective since it maps the type-2 functional $(\lambda x^1.(f(x0)))$ to f. $\square$

2.3.3 Consequences of the existence of small type-2 sets and type-3 families

Using the above definitions we have:

Proposition 2.15 (RCA_0^ω). *1. The existence of a countable type-2 set of type-1 objects is equivalent to* $(\mathcal{E}_1)$.

2. The existence of a type-3 set of type-2 objects with cardinality $\leq \beth_1$ *is equivalent to* $(\mathcal{E}_2)$.

Proof. **1.** Note that E_1 is the singleton set consisting of the function $(\lambda m.0)$, which is countable via the constant enumeration $n \mapsto (\lambda m.0)$. To see the converse, let X be a type-2 set and let $m \mapsto x_m$ be a countable enumeration of X. We first define a type-$(1 \to 1)$ functional g by the rule:

$$g(x) :=_1 \left(\lambda\langle m,n\rangle. \begin{cases} (x_0 \langle m,n\rangle), & \text{if } (xn) =_0 0 \\ (x_m \langle m,n\rangle) + 1, & \text{otherwise.} \end{cases} \right)$$

If $(xn) =_0 0$, for all n, then $g(x) =_1 x_0$. Suppose that for some n, $(xn) \neq_0 0$; then for every m we have that $g(x) \neq_1 x_m$, since:

$$((gx)\langle m,n\rangle) =_0 (x_m \langle m,n\rangle) + 1 \neq_0 (x_m \langle m,n\rangle).$$

Define E_1 to be:

$$E_1 :=_2 (\lambda x.(X(gx))).$$

2. Note that E_2 is the singleton set consisting of the functional $(\lambda x.0)$, which trivially has cardinality $\leq \beth_1$.. For the converse, let $\mathcal{F}$ be a type-3 set and let $f \mapsto X_f$ be an enumeration of $\mathcal{F}$, of cardinality $\leq \beth_1$. This case is similar to the previous case; define:

$$g(X) :=_2 \left(\lambda \langle f, g \rangle . \begin{cases} (X_0 \langle f, g \rangle), & \text{if } (Xg) =_0 0 \\ (X_f \langle f, g \rangle) + 1, & \text{otherwise,} \end{cases} \right)$$

and define E_2 to be:

$$E_2 :=_3 (\lambda X.(\mathcal{F}(gX))).$$

□

The consequences of Proposition 2.15 for studying topology in RCA_0^ω are significant. Since the topology of countable spaces tends to be simple, we would want to consider topological spaces consisting of objects of at least type 1 (if not higher)—that is, topologies on the set ${}^{\mathbb{N}}\mathbb{N}$. Proposition 2.15 says that the existence of any countable family of subsets of ${}^{\mathbb{N}}\mathbb{N}$ implies $(\mathcal{E}_2)$. In particular, if, in some model $\mathcal{M}$ of RCA_0^ω, a topology has a countable basis then it would seem that $\mathcal{M} \models (\mathcal{E}_2)$.

As noted in Corollary 2.6, the second-order consequences of $(\mathcal{E}_2)$ include full second-order comprehension. This means that the axiom $(\mathcal{E}_2)$ is quite strong.

Note that Proposition 2.15 applies only when a sequence has a corresponding type-2 or type-3 set. If one's "countable basis" is in fact just a countable enumeration of sets, then the existence of the countable basis does not necessarily imply $(\mathcal{E}_2)$. However, the topology generated by a countable enumeration of a basis would (under the usual definition) have a type-$(1 \to 2)$ enumeration—so Proposition 2.15 would still apply and we would have $(\mathcal{E}_2)$.

In fact, as discussed in Section 2.3.5, many seemingly trivial topological statements, when applied to the set ${}^{\mathbb{N}}\mathbb{N}$, are equivalent to $(\mathcal{E}_2)$. The set ${}^{\mathbb{N}}\mathbb{N}$ is very flexible, and to avoid this flexibility and try to capture the essence of reverse topology we introduce a new base theory in Chapter 3.

2.3.4 What topologies does one get from $\mathbf{RCA}_0^\omega$ + $(\mathcal{E}_1)$?

Definition 2.16. *A **topology** (on $^{\mathbb{N}}\mathbb{N}$) is a type-3 family $\mathcal{T}$ of type-2 sets satisfying:*

- $\emptyset \in \mathcal{T}$. *(Here we define $\emptyset :=_2 (\lambda x^1.0)$.)*
- $^{\mathbb{N}}\mathbb{N} \in \mathcal{T}$. *(Here we define $^{\mathbb{N}}\mathbb{N} :=_2 (\lambda x^1.1)$.)*
- *If $X, Y \in \mathcal{T}$ then $X \cap Y \in \mathcal{T}$. (Here we define $X \cap Y :=_2 (\lambda x^1.((Xx) \cdot (Yx)))$.)*
- *If $\mathcal{F} \subseteq \mathcal{T}$ is a type-3 family of type-2 sets, then $\bigcup \mathcal{F}$, defined to be:*

$$\bigcup \mathcal{F} :=_2 \left(\lambda x^1. \begin{cases} 1, & \text{if } \exists U^2 \in \mathcal{F}\,(x \in U) \\ 0, & \text{otherwise,} \end{cases} \right)$$

  exists and is in $\mathcal{T}$. (Alternatively, we may drop the requirement that $\bigcup \mathcal{F}$ exist and require only that if $\bigcup \mathcal{F}$ exists, then $\bigcup \mathcal{F} \in \mathcal{T}$. The results in this paper hold for either definition.)

The axiom $(\mathcal{E}_2)$ suffices to prove the existence of $\bigcup \mathcal{F}$ for every family $\mathcal{F}$—in general, the axiom $(\mathcal{E}_\tau)$ implies Σ^τ_∞-comprehension. However, in the absence of $(\mathcal{E}_2)$ one cannot necessarily assume that $\bigcup \mathcal{F}$ exists.

We could extend the definition of a topology to include topologies on proper subsets of $^{\mathbb{N}}\mathbb{N}$. However, the existence of a topology on a proper subset of $^{\mathbb{N}}\mathbb{N}$ implies $(\mathcal{E}_2)$, so have not examined these.

The theory RCA_0^ω + $(\mathcal{E}_1)$ suffices to prove the existence of certain topologies; however, the only topologies whose existence it proves are essentially just topologies on a countable space:

Proposition 2.17. *If $\mathcal{T}$ is a topology on $^{\mathbb{N}}\mathbb{N}$ that exists in a model $\mathcal{N}$ of RCA_0^ω + $(\mathcal{E}_1)$, as constructed in Theorem 2.5, then $\mathcal{N}$ contains a set S of elements of $^{\mathbb{N}}\mathbb{N}$ and a countable enumeration $\langle x_n : n \in \mathbb{N} \rangle$ such that:*

- *S is a topology on $\mathbb{N}$ and*

- $\mathcal{T}$ *is* $\{U \subseteq \{x_n : n \in \mathbb{N}\} : \{n : x_n \in U\} \in S\} \times \mathcal{P}\left({}^{\mathbb{N}}\mathbb{N} \setminus \{x_n : n \in \mathbb{N}\}\right)$.

In other words, the topology $\mathcal{T}$ corresponds to S on the sequence $\langle x_n : n \in \mathbb{N}\rangle$ (by the correspondence $n \mapsto x_n$) and ignores ${}^{\mathbb{N}}\mathbb{N} \setminus \{x_n : n \in \mathbb{N}\}$.

Proof. Let $(\lambda Z^2.t[Z])$ be a closed term defining $\mathcal{T}$. Without loss of generality, assume that the type-0 term t is in normal form (fully reduced); then all variables occurring in subterms of t are of type 0.

Note that the variable Z occurs in subterms of t only on the left-hand side of an application:

$$(Z\, s_i[\overline{m_i}, Z]),$$

since no constant symbols in our model take inputs of type 2. Let $(Z\, s_1), \ldots, (Z\, s_k)$ list all applications in t involving the variable Z; let $\overline{m_i}$ list the type-0 variables free in s_i.

We will prove that for every set X there is a countable sequence $\langle x_n : n \in \mathbb{N}\rangle$ of type-1 points such that for all sets Y:

$$(\lambda n.(X\, x_n)) =_1 (\lambda n.(Y\, x_n)) \implies (\mathcal{T}\, X) =_0 (\mathcal{T}\, Y).$$

In other words, any set that agrees with X on the sequence $\langle x_n : n \in \mathbb{N}\rangle$ is in $\mathcal{T}$ if and only if X is.

For each parameter X and subterm $s_i[\overline{m_i}, X]$, define the sequence $\langle x_{n,i} : n \in \mathbb{N}\rangle$ by the rule:

$$x_{\overline{m_i},i} :=_1 s_i[\overline{m_i}, X].$$

Fixing X, for each $\overline{m_i}$ the subterm $s_i[\overline{m_i}, X]$ is closed, so $x_{\overline{m_i},i}$ is in our model. The definition of $x_{\overline{m_i},i}$ is uniform in $\overline{m_i}$ and can be made in RCA_0^ω, so the sequence $\langle x_{n,i} : n \in \mathbb{N}\rangle$ is in our model.

Let $\langle x_n : n \in \mathbb{N}\rangle$ be a countable sequence consisting of all elements of $\langle x_{\overline{m_1},1}\rangle, \ldots, \langle x_{\overline{m_k},k}\rangle$. (This can be done via a computable pairing function on $\mathbb{N}$, so the sequence $\langle x_n : n \in \mathbb{N}\rangle$ is in our model.)

Now suppose that the set Y agrees with X on $\langle x_n : n \in \mathbb{N}\rangle$; then we will show, by induction on subterms of t, that $(\mathcal{T}\, X) =_0 (\mathcal{T}\, Y)$. If not, then there is a subterm $(Z s_i[\overline{m_i}, Z])$ such that for some value of $\overline{m_i}$:

$$(X s_i[\overline{m_i}, X]) \neq_0 (Y s_i[\overline{m_i}, Y]).$$

Fix the innermost such s_i; since s_i is innermost, for all subterms s_j of s_i and all $\overline{m_j}$:

$$(Y s_j[\overline{m_j}, Y]) =_0 (X s_j[\overline{m_j}, X]).$$

So, for every $\overline{m_i}$, since $s_i[\overline{m_i}, Z]$ depends only on $\overline{m_i}$ and the variable Z, we have that $s_i[\overline{m_i}, Y] =_1 s_i[\overline{m_i}, X] =_1 x_{\overline{m_i},i}$. However, by assumption, X and Y agree on all x_n and hence on all $x_{\overline{m_i},i}$—a contradiction.

Now $\mathcal{T}$ is a topology so it contains both $\emptyset$ and ${}^{\mathbb{N}}\mathbb{N}$. Let $\langle z_n : n \in \mathbb{N}\rangle$ be the countable sequence for $\emptyset$, and let $\langle y_n : n \in \mathbb{N}\rangle$ be the countable sequence for ${}^{\mathbb{N}}\mathbb{N}$. Since the sequence $\langle z_n : n \in \mathbb{N}\rangle$ comes from $\emptyset$, and each $z_n \notin \emptyset$, it must be the case that any set not containing any of the z_n's is in $\mathcal{T}$. Similarly, any set containing all of the y_n's is in $\mathcal{T}$.

Let $\langle x_n : n \in \mathbb{N}\rangle$ be the sequence consisting of all y_n's and z_n's. Now if $U \in \mathcal{T}$ and V agrees with U on $\langle x_n : n \in \mathbb{N}\rangle$ then V is also in $\mathcal{T}$, since V is the union of the open set:

$$V \setminus \{x_n : n \in \mathbb{N}\}$$

(open because it excludes all z_n's), with the intersection of U and:

$$\{x_n : n \in \mathbb{N}\} \cup (V \setminus \{x_n : n \in \mathbb{N}\})$$

(open because it includes all y_n's).

It remains for us to check that $\mathcal{T}$ yields a topology S on $\mathbb{N}$. Let S be the set of all type-1 characteristic functions f such that:

$$\{x_n : (fn) =_0 1\} \in \mathcal{T}.$$

Note that $U \in \mathcal{T} \iff \{n : x_n \in U\} \in S$. To see that S is a topology on $\mathbb{N}$, note that $\mathbb{N}$ and $\emptyset$ are clearly in S, while if characteristic functions f_A and f_B are in S, then:

$$f_{A\cap B} := (\lambda m.(f_A m) \cdot (f_B m)) \in S,$$

since the type-2 sets corresponding to A and B are in $\mathcal{T}$.

Finally, suppose $F \subseteq S$ is a family of type-1 sets; we must show that $\bigcup F$ exists and is in S. Define $\mathcal{F} \subseteq \mathcal{T}$ by:

$$\mathcal{F} :=_3 \{U^2 : (\lambda m.(U\, x_m)) \in F)\};$$

this family is definable in RCA_0^ω and hence exists in $\mathcal{N}$. Since $\mathcal{T}$ is a topology, $V :=_2 \bigcup \mathcal{F}$ exists and is open. Then:

$$\bigcup F =_1 (\lambda m.\ (V\, x_m))$$

is in S. □

The second part of the proof of Proposition 2.17 relied on the fact that $\mathcal{T}$ was a topology to show that $\mathcal{T}$ had a single countable sequence that it used to test all sets. However, the first part of the proof works for all families $\mathcal{F}$: if sets X and Y agree on the sequence used by $\mathcal{F}$ to test X, then $(\mathcal{F}\, X) =_0 (\mathcal{F}\, Y)$. This fact yields another proof that $(\mathcal{E}_1)$ does not imply $(\mathcal{E}_2)$. (The first proof comes from the conservation results.) Suppose, for a contradiction, that E_2 is definable from E_1, and apply the first part of the proof of Proposition 2.17 to the type-3 functional E_2 and the type-2 parameter $\emptyset$.

Then we get a countable sequence $\langle x_n : n \in \mathbb{N} \rangle$ such that for every set X excluding all x_n's, $(E_2\, X) =_0 (E_2\, \emptyset) =_0 1$, a contradiction.

Further, the first part of proof works even if you add arbitrary type-2 (or lower) functionals to $\mathcal{N}$. (This is because the first part of the proof requires only that $\mathcal{N}$ have no functionals that take type-2 inputs.) This fact shows that no third-order axiom can imply $(\mathcal{E}_2)$, since third-order axioms do not add functionals of types higher than 2.

2.3.5 Axioms and topologies on the space of type-1 objects

Using the definitions of (type-2) sets (of type-1 objects) given above, one can start to perform reverse topology on the space ${}^{\mathbb{N}}\mathbb{N}$ of all type-1 objects. In this section we use Definition 2.16 for topologies on ${}^{\mathbb{N}}\mathbb{N}$, the space of all type-1 objects.

Proposition 2.18 (RCA_0^ω)**.** ($\mathcal{E}_1$) *is equivalent to the statement:* Given a countable dense enumeration $\langle x_n : n \in \mathbb{N}\rangle$ together with a type-$(1 \to 1)$ metric d, there is a countable enumeration $\langle B_n : n \in \mathbb{N}\rangle$ of a basis for the space.

Proof. $\Rightarrow$: Suppose we have $\langle x_n : n \in \mathbb{N}\rangle$ and f. Define the enumeration $\langle B_n : n \in \mathbb{N}\rangle$ of a basis for the space by:

$$x \in B_{\langle n,q\rangle} \iff (d \langle x, x_m\rangle) <_{\mathbb{R}} q,$$

where q is a positive rational number. The relation $<_{\mathbb{R}}$ can be defined from E_1; see [Koh05].

$\Leftarrow$: The standard metric d on $\mathbb{R}$ is definable in RCA_0^ω, by:

$$(d \langle x, y\rangle) =_1 (\lambda n^0.|(xn) - (yn)|).$$

One can define the sequence of rationals by:

$$q \mapsto (\lambda n^0.q).$$

Suppose that there is a countable enumeration $\langle B_n : n \in \mathbb{N}\rangle$ of a basis for the standard topology on $\mathbb{R}$. In particular, some B_N contains $2_{\mathbb{R}}$ and is contained within the ball of radius $1_{\mathbb{R}}$ centered at $2_{\mathbb{R}}$. Define the sequence $n \mapsto q_n$ of dyadic rationals:

$$\begin{aligned} q_0 &:= 2 \\ q_{i+1} &:= q_i - 2^{-m_{i+1}}, \end{aligned}$$

for $m_{i+1} > m_i \geq i$ least such that $(\lambda n.(q_i - 2^{-m_{i+1}}))$, a rational, is in B_N. The fact that B_N is open (in the sense of the metric d) implies that there is always some such m_{i+1}, so RCA_0^ω proves that the recursively-defined sequence $n \mapsto q_n$ exists.

Note that $x_0 :=_1 (\lambda i.q_i)$ is a real, and that for every i, the distance between $(\lambda n.q_i)$ and x_0 is no greater than 2^{-i}. The rationals q_i decrease monotonically to x_0. Further, $x_0 \notin B_N$, since if it were then B_N would also contain all reals within some ε of x_0—which would mean that we would eventually choose some q_i to the left of x_0, a contradiction.

So B_N is not everywhere $\epsilon - \delta$-continuous, and we can apply Proposition 3.12 from [Koh05]. □

Proposition 2.19 (RCA_0^ω)**.** *The following are equivalent to* ($\mathcal{E}_2$)*:*

1. *There exists a topology for a connected space. (I.e., the only clopen sets are* $\emptyset$ *and* ${}^{\mathbb{N}}\mathbb{N}$*.)*
2. *There exists a topology with a dense, nowhere-dense set.*
3. *There exists a topology generated by a countable enumeration for a basis.*

Proof. The proofs of these equivalences are all straightforward.

1. Using $\mathcal{E}_2$, one can define the indiscrete topology $\mathcal{T}$ by:

$$X \in \mathcal{T} \iff (X =_2 \emptyset) \vee (X =_2 {}^{\mathbb{N}}\mathbb{N}).$$

Conversely, from a connected topology $\mathcal{T}$ one can define E_2 by:

$$(E_2 X) =_0 0 \iff (X \in \mathcal{T}) \wedge (({}^{\mathbb{N}}\mathbb{N} \setminus X) \in \mathcal{T}) \wedge (x_0 \notin X),$$

where x_0 is some fixed element of type 1. The set ${}^{\mathbb{N}}\mathbb{N} \setminus X$ can be defined by:

$$(\lambda x^1.({}^{\mathbb{N}}\mathbb{N}x) \wedge \neg(Xx)).$$

2. Using $\mathcal{E}_2$, one can define the indiscrete topology, as well as a singleton set $\{x_0\}$; then $\{x_0\}$ is trivially dense and nowhere-dense. Conversely, suppose $\mathcal{T}$ is a topology and $D \subseteq {}^{\mathbb{N}}\mathbb{N}$ is dense and nowhere-dense. Then one can define E_2 by:

$$(E_2 X) =_0 0 \iff (X \cap D) \in \mathcal{T}.$$

3. Using $\mathcal{E}_2$, one can define both the standard topology on $\mathbb{R}$ and an enumeration of all open intervals with rational endpoints. Conversely, if $\mathcal{T}$ is generated from a countable enumeration $\langle B_n : n \in \mathbb{N}\rangle$ for a basis, then:

$$X \in \mathcal{T} \iff \exists f^1 (X =_2 \bigcup \{B_n : (fn) =_0 1\}),$$

which means that $\mathcal{T}$ has cardinality $\leq \beth_1$. By Proposition 2.15, the existence of $\mathcal{T}$ implies ($\mathcal{E}_2$).

$\square$

Proposition 2.20 (RCA_0^ω + ($\mathcal{E}_1$)). *The following are equivalent to* ($\mathcal{E}_2$)*:*

1. *A separable topology exists.*
2. *A topology of first-category exists. (I.e., the space is a countable union of nowhere-dense sets.)*
3. $\mathrm{QF\text{-}AC}^{0,1}$ + *a second-countable topology exists.*

Proof. The proofs of these equivalences are all straightforward.

1. As in the previous proof, E_2 allows one to define the indiscrete topology, which is trivially separable. Conversely, if $\mathcal{T}$ has a countable dense enumeration $\langle x_n : n \in \mathbb{N}\rangle$ then one can define E_2 by:

$$(E_2X) =_0 0 \iff (X \in \mathcal{T}) \wedge (\forall n^0(x_n \notin X)).$$

Note that one can use E_1 to handle this definition's type-0 quantifier.

2. Over ($\mathcal{E}_2$), using the indiscrete topology, ${}^{\mathbb{N}}\mathbb{N} =_2 \{x_0\} \dot\cup ({}^{\mathbb{N}}\mathbb{N} \setminus \{x_0\}$, a finite union of nowhere-dense sets. Conversely, suppose we have a topology $\mathcal{T}$ and an enumeration $\langle D_n : n \in N\rangle$ of nowhere-dense sets such that ${}^{\mathbb{N}}\mathbb{N} =_2 \bigcup_{n\in\mathbb{N}} D_n$. Then one can define E_2 by:

$$(E_2X) =_0 0 \iff \forall n^0\,(X \cap D_n) \in \mathcal{T}.$$

3. The standard topology on $\mathbb{R}$ is second-countable, over ($\mathcal{E}_2$). Conversely, suppose $\mathcal{T}$ has a countable enumeration $\langle B_n : n \in \mathbb{N}\rangle$ for a basis. Apply $\mathrm{QF\text{-}AC}^{0,1}$ to the formula:

$$\forall n^0 \exists x^1 (x \in B_n),$$

to get a countable dense enumeration $\langle x_n : n \in \mathbb{N}\rangle$. Then apply part (1), above. □

The fact that the existence of so many basic topologies implied ($\mathcal{E}_2$) led us to consider a new approach to higher-order reverse topology, which is outlined in the next chapter. Intuitively, it is tricky to do topology in a theory in which does not have a set-equality functional. That is, E_2, which allows us to define functionals that distinguish between different sets, seems like something we would want to have. However, ($\mathcal{E}_2$) implies the second-order theory Π^1_∞-CA_0, and also allows us to quantify over spaces of size $\beth_1$. The axiom $(\mathcal{E}_2)$ is very strong.

Chapter 3

Base theory RCA_0^ω + (ATOMS) and related theories

For reverse topology, we would prefer to have a set-equality functional with weaker consequences that $(\mathcal{E}_2)$. Further, we would like to be able to consider topologies of arbitrary cardinality—not just topologies of size $\beth_1$. To that end, we apply an idea presented by Victor Harnik in [Har87] and extend the base theory RCA_0^ω by throwing in a new atomic type, α. Elements of our topological space will be of type α; sets and families of sets will be of appropriate higher types.

3.1 Definitions

3.1.1 Finite types

Our language is as before, with the addition of new finite types:

- 0 is an atomic type, as before.
- α is an atomic type, the type of "atoms."
- If σ and τ are finite types, then $(\sigma \to \tau)$ is a finite type as well.

As before, the finite types are defined inductively from the atomic types, now including α as well as 0. We abuse notation somewhat by writing "$(\alpha+1)$" for the type $(\alpha \to 0)$, "$(\alpha+2)$" for the type $((\alpha+1) \to 0)$, and so on.

With respect to topological statements, our elements or points will be of type α; our sets, open or not, will be of type $(\alpha+1)$; and our topologies will be of type $(\alpha+2)$.

If τ is a finite type made up only of "0" and "$\rightarrow$," without "α," then we say that τ is a **numerical** type.

As before, our language consists of all variables of finite types, and for each variable type we have $\exists$ and $\forall$ quantifiers of the corresponding types.

3.1.2 The base theory $\mathbf{RCA}_0^\omega$ + (ATOMS) and stronger theories

We define a new theory "RCA_0^ω + (ATOMS)," which is conservative over and implies RCA_0^ω. The axioms of our new theory are the same as those of the old, with a few exceptions:

- The axioms for the combinators Σ and Π are extended to allow the definition of (semantic) λ-abstraction for all finite types—not just those types built up from 0.

- In addition to the axioms defining type-0 equality, $=_0$, we add axioms defining type-α equality. Higher-order equality is treated extensionally, as before.

- We add a new axiom ($\mathcal{A}_0$) asserting the existence of the type-$((\alpha,\alpha) \rightarrow 0)$ functional A_0 (the type-α equality relation), as follows:

  $$\forall x^\alpha \forall y^\alpha (x =_\alpha y \leftrightarrow ((A_0\, x, y) =_0 1).$$

  Note that no analogue to A_0 is required for type-0 terms, since:

  $$m =_0 n \leftrightarrow (1 \dot{-} [(m \dot{-} n) + (n \dot{-} m)] =_0 1.$$

  The boolean values 0 and 1 are already type-0 terms.

- We add new axioms defining a type-$((\alpha,\alpha,0) \rightarrow \alpha)$ functional $\mathbf{P}$, allowing us to build finite, type-0-indexed lists:

  $$\forall x^\alpha \forall y^\alpha \forall n^0 [(n =_0 0 \rightarrow (\mathbf{P}\, x\, y\, n) =_\alpha x) \wedge (n \neq_0 0 \rightarrow (\mathbf{P}\, x\, y\, n) =_\alpha y)].$$

  Intuitively, $\mathbf{P}$ allows for functionals that return different atoms based on a boolean conditional; note that no analogue to $\mathbf{P}$ is required for type-0 terms, since pairing for type-0 objects can be recursively defined.

- We add (infinitely many) new axioms stating that there are infinitely many distinct atoms.

This is the base theory, RCA_0^ω + (ATOMS). We have analogues A_n, where $n > 0$, to the equality functions E_n for type-n elements, whose existence is asserted by corresponding axioms $(\mathcal{A}_n)$:

$$(\mathcal{A}_n) : \exists A_n^{(\alpha+n)\to 0} \forall X^{\alpha+n} (\exists y^\alpha (X(y) \neq_0 0)) \leftrightarrow (A_n(X) =_0 1),$$

(Note that, for reasons already discussed, A_0 is of type $((\alpha,\alpha) \to 0)$ rather than type $(\alpha \to 0)$.) As with E_n one can (in RCA_0^ω + (ATOMS)) use A_n to define the characteristic function of the type-$(\alpha+n)$ equality relation.

3.2 Conservation results for $\mathbf{RCA}_0^\omega$ + (ATOMS) and related theories

Proposition 3.1. *The theory* RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ *is conservative over* RCA_0^ω*.*

This means that the interaction between the $(\mathcal{A}_1)$ axiom and the numerical part of the theory RCA_0^ω + (ATOMS) is limited.

Proof. Let $\mathcal{M}$ be an arbitrary model of RCA_0^ω. Construct a new term model, $\mathcal{N}$ from $\mathcal{M}$ by first including constant symbols (of the appropriate types) for all elements of $\mathcal{M}$. Also include type-α constants for each element of some infinite set of size $\mathbb{N}$: that is, include a type-α constant corresponding to each type-0 constant. (We want $\mathbb{A}$ to have the same cardinality as $\mathbb{N}$, rather than just cardinality ω, to ensure that the resulting model $\mathcal{N}$ satisfies the quantifier-free induction schema.)

Let $\mathcal{N}$ consist of all equivalence classes of (valid) closed terms using these constant symbols as well as constant symbols defined in the new theory (*e.g.,* $\mathbf{R}_0$, A_0, A_1, etc.). Note that although the set of atoms has the same real-world cardinality as $\mathbb{N}$, the model does not necessarily know this: in particular, $\mathcal{M}$ does not contain a surjection from the set $\mathbb{A}$ of atoms onto the set $\mathbb{N}$ of all type-0 objects. (If $\mathcal{M}$ contained a surjection f from $\mathbb{A}$ onto $\mathbb{N}$, then we could define E_1 from A_1 by:

$$E_1 :=_2 (\lambda x^1.(A_1(\lambda y^\alpha.(x(fy))))),$$

contradicting conservation.)

We will show that every closed term t of numerical type is equivalent to a closed term t' all of whose subterms are also numerical. This means that the term t' uses only constant symbols from $\mathcal{M}$, and hence corresponds to an element of $\mathcal{M}$. Further, we will show that if t and t' are closed terms of numerical type, both using only constant symbols from $\mathcal{M}$, that are equivalent in $\mathcal{N}$, then they are also equivalent in $\mathcal{M}$. As in the proof of Theorem 2.5, this will show conservation by showing that the numerical part of $\mathcal{N}$ is isomorphic to $\mathcal{M}$.

As in the previous conservation proofs, once we have shown this fact it will be clear that the resulting model $\mathcal{N}$ satisfies the new theory. That the resulting model $\mathcal{N}$ satisfies QF-AC^{1,0} follows from the fact that terms of numerical type are equivalent to terms in $\mathcal{M}$. Suppose that $\Phi(f, n)$ is a quantifier-free formula, possibly with additional parameters, not displayed, of numerical or non-numerical types, such that:

$$\forall f^1 \exists n^0 (\Phi(f, n)).$$

Without loss of generality, Φ is "$t[f, n, \ldots] =_0 1$" for some term t that shares Φ's parameters. (The parameters of Φ other than f and n are hidden in the ellipsis.) For each tuple of parameters, $t[f, n, \ldots]$ is a closed term of type 0, which is a numerical type. By assumption, for each tuple of undisplayed parameters, t is equivalent to a type-$(1, 0) \to 0$ element G in the original model, $\mathcal{M}$. Since $\mathcal{M}$ satisfies QF-AC^{1,0}, it contains a type-2 functional F such that:

$$\forall f^1((G\, f\, (F f)) =_0 1).$$

So $F \in \mathcal{N}$, as well, which means that $\mathcal{N}$ satisfies QF-AC^{1,0}. That $\mathcal{N}$ also satisfies QF-IA follows from the same argument—$\mathcal{M}$ satisfies QF-IA, and the terms involved in the QF-IA schema are of numerical type.

Now let τ be an arbitrary numerical type, and let x be an arbitrary type-τ element of the model $\mathcal{N}$. Let t be a closed term defining x; we will show that there is a closed term s, all of whose subterms are of numerical type, that also defines x. From this we can conclude that x is defined by a constant from $\mathcal{M}$, and hence that every element of $\mathcal{N}$ of numerical type corresponds to an element of $\mathcal{M}$.

To get such a term s, simply rewrite t using the following rules:

- Replace type-α constant $\mathbf{c_i}$ occurring in t with a type-0 constant $\mathbf{i}$.
- Replace type-α variable x_i occurring in t with the type-0 variable n_i; and adjust the types of the combinators (that is, of λ-abstraction) accordingly.

  Note that by the same argument used in the proof of Theorem 2.5, we may assume that all variables occurring in t that are not of numerical type are of type α (since we may assume that the term t is fully reduced).
- Replace $(A_0\,(\mathbf{P}\,t_1\,t_2\,t_3)\,t_4)$ with:

  $$(E_0\,0\,,t_3)\cdot(A_0\,t_1\,t_4)+\neg(E_0\,0\,t_3)\cdot(A_0\,t_2\,t_4),$$

  where E_0 is the type-$((0,0)\to 0)$ characteristic function of the relation $=_0$, defined by:

  $$(E_0\,t_1\,t_2):=_0(1\dot{-}[(t_1\dot{-}t_2)+(t_2\dot{-}t_1)]).$$

  Note that any occurrence of the constant $\mathbf{P}$ in t must appear inside an application involving either $\mathbf{P}$ or $\mathcal{A}_0$, so applying the above rule inductively (from the outside in) eliminates all occurrences of $\mathbf{P}$.
- Replace $(A_0\,t_1\,t_2)$ with the expression:

  $$(E_0\,s_1\,s_2),$$

  where s_1 and s_2 are the type-0 terms corresponding to the type-α terms t_1 and t_2, rewritten using these rules.

Now if we were to continue naïvely we would try to replace A_1 with E_1—which does not work for proving conservation over the base theory RCA_0^ω. However, note that the closed terms defining type-$(\alpha+1)$ elements of $\mathcal{N}$ are, in general, much simpler than the type-1 elements of $\mathcal{M}$. (For example, $\mathcal{N}$ does not have addition and subtraction functions for atoms.)

Replacing A_1 is straightforward but requires a fair amount of string-wrangling. Given t, we first eliminate all occurrences of $\mathbf{P}$ by the method given above; we then eliminate all occurrences of A_1 by the method described below; finally, we use the remaining rules above to eliminate all subterms of non-numerical type.

Suppose t has a subterm of the form:

$$(A_1(\lambda y^\alpha.t'[y^\alpha, x_1^\alpha, \ldots, x_k^\alpha; \ldots])),$$

where the free variables of t' may have include variables of numerical types, not displayed (and hidden in the last ellipsis), in addition to the displayed type-α variables. By the induction hypothesis, the constant A_1 does not occur in t'.

- List all of the occurrences of A_0 in t' as $(A_0\, s_1^\alpha\, t_1^\alpha), \ldots, (A_0\, s_n^\alpha\, t_n^\alpha)$. Note that there are 2^n possibilities for values of $(A_0\, s_1^\alpha\, t_1^\alpha), \ldots, (A_0\, s_n^\alpha\, t_n^\alpha)$. By assumption, since we have eliminated all occurrences of $\mathbf{P}$, each s_i and t_i is either a type-α variable or a type-α constant symbol. Replace $(A_1(\lambda y.t))$ by the term:

$$\bigvee_{\sigma\in 2^n}\left(A_1\left(\lambda y.\left(\bigwedge_{1\le i\le n}(\sigma(i))\cdot(A_0\, s_i\, t_i)+(1-\sigma(i))\cdot\neg(A_0\, s_i\, t_i)\right)\right)\wedge t''\right)\right),$$

where t'' is derived from t' by replacing the subterm "$\mathcal{A}_0(s_i, t_i)$" with either "0" or "1," depending on the value of $\sigma(i)$. Here we are using "$\vee$" and "$\wedge$" as boolean operators on 0's and 1's:

$$(s_1 \vee s_2) :=_0 \operatorname{sign}(s_1 + s_2),$$

and:

$$(s_1 \wedge s_2) :=_0 s_1 \cdot s_2,$$

while $\sigma(i)$ is either the constant 0 or the constant 1, depending on σ and i.

Note that all subterms of the resulting term t'' are of numerical type. To save space, we will write:

$$(A_0^{\sigma(i)}\, s_i\, t_i)$$

to abbreviate:

$$(\sigma(i)) \cdot (A_0\, s_i\, t_i) + (1 - \sigma(i)) \cdot \neg(A_0\, s_i\, t_i).$$

All we have done by rewriting $(A_1(\lambda y.t'))$ in this way is to separate the term into two parts: the first of which is a finite conjunction of A_0's and $\neg A_0$'s, and the second of which contains subterms only of numerical types.

- Noting that t'' does not depend on y (since it has no non-numerical subterms), we can rewrite $(A_1(\lambda y.t'))$ again as:

$$\bigvee_{\sigma \in 2^n} \left[\left(A_1 \left(\lambda y. \bigwedge_{1 \leq i \leq n} \left(\mathcal{A}_0^{\sigma(i)}\, s_i\, t_i \right) \right) \right) \wedge t'' \right], \tag{3.1}$$

  pulling t'' outside of the A_1. We now concern ourselves solely with the finite conjunction of A_0's and $\neg A_0$'s and ignore t''.

- If neither s_i nor t_i is the variable y then we may pull the corresponding $(A_0\, s_i\, t_i)$ or $\neg(A\, s_i\, t_i)$ outside of the A_1. Without loss of generality, assume that for every variable x_j in $x_1, \ldots, x_k$, exactly one of $(A_0\, x_j\, y)$ or $\neg(A_0\, x_j\, y)$ occurs in the first part of the term in equation (3.1). Listing the pairs to which A_0 is applied as:

$$(x_1, y), \ldots, (x_k, y), (\mathbf{c_{i_{k+1}}}, y), \ldots, (\mathbf{c_{i_m}}, y), (x_{i_{m+1}}, x_{i'_{m+1}}), \ldots, (x_{i_n}, x_{i'_n}),$$

  we can rewrite $(A_1(\lambda y.t'))$ once again as:

$$\bigvee_{\sigma \in 2^n} \left[\left(A_1 \left(\lambda y. \bigwedge_{1 \leq j \leq k} \left(A_0^{\sigma(j)}\, x_j\, y \right) \wedge \bigwedge_{k < j \leq m} \left(A_0^{\sigma(j)}\, \mathbf{c_{i_j}}\, y \right) \right) \right) \cdot \right.$$
$$\left. \wedge \bigwedge_{m < j \leq n} \left(A_0^{\sigma(j)}\, x_{i_j}\, x_{i'_j} \right) \wedge t'' \right].$$

- Finally, we will eliminate the constant A_1 from $(A_1(\lambda y.t'))$ by replacing, for each σ, the subterm:

$$\left(A_1 \left(\lambda y. \bigwedge_{1 \leq j \leq k} A_0^{\sigma(j)}(x_j, y) \wedge \bigwedge_{k < j \leq m} A_0^{\sigma(j)}(\mathbf{c_{i_j}}, y) \right) \right), \tag{3.2}$$

  with an equivalent term in which A_1 does not occur.

Note that the term in equation (3.2) equals 1 if and only if there is an atom y that satisfies all of $(A_0^{\sigma(j)}\, x_j\, y)$ and $(A_0^{\sigma(j)}\, \mathbf{c_{i_j}}\, y)$ simultaneously. And these are all satisfied simultaneously if and only if:

- For all $j, j' \leq k$ such that $\sigma(j) = \sigma(j') = 1$, thus requiring that both x_j and $x_{j'}$ be equal to y, $(A_0\, x_j\, x_j') =_0 1$.
- For all $j, j' \leq k$ such that $\sigma(j) \neq \sigma(j')$, thus requiring that exactly one of x_j and $x_{j'}$ be equal to y, $(A_0\, x_j\, x_{j'}) =_0 0$.

  Note that the third possibility, $\sigma(j) = \sigma(j') = 0$, may be ignored, since the set of atoms is infinite. (So there is always a y that does not equal either x_j or $x_{j'}$.
- For all $j \leq k$ and $k < j' \leq m$ such that $\sigma(j) = \sigma(j') = 1$, thus requiring that both x_j and $\mathbf{c}_{i_{j'}}$ be equal to y, $(A_0\, x_j\, \mathbf{c_{i_{j'}}}) =_0 1$.
- For all $j \leq k$ and $k < j' \leq m$ such that $\sigma(j) \neq \sigma(j')$, thus requiring that exactly one of x_j and $\mathbf{c}_{i_{j'}}$ be equal to y, $(A_0\, x_j\, \mathbf{c}_{i_{j'}}) =_0 0$.

  Again, the third possibility, $\sigma(j) = \sigma(j') = 0$, may be ignored, since the set of atoms is infinite.

There are only finitely many of these conditions, and none of them depends on y. So we may rewrite the first part of the conjunction using only these conditions. Then the subterm to which A_1 is applied in the resulting, rewritten, term, may be pulled outside of A_1. This allows us to eliminate A_1, since if y does not occur in a term w then:

$$(A_1(\lambda y.w)) =_0 w.$$

Since the closed term s resulting from rewriting t using the above rules contains no subterms of non-numerical type, the element represented by s is in $\mathcal{M}$. Since the rewriting rules above clearly yield an equivalent term, $s \approx t$, so the element represented by t is in $\mathcal{M}$.

Finally, suppose that x and y are distinct elements of $\mathcal{M}$; we will show that they are distinct elements of $\mathcal{N}$, completing the proof. Let s be a closed term for x and let t be a closed term for y; suppose $s \approx t$ by some finite list of syntactic and semantic reductions.

Note that the only new constants we add, besides constants for atoms, are for the functionals A_0, A_1, and $\mathbf{P}$—so our only new semantic rules are for these three functionals. However, each of A_0, A_1, and $\mathbf{P}$ yields exactly one output for every input, so the proof from Theorem 2.5 works and $s \approx t$ in $\mathcal{M}$, contradicting $x \neq y$. □

However, if we also throw in the axiom $(\mathcal{A}_2)$, the resulting theory is not conservative over RCA_0^ω:

Proposition 3.2 (RCA_0^ω + (ATOMS))**.** $(\mathcal{A}_2) + (\mathcal{A}_1)$ *implies* $(\mathcal{E}_1)$.

Proof. Let f be an arbitrary type-1 function. Let $\mathcal{F}$ be the type-$(\alpha+2)$ functional defined by:

$$F \in \mathcal{F}_f \iff \forall x^\alpha \forall y^\alpha ((Fx) =_0 (Fy) \wedge (f(Fx)) \neq_0 0).$$

That is, $F \in \mathcal{F}_f$ if and only if F is a constant function and f applied to that constant is not zero.

The theory RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ proves the existence of $\mathcal{F}$, since it was defined using only type-0 equality and type-α quantification. Note that the definition of $\mathcal{F}$ is uniform in f, so in fact our theory proves the existence of the type-$(1 \to (\alpha+2))$ map $f \mapsto \mathcal{F}_f$.

Suppose there is some number $\mathbf{n} \in \mathbb{N}$ such that $(f\,\mathbf{n}) \neq_0 0$. Then RCA_0^ω + (ATOMS) proves the existence of a type-$(\alpha+1)$ functional F defined by:

$$F :=_{\alpha+1} (\lambda x^\alpha.\mathbf{n}).$$

Note that $F \in \mathcal{F}_f$. If, instead, $(f\,\mathbf{n}) =_0 0$ for every $\mathbf{n} \in \mathbb{N}$ then no type-$(\alpha+1)$ functional F is in $\mathcal{F}$, since no functional F could satisfy "$(f(Fx)) \neq_0 0$," for any x.

So we have:

$$\exists n^0((fn) \neq_0 0) \iff \exists F^{\alpha+1}((\mathcal{F}_f\, F) \neq_0 0).$$

This means we can define the functional E_1 from A_2 and the functional $f \mapsto \mathcal{F}_f$ by the rule:

$$(E_1\, f) :=_0 (A_2\, \mathcal{F}_f),$$

thus implying $(\mathcal{E}_1)$. □

Note that the previous proposition assumed both $(\mathcal{A}_2)$ and $(\mathcal{A}_1)$: although $(\mathcal{E}_2)$ implies $(\mathcal{E}_1)$, it is not the case that $(\mathcal{A}_2)$ implies $(\mathcal{A}_1)$.

Proposition 3.3 (RCA_0^ω + (ATOMS))**.** $(\mathcal{A}_2)$ *does not imply* $(\mathcal{A}_1)$.

Proof. Let $\mathcal{M}$ be a minimal ω-term model for RCA_0^ω + (ATOMS) + $(\mathcal{E}_1)$ + $(\mathcal{A}_2)$; we will show that $\mathcal{M}$ does not contain the functional A_1. (The model $\mathcal{M}$ is just a term model, built from the second-order structure (ω, ARITH) by including constant symbols for E_1 and $\mathcal{A}_2$, as in the proof of Theorem 2.5. The first-order part of the structure (ω, ARITH) is the standard ω; the second-order part consists of all functions corresponding to arithmetical sets.)

Suppose not; then A_1 is defined by a term of the form:

$$(\lambda X^{\alpha+1}.t[X]),$$

where t is a type-0 term with only X free. By assuming that t is in normal form, we have that no type-α variable occurs in t; this is the key fact.

This fact follows from the same argument used in the proof of Theorem 2.5: any type-α variable x in t must occur freely in a subterm of t, and must be bound to a λ—which must in turn be captured by a function constant of type $((\alpha \to \sigma) \to \tau)$ for some types σ and τ. However (since we do not include a constant for A_1), there is no such function constant in $\mathcal{M}$, so no type-α variable occurs in t.

Note also that if X were to occur in t on the right-hand side of an application, then the left-hand side would have to be a term of type $((\alpha \to 0) \to \tau)$ for some type τ. However, we just noted that $\mathcal{M}$ contains no function constant of any such type. So, by assuming that t is in normal form we have that X must occur only on the left-hand sides of applications, the right-hand sides of which must be type-α terms not involving variables.

Using the same technique we used in the proof of Proposition 3.1, we may assume that no $\mathbf{P}$'s occur in t. This means that the only occurrences of X in t are of the form $(X\mathbf{c_i})$ where $\mathbf{c_i}$ is some type-α constant. So any A_2's occurring in t are irrelevant: the functional defined by $(\lambda X.t[X])$ decides what number to output on input X based solely on the outputs of finitely many $(X\mathbf{c_i})$'s, where $\mathbf{c_i}$ is one of finitely many atom constants $\mathbf{c}_1, \ldots, \mathbf{c}_k$ occurring in t.

This yields a contradiction—$(\lambda X.t[X])$ cannot define $\mathcal{A}_1$, since it cannot distinguish between $\emptyset$ and the singleton set $\{x\}$ for any $x \notin \{\mathbf{c}_1, \ldots, \mathbf{c}_k\}$. □

Finally, the implication $(\mathcal{A}_2) + (\mathcal{A}_1) \Rightarrow (\mathcal{E}_1)$ is sharp:

Proposition 3.4. *The theory* RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$ *is conservative over* RCA_0^ω + $(\mathcal{E}_1)$*.*

Proof. Following the proof of Proposition 3.1, let $\mathcal{M}$ be an arbitrary model of RCA_0^ω + $(\mathcal{E}_1)$; let $\mathcal{N}$ be the term model built up from constants for every element of $\mathcal{M}$, new type-α constants $\mathbf{c}_i$, for $i \in \mathbb{N}$, and new type-$(\alpha+1)$ constants for functionals with finite support, as defined below.

We include type-$(\alpha+1)$ constant $\mathbf{F}$ in $\mathbb{N}$ if and only if there is, in $\mathcal{M}$, a finite sequence of pairs of integers, coded by $\sigma \in \mathbb{N}$, defining a finite partial function from $\mathbb{N}$ into $\mathbb{N}$ such that:

$$\mathbf{F}(\mathbf{c}_i) =_0 \begin{cases} \sigma(k+1)(1), & \text{if } \exists k^0 < (|\sigma|-1)\ (\sigma(k+1)(0) =_0 i) \\ \sigma(0)(1), & \text{otherwise.} \end{cases}$$

Here $|\sigma|$ denotes the length of the finite string coded by σ.

The code σ specifies the value of $(\mathbf{F}\,\mathbf{c}_i)$ for finitely many (in the sense of the model) i, and specifies $\sigma(0)(1)$ to be the value of $(\mathbf{F}\,x)$ for all other x. We say that the **support** of such a code σ is $\{\sigma(k+1)(0) : k < |\sigma|-1\}$. (Note that the support is a finite (in the sense of the model) list of integers—not a finite list of atoms. Note also that $\mathcal{N}$ is not necessarily aware of the correspondence between the type-0 index i and the type-α atom $\mathbf{c}_i$.)

Add the constant symbols A_0, A_1, and A_2, and close $\mathcal{N}$ under definability. As in the proof of Proposition 3.1, we will show that the submodel of $\mathcal{N}$ consisting of only (equivalence classes for) closed numerical terms of $\mathcal{N}$ is isomorphic to $\mathcal{M}$, using various rewriting rules; this part is similar to the proof of Proposition 3.1. The basic idea is that we included in $\mathcal{N}$ type-$(\alpha+1)$ constants only for functionals that correspond to type-1 functions with finite support. Quantifying over the set of all type-($\alpha+1$) objects should correspond to quantifying over the set of all type-0 objects, since type-1 functions with finite support can be coded, in RCA_0^ω, by type-0 integers.

(The potential problem is that additional type-($\alpha + 1$) functionals might be definable. As we show below, it turns out that this does not happen.)

Any closed numerical term t can be rewritten to have only numerical subterms, using $(\mathcal{E}_1)$ and the following rules:

- replace $(A_0\, x^\alpha\, y^\alpha)$ with $(E_0\, i^0\, j^0)$;
- replace $(A_1\,(\lambda x.t))$ with $(E_1\,(\lambda i.t'))$, where t' is rewritten using these rules;
- replace $(\mathbf{F}x)$ with $\tilde{\sigma}(i)$, where:

$$\tilde{\sigma}(i) := \begin{cases} \sigma(j+1)(1), & \text{if } \exists j < (|\sigma|-1)\ (\sigma(j+1)(0) =_0 i,) \\ \sigma(0)(1), & \text{otherwise,} \end{cases}$$

  and σ is the type-0 code corresponding to the type-($\alpha + 1$) constant $\mathbf{F}$.
- replace $(A_2\,(\lambda F.t))$ with $(E_1\,(\lambda\sigma.t'))$, where t' is rewritten using these rules. Here σ ranges over all type-0 codes for (type-1) finite-support functions on $\mathbb{N}$.

As in the proof of Proposition 3.1, once we have shown that the numerical part of $\mathcal{N}$ is $\mathcal{M}$, we will have shown that the axioms $\mathrm{QF\text{-}AC}^{0,1}$ and QF-IA hold in $\mathcal{N}$. Also, as in the proof of Proposition 3.1, once we have shown that each closed term of numerical type is equivalent to some element of $\mathcal{M}$, we will have that that element is unique. This is because the only new constants we add, besides constants for atoms, are A_0, A_1, A_2, $\mathbf{P}$, and constants for type-($\alpha+1$) functionals—the semantic rules for all of which yield exactly one output for every input, and which respect extensional equality. For example, we define the equivalence relation $\approx$ so that if $(Fx) =_0 (Gx)$, for all type-α inputs $x \in \mathcal{M}$, then $(A_1\, F) \approx (A_1\, G)$.

By its definition, $\mathcal{N}$ is closed under the combinators and satisfies $(\mathbb{A}_2)$ and $(\mathbb{A}_1)$. It remains for us to show that every closed type-($\alpha + 1$) term is equivalent to a type-($\alpha + 1$)

constant—thus showing that:

$$(A_2\,(\lambda F.t)) \approx (E_1\,(\lambda\sigma.t')),$$

and thus that the last rewriting rule, above, is valid.

Let $(\lambda x.t[x])$ be an arbitrary closed type-$(\alpha+1)$ term; we will show, by induction, that there is a finite sequence $\sigma \in \mathbb{N}$ corresponding to the term $(\lambda x.t[x])$, thus showing that this term is equivalent to a constant.

Applying the rewriting rules given above, we get a closed type-1 term $(\lambda n^0.t'[n])$, all of whose subterms are of numerical type. We want to show that the type-1 term $(\lambda n.t')$ corresponds to a finite code—*i.e.,* that the support of this term is finite.

In fact, the support of $(\lambda n.t')$ is the union of (a) the supports of codes in t' corresponding to type-$(\alpha+1)$ constants in t with (b) all indices i in t' corresponding to type-α constants $\mathbf{c}_i$ in t. To see this, let $\mathbf{y}$ and $\mathbf{z}$ be arbitrary type-α constants whose indices are not included in this set; we will show that $t[\mathbf{y}/x] =_0 t[\mathbf{z}/x]$.

Let φ be the permutation on the space $\mathbb{A}$ of type-α elements of $\mathcal{N}$ that swaps $\mathbf{y}$ and $\mathbf{z}$ but fixes all other elements. Let Φ be the permutation on type-$(\alpha+1)$ elements $\mathbf{F}$ of $\mathcal{N}$, induced by φ:

$$(\Phi(\mathbf{F})\,x) = (\mathbf{F}\,\varphi^{-1}(x)).$$

(Both φ and Φ are in $\mathcal{M}$, although this fact is irrelevant for the proof.)

Now if s is any type-0 subterm of t containing various free type-α variables $x_1,\ldots,x_k$, as well as various type-$(\alpha+1)$ variables $X_1,\ldots,X_m$ and type-0 variables $\overline{n}$, then:

$$\begin{aligned} &s[\mathbf{x_1}/x_1,\ldots,\mathbf{x_k}/x_k;\mathbf{X_1}/X_1,\ldots,\mathbf{X_m}/X_m;\overline{\mathbf{n}}/\overline{n}] \\ &\qquad =_0 s[\varphi(\mathbf{x_1})/x_1,\ldots,\varphi(\mathbf{x_k})/x_k;\Phi(\mathbf{X_1})/X_1,\ldots,\Phi(\mathbf{X_m})/X_m;\overline{\mathbf{n}}/\overline{n}]. \quad (3.3)\end{aligned}$$

(By assuming that s is a subterm of t we ensure that $\mathbf{y}$ and $\mathbf{z}$ occur neither as type-α constants in s nor in the support of any constant occurring in s.) The proof is by induction on terms. For simplicity we will sometimes write "$t[\mathbf{x}]$," below, in place of "$t[\mathbf{x}/x]$."

Note that variable x_i occurs in s only in primitives of the form:

- $(\mathbf{F}x_i)$, where $\mathbf{F}$ is a type-$(\alpha+1)$ constant: by assumption, neither $\mathbf{y}$ nor $\mathbf{z}$ are in $\mathbf{F}$'s support, so φ is the identity on $\mathbf{F}$'s support, and so $(\mathbf{F}\mathbf{x_i}) =_0 (\mathbf{F}\varphi(\mathbf{x_i}))$.
- $(A_0\, x_i\, \mathbf{c})$, where $\mathbf{c}$ is a type-α constant: by assumption, $\mathbf{c} \neq \mathbf{y}, \mathbf{z}$, so $(A_0\, \mathbf{x_i}\, \mathbf{c}) =_0 (A_0\, \varphi(\mathbf{x_i})\, \mathbf{c})$.
- $(A_0\, x_i\, x_j)$, where x_j is another type-α variable: since φ is a bijection, $\mathbf{x}_i =_\alpha \mathbf{x}_j$ if and only if $\varphi(\mathbf{x}_i) =_\alpha \varphi(\mathbf{x}_j)$, so $(A_0\, \mathbf{x_i}\, \mathbf{x_j}) =_0 (A_0\, \varphi(\mathbf{x_i})\, \varphi(\mathbf{x_j}))$.
- $(X_j x_i)$, where X_j is a type-$(\alpha+1)$ variable: by definition, $(\mathbf{X_j}\, \mathbf{x_i}) =_0 (\Phi(\mathbf{X_j})\varphi(\mathbf{x_i}))$.

And variable X_i occurs in s only in primitives of the form:

- $(X_i x_j)$, where X_j is a type-$(\alpha+1)$ variable: this is the previous case.
- $(A_1\, X_i)$: since φ is a bijection, $(A_1\, \mathbf{X_i}) =_0 (A_1\, \Phi(\mathbf{X_i}))$.
- $(X_i \mathbf{c})$, where $\mathbf{c}$ is a type-α constant: by assumption, $\mathbf{c} \neq \mathbf{y}, \mathbf{z}$, so $(\mathbf{X_i}\, \mathbf{c}) =_0 (\Phi(\mathbf{X_i})\, \mathbf{c})$.

We have two induction cases: $(A_1\, (\lambda x_i.s[\overline{x}; \overline{X}; \overline{n}]))$ and $(A_2\, (\lambda X_i.s[\overline{x}; \overline{X}; \overline{n}]))$. The hypothesis holds in the first case since φ is a bijection, and in the second case since Φ is a bijection. Formally, in the first case we have:

$$\begin{aligned}
(A_1\, (\lambda x_i.s[\overline{\mathbf{x}}; \overline{\mathbf{X}}; \overline{\mathbf{n}}])) &=_0 (A_1\, (\lambda x_i.s[\varphi(\mathbf{x_1}), \ldots, \varphi(\mathbf{x_i}), \ldots, \varphi(\mathbf{x_k}); \Phi(\mathbf{X_1}), \ldots, \Phi(\mathbf{X_m}); \overline{\mathbf{n}}])) \\
&=_0 (A_1\, (\lambda x_i.s[\varphi(\mathbf{x_1}), \ldots, \quad \mathbf{x_i} \quad, \ldots, \varphi(\mathbf{x_k}); \Phi(\mathbf{X_1}), \ldots, \Phi(\mathbf{X_m}); \overline{\mathbf{n}}])).
\end{aligned}$$

The second case is similar.

Applying equation (3.3) to $t[x]$, we have that $t[\mathbf{x}] =_0 t[\varphi(\mathbf{x})]$, for all $\mathbf{x} \in \mathbb{A}$, so $t[\mathbf{y}] =_0 t[\mathbf{z}]$. □

3.3 Topology

3.3.1 Topological spaces

A **topological space** $(X, \mathcal{T})$ is a type-$(\alpha+1)$ characteristic functional X together with a type-$(\alpha+2)$ characteristic functional $\mathcal{T}$ satisfying:

- If $U \in \mathcal{T}$ then $U \subseteq X$. (This is a nice requirement to have, and costs us nothing in theories including $(\mathcal{A}_1)$.)
- $\emptyset \in \mathcal{T}$ and $X \in \mathcal{T}$.
- If $U, V \in \mathcal{T}$ then $U \cap V \in \mathcal{T}$.
- If $\mathcal{F} \subseteq \mathcal{T}$ then $\bigcup \mathcal{F}$ exists and $\bigcup \mathcal{F} \in \mathcal{T}$. (Alternatively, we may drop the requirement that $\bigcup \mathcal{F}$ exist and require only that if $\bigcup \mathcal{F}$ exists, $\bigcup \mathcal{F} \in \mathcal{T}$. The results in this paper hold for either definition.)

Note that, *a priori*, it is possible, in models not satisfying $(\mathcal{A}_2)$, to have a family $\mathcal{F}$ of sets without having $\bigcup \mathcal{F} := \{x : \exists U \in \mathcal{F}(x \in U\}$. Our definition of "topology" requires that the unions of all open families exist; note that the existence of unions of open families does not necessarily imply $(\mathcal{A}_2)$.

That $(\mathcal{A}_2)$ implies the existence of the set $\bigcup \mathcal{F}$ follows from using A_2 to defining $\bigcup \mathcal{F}$, as:

$$\bigcup \mathcal{F} :=_{\alpha+1} \left(\lambda x. \begin{cases} 1, & \text{if } \exists Y^{\alpha+1} \in \mathcal{F}\left((Y\,x) =_0 1\right) \\ 0, & \text{otherwise} \end{cases} \right) .)$$

3.3.2 Product spaces

Suppose we have topological spaces $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$. Suppose we also have:

1. a type-$(\alpha, \alpha) \to \alpha$ bijection f from $X \times Y$ onto Z (*i.e.*, for all $x \in X$ and $y \in Y$, $f(x)(y) \in Z$);
2. a type-$(\alpha \to \alpha)$ function g_X from Z onto X; and
3. a type-$(\alpha \to \alpha)$ function g_Y from Z onto Y.

Then we say that Z is the **product set** $X \times Y$.

Note that the real product set $X \times Y$ is a subset of $\mathbb{A} \times \mathbb{A}$; however, the pairing function f and projection functions g_X and g_Y give us in $\mathbb{A}$, via Z, an isomorphic copy of $X \times Y$

that is a subset of $\mathbb{A}$. This is helpful, because the axioms we consider deal only with $\mathbb{A}$ and not with finite products of $\mathbb{A}$. For convenience, we use "$\langle x, y\rangle$" to refer to $f(x)(y)$; to avoid confusion, we use "(x, y)" to refer to the corresponding element of $\mathbb{A} \times \mathbb{A}$.

Suppose now that Z, the product set $X \times Y$, has a topology $\mathcal{T}_Z$ satisfying the two axioms:

$$\forall W \in \mathcal{T}_Z \, \forall x \in W \, \exists U \in \mathcal{T}_X \, \exists V \in \mathcal{T}_Y \, (x \in U \times V \subseteq_{\alpha+1} W),$$

and:

$$\forall U \in \mathcal{T}_X \, \forall V \in \mathcal{T}_Y \, \exists W \in T_Z (W =_{\alpha+1} U \times V),$$

where here "$U \times V$" refers to the subset of Z:

$$\{\langle x, y\rangle : x \in U \wedge y \in V\}.$$

Then we say $\mathcal{T}_Z$ is the **product topology** on Z $(= X \times Y)$.

Note that by definition the product topology is generated by the basis of open rectangles $U \times V$. Since "$U \times V$," as a subset of $X \times Y$, is definable in RCA_0^ω + (ATOMS), the product topology must contain all open rectangles (in the sense of the model); it may also (and usually will) contain additional open sets.

In the following sections we will take RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ to be our base theory; inside each model of this theory, a product topology is unique. To see this, suppose that both $\mathcal{T}_Z$ and $\mathcal{T}_Z'$ are topologies satisfying the two axioms above and generated by the same open rectangles. To obtain a contradiction, let $W' \in \mathcal{T}_Z'$ be an set open in the topology $\mathcal{T}_Z'$ but not in the topology $\mathcal{T}_Z$. Then the open family:

$$\mathcal{F} := \{W \in \mathcal{T}_Z : W \subseteq_{\alpha+1} W'\},$$

exists, since it can be defined using A_1. Since both $\mathcal{T}_Z$ and $\mathcal{T}_Z'$ are product topologies, they both contain all open rectangles, so $\bigcup \mathcal{F} =_{\alpha+1} W'$, and so $W' \in \mathcal{T}_Z$, which contradicts our assumption that it was not.

3.3.3 Compact spaces

A **sequence of sets** is a type-$(0 \to (\alpha+1))$ functional $\langle X_n : n \in \mathbb{N}\rangle$. A **finite sequence of sets** is a sequence of sets, together with a number $N \in \mathbb{N}$; we write this $\langle X_n : n < N\rangle$. We interpret $\langle X_n : n < N\rangle$ to be the (finite) subsequence of $\langle X_n : n \in \mathbb{N}\rangle$ consisting of the first N sets.

Note that a sequence of sets is a type-$(0 \to (\alpha+1))$ functional and thus does not fit nicely into our hierarchy of axioms $(\mathcal{A}_1)$, $(\mathcal{A}_2)$, etc. However, a finite sequence of sets is equivalent to a type-$(\alpha + 1)$ functional $F_{X,N}$—which does fit nicely into our axiom hierarchy—of the form:

$$F_{X,N}(x) := \langle X_0(x), X_1(x), \ldots, X_{N-1}(x)\rangle,$$

allowing us to use this latter formula as the definition. (Here we use $\langle \ldots \rangle$ to refer to some type-0 pairing/finite-list functional definable in RCA_0^ω.) In the former definition, given a type-0 index $n < N$, the finite sequence of sets returns the nth set in the sequence; in the latter definition, given a type-α element x, the finite sequence of sets returns an N-tuple of zeros and ones giving the element x's membership in all N sets.

One can pass between the two definitions via two functionals that can be defined in RCA_0^ω + (ATOMS). (This is just the standard trick of hiding a finite number (in the sense of the model) in the output of a type-$(\sigma \to 0)$ functional.)

Lemma 3.5 (RCA_0^ω + (ATOMS))**.** *There are a type-$(((0 \to (\alpha+1)), 0) \to (\alpha+1))$ functional G and a type-$((\alpha + 1) \to ((0 \to (\alpha + 1)), 0))$ functional H such that:*

$$G(\langle X_n : n < N\rangle) = F_{X,N},$$

and:

$$H(F_{X,N}) = \langle X_n : n < N\rangle.$$

Proof. Define G by primitive recursion, using helper function G':

$$\begin{aligned} G'\ (\langle X_n : n < N\rangle)\ (x, 0) &:=_0 \langle\rangle;\ \text{(the empty string)} \\ G'\ (\langle X_n : n < N\rangle)\ (x, n + 1) &:=_0 G'(x, n)^\frown\langle X_n(x)\rangle. \end{aligned}$$

(The functional G' can be defined using R_0.) Then define:

$$G(\langle X_n : n < N\rangle)(x) := G'(\langle X_n : n < N\rangle)(x, N),$$

and note that $F_{X,N}$ is $G(\langle X_n : n < N\rangle)$.

Define:

$$H(F_{X,N})(n)(x) :=_0 (\langle Y_n : n \in \mathbb{N}\rangle, M),$$

where $M = |(F_{X,N}(x_0))|$, the length of the finite type-0 string $(F_{X,N}(x_0))$, for some fixed x_0, and $\langle Y_n : n \in \mathbb{N}\rangle$ is given by:

$$\langle Y_n : n \in \mathbb{N}\rangle(m)(x) = \begin{cases} (F_{X,N}(x))_m & \text{if } m < M; \\ 0 & \text{otherwise.} \end{cases}$$

Since the operations on finite strings of numbers can be defined in RCA_0^ω, RCA_0^ω + (ATOMS) proves the existence of the functional H. □

A topological space $(X, \mathcal{T})$ is **compact** if and only if for every open cover $\mathcal{F} \subseteq \mathcal{T}$ there is a finite sequence of sets $\langle U_n : n < N\rangle$ such that:

- $\forall n\, (U_n \in \mathcal{F})$; and
- $\forall x \in X\, \exists n\, (x \in U_n)$.

In other words, $(X, \mathcal{T})$ is compact iff every open cover has a finite subcover.

3.4 Some reverse-mathematical results in $\mathbf{RCA}_0^\omega$ + (ATOMS) and related theories

3.4.1 The product of two compact spaces

It is a theorem of ZF that the product of any two compact spaces is compact. Tychonoff's Theorem, which requires the Axiom of Choice, further states that *any* product, possibly infinite, of compact spaces is compact. The reverse-mathematical strength of the latter theorem has been studied (it is equivalent, over ZF, to the Axiom of Choice; see, *e.g.*, [Jec73], problem 2.6.9, on p. 26), but to my knowledge no one has studied the former theorem in the context of reverse mathematics.

In general, ZF has served as a suitable base theory for the reverse-mathematical examination of choice principles. However, the former theorem follows from a comprehension axiom, which means that it holds in ZF. Traditional, second-order, reverse mathematics provides a weaker base theory in which one can study the consequences of comprehension axioms. However, "compactness" is a third-order (over atoms) concept and thus cannot be expressed (let alone studied) in traditional, second-order reverse mathematics, except through indirect codes. The former theorem is a good candidate for higher-order reverse mathematics.

The proof that the product of two compact spaces is compact given in Munkres's *Topology* [Mun75] goes through in RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$.

Proposition 3.6 (RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$)**.** *The product of two compact spaces is compact.*

Proof. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be compact spaces; let $(X \times Y, \mathcal{T})$ be a product space for them. Let $\mathcal{F} \subseteq \mathcal{T}$ be an arbitrary open cover of $X \times Y$.

For each $x \in X$, the family $\mathcal{F}_x$ of open subsets of Y:

$$\mathcal{F}_x := \{V \in \mathcal{T}_Y : \exists W \in \mathcal{F}\,(x \in W \wedge V \subseteq_{\alpha+1} \{y : \langle x, y\rangle \in W\}\,)\,\},$$

definable using A_2, is an open cover of Y. Since Y is compact, the cover $\mathcal{F}_x$ has a finite subcover $\langle V_n : n < N\rangle$. By the definition of $\mathcal{F}_x$, for every $n < N$ there exists a subset $W_n \in \mathcal{T}$ of the product space $X \times Y$ such that V_n is a subset of the intersection of W_n with the fiber $\{x\} \times Y$. Applying $(\mathcal{A}_2)$ and QF-IA, given $\langle V_n : n < N\rangle$ there is a corresponding finite sequence $\langle W_n : n < N\rangle$ of elements of $\mathcal{F}$ covering the fiber $\{x\} \times Y$.

So we have:

$$\forall x\, \exists \langle W_n : n < N\rangle\ ((\forall n\,(W_n \in \mathcal{F})) \wedge (\forall y\, \exists n(\langle x, y\rangle \in W_n))).$$

This says that every fiber $\{x\} \times Y$ is covered by a finite subcollection of $\mathcal{F}$. Following Munkres we now project $\langle W_n : n < N\rangle$ onto X and take the intersection of the projection.

The axiom $(\mathcal{A}_2)$ is sufficient to prove that the intersection of finitely many open sets is open; since the intersection contains x, it must also be non-empty. Formally, we have:

$$\forall x\, \exists U \in \mathcal{T}_X\, \exists \langle W_n : n < N \rangle \text{ as above } \left(U \subseteq_{\alpha+1} \bigcap_{n<N} \{z : \exists y\, (\langle z, y \rangle \in W_n)\} \right).$$

(Clearly $(\mathcal{A}_1)$ and $(\mathcal{A}_2)$ suffice to prove the above logical statement.) So the collection $\mathcal{F}' \subseteq \mathcal{T}_X$ of all such open subsets U of X is an open cover of X. Since X is compact, $\mathcal{F}'$ has a finite subcover $\langle U_n : n < M \rangle$.

Then for each $k < M$ there is a finite sequence $\langle W_n^k : n < N_k \rangle$ the intersection of whose projection onto X contains V_k and all of whose elements are in $\mathcal{F}$. This gives us a finite subcover of $X \times Y$, proving compactness. □

However, in the absence of $(\mathcal{A}_2)$, the product of two compact spaces need not be compact.

3.4.2 $\mathbf{RCA}_0^\omega$ + $(\mathcal{E}_1)$ + (ATOMS) + $(\mathcal{A}_1)$ does not imply that the product of two compact spaces is compact

We construct a minimal term model $\mathcal{M}$ (which is an ω-model) of RCA_0^ω + (ATOMS) + $(\mathcal{E}_1)$ + $(\mathcal{A}_1)$, using the standard term model construction as in the conservation proofs. As in the proof of Theorem 2.5, we include a constant symbol for μ (rather than E_1); in addition, we include the following constants:

- Countably many type-α constants $\mathbf{d}_i$, for $i \in \omega$, representing the elements of our topological spaces and an "error" atom.

- Three type-$(\alpha+1)$ constants $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ representing characteristic functions of three disjoint, countably infinite sets of atoms.

  The constants $\mathbf{A}$ and $\mathbf{B}$ are intended to represent disjoint copies of $\mathbb{Q} \cap (0,1)$, between which the model has no bijection. The constant $\mathbf{C}$ is intended to represent a copy of of the product set $\mathbf{A} \times \mathbf{B}$ (as defined in section 3.3.2).

- Two type-$(\alpha \to \alpha)$ constants, $\pi_{\mathbf{A}}$ and $\pi_{\mathbf{B}}$, representing projections from $\mathbf{C}$ onto $\mathbf{A}$ and $\mathbf{B}$, and one type-$((\alpha,\alpha) \to \alpha)$ constant, $\langle \cdot, \cdot \rangle$, representing the pairing function from $\mathbf{A} \times \mathbf{B}$ onto $\mathbf{C}$. These constants satisfy the axioms from section 3.3.2. When given an invalid input, these three functions return the "error" atom. (*E.g.*, $(\pi_{\mathbf{B}}\, x) =_\alpha$ error $\iff x \notin \mathbf{C}$.)

- For each open interval contained in $\mathbb{Q} \cap (0,1)$ (with rational or irrational endpoints), a type-$(\alpha+1)$ constant for a copy $\subseteq \mathbf{A}$ and a copy $\subseteq \mathbf{B}$. Similarly, for each open rectangle contained in $(\mathbb{Q} \cap (0,1))^2$, a type-$(\alpha+1)$ constant for a copy $\subseteq \mathbf{C}$. These constants will be the basic open sets for the respective topological spaces.

- Three type-$(\alpha+2)$ constants $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, and $\mathcal{T}_{\mathbf{C}}$, with the interpretation that for every type-$(\alpha+1)$ functional U in the model, $U \in \mathcal{T}_{\mathbf{A}}$ if and only if (a) $U \subseteq_{\alpha+1} \mathbf{A}$ and (b) U corresponds to an open subset of $\mathbb{Q} \cap (0,1)$, in the standard topology; and similarly for $\mathcal{T}_{\mathbf{B}}$ and $\mathcal{T}_{\mathbf{C}}$. We will show, in Lemma 3.10, that all open sets U are finite unions of open intervals or rectangles.

- One type-$(\alpha+2)$ constant $\mathcal{F}$ whose interpretation is as follows. Let $c_n = 1 - (1/\pi^n)$, for $n \in \omega$, defined outside of the model. Then $\{c_n : n > 0\}$ is a sequence of irrational numbers converging to 1. $\mathcal{F}$ represents the subfamily of $\mathcal{T}_{\mathbf{C}}$ consisting of, for each n, the (equivalence classes of terms corresponding to the) three clopen rectangles: $(c_n, c_{n+1}) \times (c_n, c_{n+1})$, $(c_{n+1}, 1) \times (c_n, c_{n+1})$, and $(c_n, c_{n+1}) \times (c_{n+1}, 1)$. (In other words, for each n the family $\mathcal{F}$ contains all but the upper-right rectangle $(c_{n+1}, 1) \times (c_{n+1}, 1)$.)

  Clearly $\mathcal{F}$ covers $\mathbf{C}$, and just as clearly it does not contain a finite subcover. So, if we can show that $\mathcal{M} \models$ "$\mathcal{T}_{\mathbf{C}}$ is a topology," we have that $\mathcal{M} \models$ "$\mathcal{T}_{\mathbf{C}}$ is a topology that is not compact."

Note that if $\mathcal{T}_A$, $\mathcal{T}_B$, and $\mathcal{T}_C$ are topologies then $\mathcal{T}_C$ is the product topology on $\mathbf{C} = \mathbf{A} \times \mathbf{B}$, since it is generated by open rectangles from $\mathbf{A} \times \mathbf{B}$.

The term model consists of (equivalence classes of) closed terms involving the various constants listed above. We have the following lemma, which is informative if not especially relevant to the subsequent proof:

Lemma 3.7. *In the resulting term model, if f is any type-$(\alpha \to \alpha)$ function such that $f \restriction \mathbf{A}$ maps $\mathbf{A}$ into $\mathbf{B}$, then $f \restriction \mathbf{A}$ has finite range.*

(Note that, by symmetry, the same holds with "**A**" and "**B**" swapped.) This lemma says that although **A** and **B** represent copies of the same set, the model does not know this.

Proof. Let f be an arbitrary function satisfying the lemma's hypotheses; let:

$$(\lambda x^{\alpha}.\mathbf{P}(s[x]\, t[x]\, \mathbf{A}x))$$

be a type-$(\alpha \to \alpha)$ term defining f. (The idea here is that $t[x]$ defines f on domain **A**, while $s[x]$ defines f outside of **A**. We care only about $t[x]$.)

Recall that the constants included in our term model that output elements of type α are:

- the type-$(\alpha \to \alpha)$ projection functions $\pi_{\mathbf{A}}$ and $\pi_{\mathbf{B}}$,
- the type-$((\alpha, \alpha) \to \alpha)$ pairing function $\langle \cdot, \cdot \rangle$, and
- the type-$((\alpha, \alpha, 0) \to \alpha)$ finite-list function **P**.

(Of these, only $\pi_{\mathbf{B}}$ and **P** map into **B**, since **A**, **B**, and **C** are all pairwise disjoint.) If a type-α term does not begin with one of these four constants, then it must be some type-α constant **d** or some type-α variable y.

We first show that we may ignore the function $\pi_{\mathbf{B}}$. For $(\pi_{\mathbf{B}}\, t'[x])$ to yield an output in **B** the input $t'[x]$ must be in **C**. (Otherwise, $\pi_{\mathbf{B}}$ yields "error.") If $t'[x] \in \mathbf{C}$ then one of the following must be true:

- $t'[x] \equiv \mathbf{c}$, for some constant $\mathbf{c} \in \mathbf{C}$,
- $t'[x] \equiv \langle t_1[x], t_2[x] \rangle$, or
- $t'[x] \equiv (\mathbf{P}\, t_1[x]\, t_2[x]\, t_3[x])$, where t_1 and t_2 are of type α, and t_3 is of type 0.

In the first case we can replace "$(\pi_{\mathbf{B}}\, \mathbf{c})$" with "**b**," where **b** is the type-α constant in **B** such that $\mathbf{b} =_{\alpha} (\pi_{\mathbf{B}}\, \mathbf{c})$. (So $\mathbf{b} \approx (\pi_{\mathbf{B}}\, \mathbf{c})$.) In the second case we can replace $\pi_{\mathbf{B}} \langle t_1[x], t_2[x] \rangle$ with the equivalent subterm $t_2[x]$. In the third case we can rewrite $(\pi_{\mathbf{B}}\, (\mathbf{P}\, t_1[x]\, t_2[x]\, t_3[x]))$ as $(\mathbf{P}\, (\pi_{\mathbf{B}}\, t_1[x])\, (\pi_{\mathbf{B}}\, t_2[x])\, t_3[x])$, at which point we can apply induction.

Since we may assume that $t[x]$ is not of the form $(\pi_{\mathbf{B}}\, t'[x])$, we have that $t[x]$ is either some constant $\mathbf{b}$, the variable x, or of the form $(\mathbf{P}\, t_1[x]\, t_2[x]\, t_3[x])$.

Note that $f(\mathbf{A}) \subseteq \mathbf{B}$, so f is not the identity—which means that $t[x] \not\equiv x$. If $t[x] \equiv \mathbf{b}$ for some type-α constant $\mathbf{b} \in \mathbf{B}$ then we're done. Otherwise, we are left with exactly one possibility: $t \equiv (\mathbf{P}\, t_1[x]\, t_2[x]\, t_3[x])$, where t_1 and t_2 are of type α, and t_3 is of type 0. Applying induction, we can repeat the above argument for both $t_1[x]$ and $t_2[x]$ to show that each of t_1 and t_2 is either a constant or of the form $(\mathbf{P}\, t_1'[x]\, t_2'[x]\, t_3'[x])$); this completes the proof. □

The concepts used in the proof of the following technical lemma also appear in subsequent proofs:

Lemma 3.8. *Let $\mathcal{G}$ be one of the five type-$(\alpha+2)$ constants A_1, $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, $\mathcal{T}_{\mathbf{C}}$, or $\mathcal{F}$. Let t be a type-0 term with type-α variables y and $x_1, \ldots, x_k$ free, as well as type-0 variables $n_1, \ldots, n_m$ free, and suppose that none of the five type-$(\alpha+2)$ constants occurs in t. Then the term:*

$$(\mathcal{G}\, (\lambda y.t[y; x_1, \ldots, x_k; n_1, \ldots, n_m]))$$

is equivalent to a term in which none of the five constants occurs.

Note that this lemma relies on our choice of model—in particular, on the fact that we include type-$(\alpha+1)$ constants only for open and clopen intervals in $\mathbf{A}$ and $\mathbf{B}$ and rectangles in $\mathbf{C}$.

Proof. Using the argument from Proposition 3.1, we may assume that $\mathbf{P}$ does not occur in t.

The constants, other than $\mathbf{P}$, that take inputs of type α and return output of type α are $\pi_{\mathbf{A}}$, $\pi_{\mathbf{B}}$, and $\langle \cdot, \cdot \rangle$. Not all syntactically-valid combinations involving these constants are semantically valid: for example, $\langle x_1, (\pi_{\mathbf{A}}\, y) \rangle$ is invalid (that is, it always returns "error"). Also, certain combinations are equivalent to simpler terms: for example, $(\pi_{\mathbf{B}}\, \langle y, x_2 \rangle)$ is equivalent to the simpler, normal-form, term x_2; while $(\pi_{\mathbf{A}}\, \mathbf{c})$, where $\mathbf{c} \in \mathbf{C}$ is a type-α constant, can be replaced by a type-α constant for an element of $\mathbf{a}$.

The only combinations involving these three constants that we need to consider are: (1) $(\pi_{\mathbf{A}}\, x)$ and (2) $(\pi_{\mathbf{B}}\, x)$, where x is a type-α variable; and (3) $\langle s, s' \rangle$, where s is a type-α constant $\mathbf{c}$, a type-α variable x, or $\pi_{\mathbf{A}} x$; and s' is either $\mathbf{c}'$, x', or $(\pi_{\mathbf{B}}\, x')$.

The only other constants that take inputs of type α are the type-$(\alpha + 1)$ constants A_0 and those constants representing open intervals/rectangles.

We use "v," below, to denote a type-α term that is (1) a constant, (2) a variable other than y, or (3) $\pi_{\mathbf{A}}$ or $\pi_{\mathbf{B}}$ applied to a variable other than y. Without loss of generality y occurs only in the following primitive subterms of t:

- **Primitive subterms involving A_0:**
  - $(A_0\, y\, v)$.
  - $(A_0\, (\pi_{\mathbf{A}}\, y)\, v)$ or $(A_0\, (\pi_{\mathbf{B}}\, y)\, v)$.

  (Note that we need not consider $(A_0 \langle s, t \rangle\, u)$, since it is equivalent to $(A_0\, s\, (\pi_{\mathbf{A}}\, u)) \wedge (A_0\, t\, (\pi_{\mathbf{B}}\, u))$.)
- **Primitive subterms involving a type-$(\alpha + 1)$ constant D:**
  - $(\mathbf{D}\, y)$.
  - $(\mathbf{D}\, (\pi_{\mathbf{A}}\, y))$ or $(\mathbf{D}\, (\pi_{\mathbf{B}}\, y))$.
  - $(\mathbf{D}\, \langle y, v \rangle)$ or $(\mathbf{D}\, \langle v, y \rangle)$. (Note that in the first case v cannot be $\pi_{\mathbf{A}}$ applied to a variable; in the second case v cannot be $\pi_{\mathbf{B}}$ applied to a variable.)
  - $(\mathbf{D}\, \langle (\pi_{\mathbf{A}}\, y), v \rangle)$ or $(\mathbf{D}\, \langle v, (\pi_{\mathbf{B}}\, y) \rangle)$. (The same note applies here.)

Since, by assumption, no type-$((\alpha+1) \to \sigma)$ constants occur in t we may assume that every type-α variable x that occurs in t is free in t. (If x were not free in t then it would be bound to a λ, and that λ-expression would have to be captured by a constant of type-$((\alpha+1) \to \sigma)$, for some type σ. But no such constant can occur in t.)

Although t is an arbitrary type-0 term, its dependence on the parameter y is captured by a finite tuple of 0's and 1's, showing for each primitive subterm occurring in t whether that primitive returns 0 or 1 on input y. If there are n such primitives occurring in t then there are only 2^n possible values for this tuple.

Let $P_1, \ldots, P_n$ list all such primitives occurring in t; let $\sigma \in 2^n$ be an arbitrary finite string of zeros and ones. As in the proof of Proposition 3.1, we write P_i^1 for P_i and P_i^0 for $\neg P_i$. Let s_σ be the type-0 term corresponding to the finite conjunction of primitives:

$$s_\sigma := \bigwedge_{1 \le i \le n} P_i^{\sigma(i)}.$$

(Here we define $\wedge$ and $\vee$ as in the proof of Proposition 3.1—s_σ is a type-0 term.)

Let t_σ, corresponding to s_σ, be derived from t by replacing each primitive P_i occurring in t with the constant $\mathbf{0}$, if $\sigma(i) = 0$, and the constant $\mathbf{1}$, if $\sigma(i) = 1$. Since every type-α variable in t is free in t, the only type-α variables occurring in s_σ are y and $x_1, \ldots, x_k$. So we can rewrite the type-0 term t as:

$$t[y; \overline{x}; \overline{n}] \equiv \dot{\bigvee_{\sigma \in 2^n}} (s_\sigma[y; \overline{x}] \wedge t_\sigma[\overline{x}; \overline{n}]).$$

(Note that the disjunction $\bigvee_{\sigma \in 2^n}$ is finite.)

Now y does not occur in t_σ, while the n_i's do not occur in s_σ. Also, for each assignment $\mathbf{y}$ to y and $\overline{\mathbf{x}}$ to $\overline{x}$ there is exactly one tuple σ such that $s_\sigma[\mathbf{y}; \overline{\mathbf{x}}] =_0 1$ (which is why the disjunction has a dot over it). So we have:

$$(\mathcal{G}\,(\lambda y.t[y; \overline{x}; \overline{n}])) \equiv \left(\mathcal{G}\,\left(\lambda y. \dot{\bigvee_{\sigma \in 2^n}} (s_\sigma[y; \overline{x}] \wedge t_\sigma[\overline{x}; \overline{n}])\right)\right).$$

Now our five type-$(\alpha + 2)$ constants distinguish only between the cases $t_\sigma =_0 0$ and $t_\sigma \neq_0 0$. Let $\tau \in 2^{(2^n)}$ be an arbitrary tuple, where $\tau(\sigma)$ is intended to correspond to the value of t_σ. Abusing notation somewhat, we also use "$\tau(\sigma)$" to denote the type-0 constant, either 0 or 1, corresponding to $\tau(\sigma)$. Using the above technique, we can move t_σ outside of $\mathcal{G}$:

$$\left(\mathcal{G}\left(\lambda y. \dot{\bigvee_{\sigma \in 2^n}} (s_\sigma[y;\overline{x}] \wedge t_\sigma[\overline{x};\overline{n}])\right)\right)$$

$$\equiv \dot{\bigvee_{\tau \in 2^{(2^n)}}} \left(\left[\bigwedge_{\sigma \in 2^n} t_\sigma^{\tau(\sigma)}[\overline{x};\overline{n}]\right] \wedge \left(\mathcal{G}\left(\lambda y. \dot{\bigvee_{\sigma \in 2^n}} (s_\sigma[y;\overline{x}] \wedge \tau(\sigma))\right)\right)\right)$$

$$\equiv \dot{\bigvee_{\tau \in 2^{(2^n)}}} \left(\left[\bigwedge_{\sigma \in 2^n} t_\sigma^{\tau(\sigma)}[\overline{x};\overline{n}]\right] \wedge \left(\mathcal{G}\left(\lambda y. \dot{\bigvee_{\sigma \in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y;\overline{x}]\right)\right)\right).$$

(As above, we use t_σ^1 to denote t_σ and t_σ^0 to denote $\neg t_\sigma$.) What is left inside $\mathcal{G}$ is a finite disjunction of $s_\sigma[y;\overline{x}]$ terms. In other words, in the rewritten term $\mathcal{G}$ is applied to a finite boolean combination of set constants and singletons $\{x_i\}$.

Recall that there are two sort of primitives in s_σ: the A_0 primitives, which relate y to type-α constants and variables, and the $\mathbf{D}$ primitives, which relate y to set constants for open intervals and open rectangles. For each tuple σ, we may write the $\mathbf{D}$-part of s_σ as the finite intersection of open intervals, open rectangles, and their complements. It makes sense to combine the A_0 primitives relating y to type-α constants and the $\mathbf{D}$ primitives—the constant parts of the s_σ's—into grids of the form:

- For $\mathbf{A}$, let $\mathbf{A}_1, \ldots, \mathbf{A}_{n_A} \subseteq_{\alpha+1} \mathbf{A}$ be (constants for) open intervals and $\mathbf{a}_1, \ldots, \mathbf{a}_{m_A} \in \mathbf{A}$ be distinct (constants for) elements not in any $\mathbf{A}_i$, such that:

  $$\mathbf{A} = \dot{\bigcup}_{1 \leq i \leq n_A} \mathbf{A}_i \,\dot{\cup}\, \{\mathbf{a}_i : 1 \leq i \leq m_A\},$$

  while any subset of $\mathbf{A}$, as well as the projection onto $\mathbf{A}$ of any subset of $\mathbf{C}$, described by the constant part of any s_σ, is the disjoint union of $\mathbf{A}_i$'s and $\{\mathbf{a}_i\}$'s.

- For $\mathbf{B}$, define $\mathbf{B}_1, \ldots, \mathbf{B}_{n_B} \subseteq_{\alpha+1} \mathbf{B}$ and $\mathbf{b}_1, \ldots, \mathbf{b}_{m_B} \in \mathbf{B}$ similarly.

- For $\mathbf{C}$, partition the set into open rectangles $\mathbf{A}_i \times \mathbf{B}_j$, points $\langle \mathbf{a}_i, \mathbf{b}_j \rangle$, and open segments $\mathbf{A}_i \times \{\mathbf{b}_j\}$ and $\{\mathbf{a}_i\} \times \mathbf{B}_j$, all sets pairwise disjoint, such that any subset of $\mathbf{C}$ described by the constant part of any s_σ is a disjoint union of certain of these sets.

We may now assume, without loss of generality, that any type-α or type-$(\alpha+1)$ constant occurring in t is one of the $\mathbf{A}_i$'s, $\mathbf{B}_i$'s, $\mathbf{a}_i$'s, or $\mathbf{b}_i$'s. This simplifies our cases.

(In the lemmas following we will repeatedly rewrite some fixed type-0 term t so that all set constants appearing in t are components of a grid of this form; a term's grid is its key feature.)

Note that a primitive of the form $(\mathbf{D}_i\,\langle y, v\rangle)$, where, for example, $\mathbf{D}_i$ is $\mathbf{A}_j \times \mathbf{B}_k$, is equivalent to the conjunction $(\mathbf{A}_j\,y) \wedge (\mathbf{B}_k\,v)$. We will address grid components $\mathbf{D}_i \subseteq \mathbf{C}$ through their projected components—in this case, $\mathbf{A}_j$ and $\mathbf{B}_k$.

We noted above that the term to which the type-$(\alpha+2)$ constant $\mathcal{G}$ is applied is of the form:

$$(\lambda y.t'[y;\overline{x}]) := \left(\lambda y. \dot{\bigvee_{\sigma \in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y;\overline{x}]\right),$$

and we considered the constant part of each s_σ separately from the part that depended on $\overline{x}$. The constant part, by assumption, identifies a particular component $\mathbf{D}_i$ of the grid, where $\mathbf{D}_i$ may be a subset of $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{C}$. Fix grid component $\mathbf{D}_i$, and consider those finite strings σ of zeros and ones such that:

$$(\lambda y.s_\sigma[y;\overline{x}]) \subseteq_{\alpha+1} \mathbf{D}_i.$$

Let $\Sigma(\mathbf{D}_i)$ be the set of all such σ's; then the portion of $(\lambda y.t'[y;\overline{x}])$ that lies within the component $\mathbf{D}_i$ is:

$$\left(\lambda y. \dot{\bigvee_{\sigma \in \Sigma(\mathbf{D}_i)}} s_\sigma[y;\overline{x}]\right).$$

So, within $\mathbf{D}_i$ the set defined by the term $(\lambda y.t'[y;\overline{x}])$ is, for each assignment $\overline{\mathbf{x}}$ of $\overline{x}$, some combination of singletons $\{\mathbf{x}_i\}$ and complements of singletons $\mathbf{D}_i \setminus \{\mathbf{x}_i\}$.

That is, within each grid component $\mathbf{D}_i \subseteq \mathbf{A}$ or $\mathbf{B}$, and for each $\overline{\mathbf{x}}$, the set $Y' \subseteq_{\alpha+1} \mathbf{D}_i$ defined by the term $(\lambda y.t'[y;\overline{\mathbf{x}}])$ is empty, finite, or cofinite:

- If $\mathbf{D}_i$ is a singleton, then Y' is **finite** unless $\Sigma(\mathbf{D}_i)$ contains only σ's specifying either $\neg(A_0\,y\,x_j)$, where $\mathbf{D}_i =_{\alpha+1} \{x_j\}$, or $(A_0\,y\,x_j)$, where $\mathbf{D}_i \neq_{\alpha+1} \{x_j\}$—in which case it is **empty**.

- If $\mathbf{D}_i$ is an open interval $\subseteq \mathbf{A}$, then Y' is **empty** if $\Sigma(\mathbf{D}_i)$ contains only σ's specifying $(A_0\, y\, x_j)$, where $x_j \notin \mathbf{D}_i$.

  The set Y' is **finite** if Y' is not empty but $\Sigma(\mathbf{D}_i)$ contains only σ's of the sort listed in the previous case, together with one or more σ's specifying $(A_0\, y\, x_j)$, where $x_j \in \mathbf{D}_i$.

  Otherwise, it is **cofinite**.

If $\mathbf{D}_i \subseteq \mathbf{C}$—say $\mathbf{D}_i$ is $\mathbf{A}_j \times \mathbf{B}_k$—then we need to consider its two projected components $\mathbf{A}_j$ and $\mathbf{B}_k$ simultaneously.

We now consider each of the five type-$(\alpha + 2)$ constants in turn:

1. $\mathcal{G}$ **is** A_1**:** This is the easiest case, since A_1 distributes across disjunctions, so:

$$
\begin{aligned}
&(A_1\,(\lambda y.t'[y;\overline{x}])) =_0 1 \\
&\qquad \iff \text{ for some } \mathbf{D}_i,\ (A_1\,(\lambda y.t''[y;\overline{x}])) := \left(A_1 \left(\lambda y. \bigvee_{\sigma\in\Sigma(\mathbf{D}_i)} s_\sigma[y;\overline{x}]\right)\right) =_0 1.
\end{aligned}
$$

   Fix component $\mathbf{D}_i$ such that $\Sigma(\mathbf{D}_i)$ is non-empty. Suppose $\mathbf{D}_i \subseteq \mathbf{A}$; from the discussion above there are only two cases: $\mathbf{D}_i$ is a singleton, or $\mathbf{D}_i$ is an open interval. In either case, whether the set described by $(\lambda y.t''[y;\overline{x}])$ is empty can be determined by comparing the (finitely many) x_j's to $\mathbf{D}_i$. So we can eliminate the constant A_1

2. $\mathcal{G}$ **is** $\mathcal{T}_{\mathbf{A}}$**:** Assume that $(\lambda y.t'[y;\overline{x}])$ describes a subset of $\mathbf{A}$. The set Y described by the term $(\lambda y.t'[y;\overline{x}])$ is open if and only if (1) for every grid component $\mathbf{A}_i$, the set Y' defined by:

$$
\left(\lambda y. \dot{\bigvee_{\sigma\in\Sigma(\mathbf{D}_j)}} s_\sigma[y;\overline{x}]\right),
$$

   is either empty or cofinite, and (2) for every grid component $\{\mathbf{a}_i\}$, if the corresponding set $Y' \subseteq \{\mathbf{a}_i\}$ is not empty then the sets corresponding to the two grid components $\mathbf{A}_j$ and $\mathbf{A}_{j+1}$ adjacent to $\{\mathbf{a}_i\}$ are cofinite. (The idea behind the second requirement is that for the set Y to be open, any point $\mathbf{a}_i \in Y$ must be covered by an open interval contained in Y.)

Whether these conditions hold can be determined by comparing the (finitely many) x_j's to the finitely many $\mathbf{D}_i$'s, so we can eliminate the constant $\mathcal{T}_{\mathbf{A}}$.

3. **$\mathcal{G}$ is $\mathcal{T}_{\mathbf{B}}$:** This case is similar to the previous case.

4. **$\mathcal{G}$ is $\mathcal{T}_{\mathbf{C}}$:** This case is similar to the previous case, but more complicated since we must now consider the projection of a given component $\mathbf{D}_i \subseteq$ the product space $\mathbf{C}$ onto both $\mathbf{A}$ and $\mathbf{B}$. Here, the set $Y \subseteq \mathbf{C}$ is open if and only if (1) for every rectangular component $\mathbf{A}_j \times \mathbf{B}_k$, the projections of the set $Y' \subseteq \mathbf{A}_j \times \mathbf{B}_k$ onto both $\mathbf{A}$ and $\mathbf{B}$ are cofinite; (2) for every segment component $\{\mathbf{a}_j\} \times \mathbf{B}_k$ or $\mathbf{A}_j \times \{\mathbf{b}_k\}$, if the corresponding set Y' is not empty then the sets corresponding to the two rectangle components adjacent to the segment are cofinite; and (3) for every singleton component $\{\langle \mathbf{a}_j, \mathbf{b}_k \rangle\}$, if the corresponding Y' is non-empty then it is abutted by four cofinite subsets of segments: two with the same $\mathbf{A}$-projections and two with the same $\mathbf{B}$-projections.
   (The idea here is that for Y to be open, any point $\langle \mathbf{a}_j, \mathbf{b}_k \rangle \in Y$ must be covered by an open rectangle contained in Y, as must any open segment $\{\mathbf{a}_j\} \times \mathbf{B}_k$ or $\mathbf{A}_j \times \{\mathbf{b}_k\}$.)
   Although this case is more complicated (it deals with two dimensions rather than one), the argument is the same, so we can eliminate the constant $\mathcal{T}_{\mathbf{C}}$.

5. **$\mathcal{G}$ is $\mathcal{F}$:** Note that the set Y defined by $(\lambda y.t'[y;\overline{x}])$ is in $\mathcal{F}$ only if Y is an open rectangle $\subseteq \mathbf{C}$. So, $(\mathcal{F}\,(\lambda y.t'[y;\overline{x}])) =_0 1$ if and only if (1) there is some finite index set I such that $\bigcup_{i \in I} \mathbf{D}_i \in \mathcal{F}$; (2) for all $i \in I$, the set $Y' \subseteq \mathbf{D}_i$ is cofinite (since it cannot be finite, and by definition is not empty); and (3) in fact, for all $i \in$ the set Y' is $\mathbf{D}_i$.
   The last condition holds if and only if for each tuple $\sigma \in \Sigma(\mathbf{D}_i)$ specifying $\neg(A_0\,y\,x_j)$ for some $x_j \in \mathbf{D}_i$, there is another $\sigma \in \Sigma(\mathbf{D}_i)$ that permits or specifies $(A_0\,y\,x_j)$—and similarly for projections of y. In other words, if some tuple σ specifies that part of $\mathbf{D}_i$ should be excluded, then there must be another σ that fills that part in.
   Whether these conditions hold can be determined from finitely many comparisons of the x_j's to the projected components of the $\mathbf{D}_i$'s, so we can eliminate the constant $\mathcal{F}$.

□

We will reuse several of the concepts introduced in the previous lemma. For convenience, we formally define the following different types of grid components:

Definition 3.9. *An* ***open rectangle*** *contained in* $\mathbf{C}$ *is a set of the form* $\mathbf{A}' \times \mathbf{B}'$*, where* $\mathbf{A}' \subseteq \mathbf{A}$*,* $\mathbf{B}' \subseteq \mathbf{B}$*, and both* $\mathbf{A}'$ *and* $\mathbf{B}'$ *are open in their respective topologies.*

An ***open segment*** *contained in* $\mathbf{C}$ *is a set of the form* $\{\mathbf{a}'\} \times \mathbf{B}'$ *or* $\mathbf{A}' \times \{\mathbf{b}'\}$*, where* $\mathbf{a}' \in \mathbf{A}$*,* $\mathbf{b}' \in \mathbf{B}$*,* $\mathbf{A}' \subseteq \mathbf{A}$*, and* $\mathbf{B}' \subseteq \mathbf{B}$*, and* $\mathbf{A}'$ *and* $\mathbf{B}'$ *are open in their respective topologies.*

An ***open slice*** *contained in* $\mathbf{C}$ *is the full open segment* $\{\mathbf{a}'\} \times \mathbf{B}$ *or* $\mathbf{A} \times \{\mathbf{b}'\}$*.*

An ***open tube*** *contained in* $\mathbf{C}$ *is the full open rectangle* $\mathbf{A}' \times \mathbf{B}$ *or* $\mathbf{A} \times \mathbf{B}'$*.*

The definition of "open rectangle" used here is the abstract topological notion: $\mathbf{A}'$ and $\mathbf{B}'$ need not be intervals; in practice, however, we will choose our grids so that $\mathbf{A}'$ and $\mathbf{B}'$ almost always are.

Using Lemma 3.8, we can show that all sets in the model $\mathcal{M}$ are finite disjoint unions of grid components:

Lemma 3.10. *In the model* $\mathcal{M}$*, all type-*$(\alpha+1)$ *functionals* $X \subseteq_{\alpha+1} \mathbf{A}$ *or* $\mathbf{B}$ *are finite disjoint unions of singletons and open (possibly clopen) intervals.*

All type-$(\alpha+1)$ *functionals* $X \subseteq_{\alpha+1} \mathbf{C}$ *are finite disjoint unions of singletons, open segments, and open rectangles.*

Proof. Suppose $X \subseteq_{\alpha+1} \mathbf{A}$; then X is defined by a closed term:

$$(\lambda x^{\alpha}.(\mathbf{A}x) \wedge t[x]),$$

where t is a type-0 term with only x free. (As in the proofs of Lemmas 3.1 and 3.8, we use $\wedge$ and $\vee$ as abbreviations for the appropriate numerical expression.)

Apply Lemma 3.8 inductively to t to get an equivalent term t' that contains no type-$(\alpha+2)$ constants and hence no type-α variables, and in which all type-$(\alpha+1)$ constants are disjoint $\mathbf{A}_i$'s contained in $\mathbf{A}$ and $\mathbf{B}_j$'s contained in $\mathbf{B}$. Then x occurs in subterms of t' only of the form:

- $(A_0\, x\, \mathbf{a}_i)$, where $\mathbf{a}_i$ is a type-α constant; or
- $(\mathbf{A}_i\, x)$, where $\mathbf{A}_i$ is a type-$(\alpha+1)$ constant for an open interval $\subseteq \mathbf{A}$.

Both of these subterms return, for each x, either 0 or 1. There are only finitely many such subterms, and for each tuple of outputs from these subterms the term t' returns either 0 or 1.

So X is a finite boolean combination of open intervals and singletons. Note that the complement of an open interval $\mathbf{U}$ is the disjoint union of 0, 1, or 2 open intervals and 0, 1, or 2 singletons (depending on whether $\mathbf{U}$ is also closed, half-closed, or not closed). So X is a disjoint union of open intervals and singletons. The proof for $X \subseteq_{\alpha+1} \mathbf{B}$ is similar.

Finally, suppose $X \subseteq_{\alpha+1} \mathbf{C}$; then X is defined by a closed term:

$$(\lambda x.(\mathbf{C}x) \wedge t[x]).$$

Applying Lemma 3.8, as in the previous case, we may assume that t contains no type-α variables and that the only type-$(\alpha+1)$ constants occurring in t are disjoint $\mathbf{A}_i$'s and $\mathbf{B}_j$'s. Then x occurs in subterms of t only of the form:

- $(A_0\, (\pi_{\mathbf{A}}\, x)\, \mathbf{a}_i)$, where $\mathbf{a}_i$ is a type-α constant;
- $(A_0\, (\pi_{\mathbf{B}}\, x)\, \mathbf{b}_i)$, where $\mathbf{b}_i$ is a type-α constant;
- $(\mathbf{A}_i\, (\pi_{\mathbf{A}}\, x))$, where $\mathbf{A}_i$ is a type-$(\alpha+1)$ constant for an open interval $\subseteq \mathbf{A}$; or
- $(\mathbf{B}_j\, (\pi_{\mathbf{B}}\, x))$, where $\mathbf{B}_j$ is a type-$(\alpha+1)$ constant for an open interval $\subseteq \mathbf{B}$.

So X is a finite boolean combination of singletons $\{\langle \mathbf{a}_i, \mathbf{b}_j \rangle\}$, slices $\{\mathbf{a}_i\} \times \mathbf{B}$ and $\mathbf{A} \times \{\mathbf{b}_i\}$, and open tubes $\mathbf{A}_i \times \mathbf{B}$ and $\mathbf{A} \times \mathbf{B}_i$—and hence a disjoint union of singletons, open segments, and rectangles. □

In particular, the only open sets in our model are finite disjoint unions of open intervals or open rectangles. Note that the three families $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, and $\mathcal{T}_{\mathbf{C}}$ trivially satisfy the axioms to be a topology except perhaps closure under arbitrary unions.

Lemma 3.11. *In $\mathcal{M}$, $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, and $\mathcal{T}_{\mathbf{C}}$ are closed under arbitrary unions—e.g., if $\mathcal{H} \in \mathcal{M}$, $\mathcal{H} \subseteq \mathcal{T}_{\mathbf{A}}$ then the set:*

$$\bigcup \mathcal{H} := \{x^\alpha : \exists X^{\alpha+1} \in \mathcal{H}\ (x \in X)\}$$

exists and is a finite disjoint union of open (possibly clopen) intervals and singletons.

So $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, and $\mathcal{T}_{\mathbf{C}}$ are topologies.

Proof. We prove the lemma first for families $\mathcal{H}$ of subsets of $\mathbf{A}$. The proof for families of subsets of $\mathbf{B}$ is similar; however, the proof for families of subsets of $\mathbf{C}$ is more complicated.

Let:

$$(\lambda X.\, ((X \subseteq_{\alpha+1} \mathbf{A}) \wedge t[X]))$$

be a term defining the family $\mathcal{H}$. As in the previous proof, partition the set and atom constants occurring in t into a grid generated by the finitely many open intervals $\mathbf{A}_1, \ldots, \mathbf{A}_{M_A}$, $\mathbf{B}_1, \ldots, \mathbf{B}_{M_B}$ and the finitely many singletons $\{\mathbf{a}_1\}, \ldots \{\mathbf{a}_{N_A}\}$, $\{\mathbf{b}_1\}, \ldots \{\mathbf{b}_{N_B}\}$.

By Lemma 3.10, every set constant $\mathbf{D}$ appearing in this term defining $\mathcal{H}$ is a disjoint union of components of the form $\mathbf{A}_i$ or $\{\mathbf{a}_i\}$, if $\mathbf{D} \subseteq_{\alpha+1} \mathbf{A}$; $\mathbf{B}_i$ or $\{\mathbf{b}_i\}$, if $\mathbf{D} \subseteq_{\alpha+1} \mathbf{B}$; or $\mathbf{A}_i \times \mathbf{B}_j$, $\{\mathbf{a}_i\} \times \mathbf{B}_j$, $\mathbf{A}_i \times \{\mathbf{b}_j\}$, or $\{\langle \mathbf{a}_i, \mathbf{b}_j \rangle\}$, if $\mathbf{D} \subseteq_{\alpha+1} \mathbf{C}$. By rewriting t, we may assume without loss of generality that these are the only type-α and type-$(\alpha+1)$ constants occurring in t.

Suppose $X_0 \in \mathcal{H}$ is an arbitrary subset of $\mathbf{A}$ defined by a closed term $(\lambda x.s[x])$. We will show that if $X_0 \cap \mathbf{A}_i \neq \emptyset$ then $\bigcup \mathcal{H} \supseteq \mathbf{A}_i$, by showing that $\mathcal{H}$ contains sets covering $\mathbf{A}_i$. This will show that $\bigcup \mathcal{H}$ is a finite union of intervals $\mathbf{A}_i$ and singletons $\{\mathbf{a}_i\}$ from the grid.

We will do this by showing that there is a class of permutations φ on the type-α elements of $\mathcal{M}$, giving rise to a class of partial maps Φ on the type-$(\alpha+1)$ such that:

$$t[X_0] =_0 t[\Phi(X_0)],$$

so $X_0 \in \mathcal{H} \iff \Phi(X_0) \in \mathcal{H}$. (Note that φ and Φ are defined outside the model. Note also that Φ need not be a permutation—in general, the map Φ need not be defined for all inputs $Y \in \mathcal{M}$, because the set induced by φ on Y need not be in $\mathcal{M}$.)

We will also show that this class is large enough such that for every $x \in \mathbf{A}_i$ there is a Φ in the class with $x \in \Phi(X_0)$. We restrict the class so that φ and Φ preserve subterms t' of t, but we need to ensure that the class is large enough to allow us to cover $\mathbf{A}_i$. Wo do the same in the proof of Lemma 3.17, but in Lemmas 3.12 and 3.18 we will also want to cover the particular grid component by only finitely many $\Phi(X_0)$'s, which is even more difficult.

We will use induction to show that φ and Φ preserve subterms of t. Provided that φ does not map an element of one grid component into another, the base cases, where $t' \subset t$ is an A_0- or $\mathbf{D}$-primitive, trivially hold. (The "$\mathbf{D}$" in a $\mathbf{D}$-primitive is a grid component, and φ preserves grid components, while the fact that φ is a permutation means that it preserves A_0-primitives as well.)

The induction cases will be to handle the various type-$(\alpha+2)$ families for which constants are included in the model. The idea here is that a subterm $t' \subset t$ will have one free type-$(\alpha + 1)$ variable (for X_0); the rest of its free variables will be of type α. Suppose φ and Φ preserve $t'' \subset t'$, and suppose t' is:

$$(\mathcal{G}\,(\lambda y.t''[y; X; \overline{x}; \overline{n}])).$$

Then φ moves $\overline{\mathbf{x}}$ around and Φ takes X_0 to $\Phi(X_0)$, but y is not touched. So although, for all parameters $\mathbf{y}$, $\overline{\mathbf{x}}$, and $\overline{\mathbf{n}}$, we have:

$$t''[\varphi(\mathbf{y}); \Phi(X_0)/X; \varphi(\overline{\mathbf{x}})/\overline{x}; \overline{\mathbf{n}}/\overline{n}] =_0 t''[\mathbf{y}; X_0/X; \overline{\mathbf{x}}/\overline{x}; \overline{\mathbf{n}}/\overline{n}],$$

the set:

$$Y := (\lambda y.t''[y; X_0/X; \overline{\mathbf{x}}/\overline{x}; \overline{\mathbf{n}}/\overline{n}])$$

differs from:

$$\Phi(Y) = (\lambda y.t''[y; \Phi(X_0)/X; \varphi(\overline{\mathbf{x}})/\overline{x}; \overline{\mathbf{n}}/\overline{n}])$$

because y is no longer permuted by φ.

In general, we will want to show that for our class of partial maps Φ, $(\mathcal{G}Y) =_0 (\mathcal{G}\Phi(Y))$ for all of our type-$(\alpha + 2)$ constants $\mathcal{G}$ and every set Y corresponding to a subterm t' of t. Note that, since φ is a bijection, this trivially holds for $\mathcal{G} = A_1$.

For this lemma, we may let φ be in the class of all continuous deformations on $\mathbf{A}_i$ that are the identity outside of $\mathbf{A}_i$. Let Φ be the partial map on sets induced by φ, defined by:

$$x \in \Phi(X) \iff \varphi(x)^{-1} \in X.$$

Then $\Phi(X)$ is defined if and only if the set $\{x : \varphi^{-1}(x) \in X\}$ is in the model $\mathcal{M}$. For this lemma, the map Φ induced by the continuous deformation φ is a permutation on sets of $\mathcal{M}$.

Let t' be an arbitrary type-0 subterm of t; we will use induction on terms to show that:

$$t'[X_0/X; \overline{\mathbf{x}}/\overline{x}; \overline{\mathbf{n}}/\overline{n}] =_0 t'[\varphi(\Phi(X_0)/X; \varphi(\overline{\mathbf{x}})/\overline{x}; \overline{\mathbf{n}}/\overline{n}].$$

As noted above, the base cases are trivial:

- $t' \equiv (A_0\, x_i\, x_j)$: Since φ is a bijection, for every $\mathbf{x}_i$ and $\mathbf{x}_j$ we have that $(A_0\, \mathbf{x}_i\, \mathbf{x}_j) =_0 (A_0\, \varphi(\mathbf{x}_i)\, \varphi(\mathbf{x}_j))$.
- $t' \equiv (A_0\, \mathbf{a}_i\, x_j)$: Since φ preserves all grid components, including singletons, for every $\mathbf{x}_j$ we have that $(A_0\, \mathbf{c}_i\, \mathbf{x}_j) =_0 (A_0\, \mathbf{c}_i\, \varphi(\mathbf{x}_j))$.
- $t' \equiv (\mathbf{A}_i\, x_j)$: Since φ preserves all grid components, for every $\mathbf{x}_j$ we have that $(\mathbf{A}_i\, \mathbf{x}_j) =_0 (\mathbf{A}_i\, \varphi(\mathbf{x}_j))$.
- $t' \equiv (X\, x_i)$: By the definition of Φ, for every $\mathbf{x}_i$ we have that $(X_0\, \mathbf{x}_i) =_0 (\Phi(X_0)\, \varphi(\mathbf{x}_i))$.

As in the proof of Lemma 3.8, the general induction step is:

$$(\mathcal{G}\, (\lambda y. t'[y; X; \overline{x}; \overline{n}])),$$

where $\mathcal{G}$ is one of our five type-$(\alpha + 2)$ constants. Note that we cannot simply apply Lemma 3.8 here because t' has a free type-$(\alpha + 1)$ variable, X, to which we will assign either X_0 or $\Phi(X_0)$. However, as in Lemma 3.8 we can rewrite t' so that we need to consider only:

$$(\mathcal{G}\, (\lambda y. t'[y; X; \overline{x}])) :\equiv \left(\mathcal{G}\, \left(\lambda y. \dot{\bigvee_{\sigma \in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y; X; \overline{x}]\right)\right),$$

where each s_σ is a finite conjunction of primitives involving A_0, set and atom constants for components from the grid, and the set variable X.

The only potential problem with the permutation φ concerns X-primitives involving y and A_0-primitives involving y and type-α variables x_i. We have five cases:

- $\mathcal{G}$ **is** A_1**:** As in the proof of Proposition 3.4, the hypothesis holds because φ is a bijection. Suppose there is a $\mathbf{y}$ such that $t'[\mathbf{y}/y; X_0/X; \overline{\mathbf{x}}/\overline{x}] \neq_0 0$. Then, by induction hypothesis, $t'[\varphi(\mathbf{y})/y; \Phi(X_0)/X; \varphi(\overline{\mathbf{x}}/\overline{x})] \neq_0 0$. Letting $\mathbf{y}' = \varphi^{-1}(\mathbf{y})$, we get there there is a $\mathbf{y}'$ such that $t'[\mathbf{y}'/y; \Phi(X_0)/X; \varphi(\overline{\mathbf{x}}/\overline{x})] \neq_0 0$. The converse is similar.

- $\mathcal{G}$ **is** $\mathcal{T}_{\mathbf{A}}$**:** As in the proof of Lemma 3.8, a set $Y \subseteq \mathbf{A}$ defined by $(\lambda y.t'[y; X_0; \overline{x}])$ is open if and only if (1) $Y \cap \mathbf{A}_j$ is open, for each component $\mathbf{A}_j$, and (2) for every component $\{\mathbf{a}_j\}$ such that $\mathbf{a}_j \in Y$, Y contains a neighborhood of $\mathbf{a}_j$.

  Notice that these criteria are affected by the permutation φ only inside the component $\mathbf{A}_i$, and that φ preserves open sets as well as neighborhoods of $\mathbf{A}_i$'s endpoints (if any). So the hypothesis holds.

- $\mathcal{G}$ **is** $\mathcal{T}_{\mathbf{B}}$**:** The hypothesis holds trivially because φ is the identity on $\mathbf{B}$, and, since φ is a bijection outside of $\mathbf{B}$, there is some $\mathbf{y} \notin \mathbf{B}$ such that $t'[\mathbf{y}/y; X_0/X; \overline{\mathbf{x}}/\overline{x}] =_0 1$ if and only if there is some $\mathbf{y}' \notin \mathbf{B}$ such that $t'[\mathbf{y}'; \Phi(X_0)/X; \varphi(\overline{\mathbf{x}})/\overline{x}] =_0 1$.

- $\mathcal{G}$ **is** $\mathcal{T}_{\mathbf{C}}$**:** Let $Y \subseteq \mathbf{C}$ be the set defined by $(\lambda y.t'[y; X_0; \overline{x}])$. Note that Y is open if and only if (1) $Y \cap (\mathbf{A}_j \times \mathbf{B}_k)$ is open, for each such rectangular component of the product space; (2) the projection of $Y \cap (\{\mathbf{a}_j\} \times \mathbf{B}_k)$ onto $\mathbf{B}$ is open, and similarly for the projection of $Y \cap (\mathbf{A}_j \times \{\mathbf{b}_k\})$ onto $\mathbf{A}$; (3) every singleton component $\{\langle \mathbf{a}_j, \mathbf{b}_k \rangle\}$ is abutted by (relatively) open neighborhoods in the four adjacent segment components; and (4) every (relatively) open subset of a segment component is abutted by two open neighborhoods in its two adjacent rectangular components.

  (Notice that this case is more complicated than in the proof of Proposition 3.4, since we now have an arbitrary open set X_0 in addition to the grid components.) Recall that φ affects only the $\mathbf{A}$-dimension of the product space $\mathbf{C}$, leaving the $\mathbf{B}$-dimension alone. We have shown that the hypothesis holds for $\mathcal{T}_{\mathbf{A}}$, so the hypothesis holds for $\mathcal{T}_{\mathbf{C}}$ as well.

- $\mathcal{G}$ **is** $\mathcal{F}$**:** First, by refining the grid we may assume without loss of generality that we have chosen our grid components $\mathbf{A}_1, \ldots, \mathbf{A}_{M_A}$ and $\mathbf{B}_1, \ldots, \mathbf{B}_{M_B}$ such that if $Y \subseteq \mathbf{C}$ is in $\mathcal{F}$ but is not contained in the upper-right rectangle $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$, then Y is the disjoint union of certain grid components. So if $\mathbf{A}_i \neq \mathbf{A}_{M_A}$ then the hypothesis holds since φ is a bijection on $\mathbf{A}_i$ and the identity everywhere else.

  Suppose $\mathbf{A}_i = \mathbf{A}_{M_A}$, so $\mathbf{A}_i \times \mathbf{B}_{M_B}$ is the upper-right corner of the product set $\mathbf{C}$. Let $Y \subseteq \mathbf{C}$ be the set defined by $(\lambda y.t'[y; X_0; \overline{x}])$. The hypothesis holds (trivially) if $Y \cap (\mathbf{A}_i \times \mathbf{B}_{M_B}) = \emptyset$, since φ is a bijection on $\mathbf{A}_i$ and the identity everywhere else, and in this case $Y \in \mathcal{F}$ if and only if Y is the disjoint union of certain grid components.

  If Y intersects some other component as well as the upper-right corner then $Y \notin \mathcal{F}$, so (trivially) $\Phi(Y) \notin \mathcal{F}$, so the hypothesis holds.

  Finally, suppose that Y is contained in the upper-right corner $\mathbf{A}_i \times \mathbf{B}_{M_B}$. Then $Y \notin \mathcal{F}$, because by assumption t' contains no constants for proper subsets of $\mathbf{B}_{M_B}$—which means that the projection of $Y \cap (\mathbf{A}_i \times \mathbf{B}_{M_B})$ must be a finite or cofinite subset of $\mathbf{B}_{M_B}$, so $(\mathcal{F}\,(\lambda y.t'[y; X_0/X; \overline{\mathbf{x}}/\overline{x}])) =_0 (\mathcal{F}\,(\lambda y.t'[y; \Phi(X_0)/X_0; \varphi(\overline{\mathbf{x}})/\overline{x}])) =_0 0$.

This completes our proof for families $\mathcal{H}$ of subsets of the space $\mathbf{A}$; the proof for families of subsets of $\mathbf{B}$ is similar.

Now suppose that $\mathcal{H}$ is a family of subsets of the product space $\mathbf{C}$. We use the same idea as in the proof for families of subsets of $\mathbf{A}$, except now we would like to let φ be in the class of all permutations of the form:

$$(\varphi\,\langle x, y\rangle) =_\alpha \langle(\varphi_{\mathbf{A}}\,x), (\varphi_{\mathbf{B}}\,y)\rangle,$$

where $\varphi_{\mathbf{A}}$ and $\varphi_{\mathbf{B}}$ are continuous deformations on some $\mathbf{A}_i$ and $\mathbf{B}_j$, respectively, and the identity everywhere else. Note that, as in the previous cases, the induced map Φ is a permutation on sets of the model $\mathcal{M}$.

There are two new issues.

First, the permutation φ affects not only a single grid component, say $\mathbf{A}_i \times \mathbf{B}_j$, but also the entire open tubes $\mathbf{A}_i \times \mathbf{B}$ and $\mathbf{A} \times \mathbf{B}_j$. (If φ is defined for an open segment $\{\mathbf{a}_i\} \times \mathbf{B}_j$ or $\mathbf{A}_i \times \{\mathbf{b}_j\}$ then φ affects, respectively, the open tube $\mathbf{A} \times \mathbf{B}_j$ or $\mathbf{A}_i \times \mathbf{B}$.)

Second, the constant $\mathcal{F}$ has infinite support on the the upper-right corner $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$, so we cannot just continuously deform the open tubes $\mathbf{A}_{M_A} \times \mathbf{B}$ or $\mathbf{A} \times \mathbf{B}_{M_B}$. (Any such deformation may cause a set to enter or leave $\mathcal{F}$.) So we will require that φ be in the class described above—unless $\mathbf{A}_i = \mathbf{A}_{M_A}$ or $\mathbf{B}_j = \mathbf{B}_{M_B}$, which we will handle separately.

As before, we assume that the grid is sufficiently fine that if $Y \in \mathcal{F}$ then either $Y \subseteq (\mathbf{A}_{M_A} \times \mathbf{B}_{M_B})$ or Y is the disjoint union of grid components.

Now suppose $\mathbf{A}_i = \mathbf{A}_{M_A}$ or $\mathbf{B}_j = \mathbf{B}_{M_B}$; notice that both open tubes $\mathbf{A}_{M_A} \times \mathbf{B}$ and $\mathbf{A} \times \mathbf{B}_{M_B}$ are actually clopen. Recall that $\mathcal{F}$ consists of, for each $n \in \omega$, three clopen rectangles $(c_n, c_{n+1}) \times (c_n, c_{n+1})$, $(c_{n+1}, 1) \times (c_n, c_{n+1})$, and $(c_n, c_{n+1}) \times (c_{n+1}, 1)$. (The fourth rectangle, $(c_{n+1}, 1) \times (c_{n+1}, 1)$ is missing from $\mathcal{F}$.)

If $\mathbf{A}_i = \mathbf{A}_{M_A}$ or $\mathbf{B}_j = \mathbf{B}_{M_B}$, or both, then we take the class of all permutations φ that are the identity function outside of $\mathbf{A}_i \times \mathbf{B}_j$, and such that:

1. $\varphi_{\mathbf{A}}$ and $\varphi_{\mathbf{B}}$ are continuous deformations of some intervals (c_m, c_{m+1}) or $(c_{m+1}, 1)$, and (c_n, c_{n+1}) or $(c_{n+1}, 1)$, contained in $\mathbf{A}_i = \mathbf{A}_{M_A}$ and $\mathbf{B}_j = \mathbf{B}_{M_B}$, respectively;
2. For every $m, n \in \omega$, $\varphi_{\mathbf{A}}$ and $\varphi_{\mathbf{B}}$ may also swap any of the three rectangles $(c_m, c_{m+1}) \times (c_m, c_{m+1})$, $(c_m, c_{m+1}) \times (c_{m+1}, 1)$, or $(c_{m+1}, 1) \times (c_m, c_{m+1})$ for any of the three rectangles $(c_n, c_{n+1}) \times (c_n, c_{n+1})$, $(c_n, c_{n+1}) \times (c_{n+1}, 1)$, or $(c_{n+1}, 1) \times (c_n, c_{n+1})$.

If $\mathcal{H}$ contains some set X_0 intersecting the upper-right corner of the grid, $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$, then the first condition allows us to cover any element of $\mathcal{F}$ that is contained in $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$ by permutations $\Phi(X_0)$; the second condition allows us to apply the first condition to any other element of $\mathcal{F}$ that is contained in $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$, but which X_0 may not intersect. These two conditions allow us to show that if $\mathcal{H}$ contains an element X_0 that intersects $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$, then $\bigcup \mathcal{H} \supseteq (\mathbf{A}_{M_A} \times \mathbf{B}_{M_B})$. This completes our definition of the class of permutations φ.

Note that none of the permutations φ on $\mathbf{C}$ affects $\mathcal{F}$. Also, since all of the permutations are bijections, none affects A_1. It remains for us to show that φ does not affect $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, or $\mathcal{T}_{\mathbf{C}}$.

As before, the proof is by induction on terms. The base cases are the same as before; the induction step is:

$$(\mathcal{G}\,(\lambda y.t'[y;X;\overline{x}])) :\equiv \left(\mathcal{G}\left(\lambda y. \dot{\bigvee_{\sigma\in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y;X;\overline{x}]\right)\right),$$

where $\mathcal{G}$ is one of our five type-$(\alpha+2)$ constants. We have five cases:

- **$\mathcal{G}$ is A_1:** As before, the hypothesis holds because φ is a bijection.
- **$\mathcal{G}$ is $\mathcal{T}_{\mathbf{A}}$:** As before, a continuous deformation φ inside of an open tube corresponding to a grid component (or to a subinterval of a grid component) does not affect the criteria for a set $Y \subseteq \mathbf{A}$ defined by $(\lambda y.t'[y;X_0;\overline{x}])$ to be open.

  Now suppose φ swaps two clopen rectangles in $\mathcal{F}$, contained in $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$. Note that, in general, if $D_1, \ldots, D_n$ is a clopen partition of some set X then a set $U \subseteq X$ is open if and only if for all i, $U \cap D_i$ is open. So any such φ preserves open sets, and the hypothesis holds.
- **$\mathcal{G}$ is $\mathcal{T}_{\mathbf{B}}$:** Similar to the previous case.
- **$\mathcal{G}$ is $\mathcal{T}_{\mathbf{C}}$:** Similar to the previous two cases, but more complicated. Note that a set $Y \subseteq \mathbf{C}$ defined by $(\lambda y.t'[y;X_0;\overline{x}])$ consists only of finitely many open rectangles, open segments, and singletons. So Y is open if and only if the projections (onto $\mathbf{A}$ and $\mathbf{B}$, respectively) of all of its fibers are open. The previous two cases show that φ preserves openness in $\mathbf{A}$ and $\mathbf{B}$, so the hypothesis holds in this case as well.
- **$\mathcal{G}$ is $\mathcal{F}$:** As noted above, continuous deformations of open tubes other than $\mathbf{A}_{M_A} \times \mathbf{B}$ or $\mathbf{A} \times \mathbf{B}_{M_B}$, continuous deformations of open tubes corresponding to elements of $\mathcal{F}$ (contained in $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$), and maps swapping elements of $\mathcal{F}$ (contained in $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$) do not affect $\mathcal{F}$. So the hypothesis holds.

□

So the three families $\mathcal{T}_{\mathbf{A}}$, $\mathcal{T}_{\mathbf{B}}$, and $\mathcal{T}_{\mathbf{C}}$ are topologies. Finally, we show that the first two are compact (the third is not, as witnessed by its open cover $\mathcal{F}$):

Lemma 3.12. *If $\mathcal{H} \in \mathcal{M}$, $\mathcal{H} \subseteq \mathcal{T}_{\mathbf{A}}$ is an open cover of $\mathbf{A}$ then there is a finite subfamily $\langle U_n : n < N \rangle$ of $\mathcal{H}$ that also covers $\mathbf{A}$.*

By symmetry, the same holds for $\mathbf{B}$ and $\mathcal{T}_{\mathbf{B}}$.

Proof. Let $\mathcal{H}$ be an arbitrary open cover, and suppose $\mathcal{H}$ is defined by the closed term:

$$(\lambda X.(X \subseteq_{\alpha+1} \mathbf{A} \wedge t[X])).$$

As in the proof of Lemma 3.11, we may assume the only type-$(\alpha+1)$ constants occurring in t are $\mathbf{A}_1, \ldots, \mathbf{A}_{M_A}$, $\mathbf{B}_1, \ldots, \mathbf{B}_{M_B}$, $\{\mathbf{a}_1\}, \ldots, \{\mathbf{a}_{N_A}\}$, and $\{\mathbf{b}_1\}, \ldots, \{\mathbf{b}_{N_B}\}$, where these constants generate grids on the three spaces $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ with the same properties as in that lemma.

If we simply apply the proof of Lemma 3.11, we may require infinitely many sets in $\mathcal{H}$ to cover $\mathbf{A}$—the problem is that $\mathbf{A}$ has intervals without endpoints, and a continuous deformation of an open set strictly contained inside one of these intervals can get arbitrarily close to the missing endpoint, but cannot reach it. To prove this lemma we will need to define a larger class of permutations φ.

Now $\mathcal{H}$ covers $\mathbf{A}$, so for each $1 \leq i \leq N_A$ there is an open set $U_i \in \mathcal{H}$ with $\mathbf{a}_i \in U_i$; fix a finite sequence of U_i's. We have a finite subcover of the set $\{\mathbf{a}_1, \ldots, \mathbf{a}_{N_A}\}$, but it remains for us to cover the $\mathbf{A}_i$'s with finitely many elements of $\mathcal{H}$.

Suppose $\mathbf{A}_i$ has endpoints $\mathbf{a}_j$ and $\mathbf{a}_{j+1}$ in $\mathbf{A}$; then the continuous deformations φ we used in the proof of Lemma 3.11 suffice to show that we can cover $\{\mathbf{a}_j\} \cup \mathbf{A}_i \cup \{\mathbf{a}_{j+1}\}$ by finitely many elements of $\mathcal{H}$. The problem is covering an interval $\mathbf{A}_i$ with an endpoint not in $\mathbf{A}$. (Recall that $\mathbf{A}$ is a version of $\mathbb{Q} \cap (0,1)$, so it lacks irrational endpoints as well as 0 and 1.)

If one of $\mathbf{A}_i$'s endpoints is missing from $\mathbf{A}$ then to cover $\mathbf{A}_i$ by finitely many elements of $\mathcal{H}$ there has to be an element of $\mathcal{H}$ containing an open neighborhood of that missing point.

For example, to cover $\mathbf{A}_{M_A}$, the rightmost component, by finitely many sets, $\mathcal{H}$ must contain an open set with a neighborhood of 1. If a set $X_0 \in \mathcal{H}$ does not contain a neighborhood of 1, then no continuous deformation φ of $\mathbf{A}_{M_A}$ can induce $\Phi(X_0)$ to contain such a neighborhood.

The key is that $\mathcal{H}$ not only cannot distinguish between X_0 and a continuous deformation $\Phi(X_0)$ of X_0, as in Lemma 3.11, but also cannot distinguish between X_0 and $(X_0 \setminus \mathbf{J}) \cup \mathbf{I}$ where $\mathbf{I}$ is any clopen interval contained in $\mathbf{A}_i$, disjoint from and not adjacent to X_0, and $\mathbf{J}$ is any clopen interval contained in $X_0 \cap \mathbf{A}_i$.

Suppose $\mathbf{A}_i$ has one or two endpoints not in $\mathbf{A}$, and suppose that X_0 does not contain an interval adjacent to one of $\mathbf{A}_i$'s missing endpoints. We extend our class of permutations φ to include, along with continuous deformations of $\mathbf{A}_i$, maps that extract a small clopen subinterval from the inside of an open interval $U \subseteq X_0 \cap \mathbf{A}_i$, and place it adjacent to one of $\mathbf{A}_i$'s missing endpoints, so that the extracted clopen subinterval does not intersect and is not adjacent to X_0.

This larger class of permutations still preserves type-0 subterms of t because, as noted in Lemma 3.11, a set $Y \subseteq \mathbf{A}$ is open if and only if (1) $Y \cap \mathbf{A}_i$ is open, for all i, and (2) if $\mathbf{a}_i \in X$ then Y contains a neighborhood of $\mathbf{a}_i$. Clearly adding clopen subsets of $\mathbf{A}_i$ does not affect the latter criterion; and of course $(X_0 \setminus \mathbf{J}) \cup \mathbf{I}$ is open if and only if X_0 is open, since $\mathbf{I}$ and $\mathbf{J}$ are both closed and open.

So the induction argument from Lemma 3.11 works here as well. Note that the constant $\mathcal{F}$ could possibly distinguish between X_0 and $(X_0 \setminus \mathbf{J}) \cup \mathbf{I}$—if X_0, $\mathbf{I}$, and $\mathbf{J}$ were elements of the product space $\mathbf{C}$. But they are not, so the same argument from Lemma 3.11 holds here as well: for a set Y defined by $(\lambda y.t'[y; X_0; \overline{x}])$ to be in $\mathcal{F}$, it cannot be contained in the upper-right component $\mathbf{A}_{M_A} \times \mathbf{B}_{M_B}$; so it must be a union of grid components, and hence is unaffected by the permutation φ. □

From the above we conclude:

Proposition 3.13 (RCA_0^ω + (ATOMS))**.** *The axiom* $(\mathcal{A}_1)$ *does not prove that the product of two compact spaces is compact.*

Proposition 3.19 shows, in the other direction, that "the product of two compact spaces is compact" does not imply $(\mathcal{A}_2)$. Formally, as discussed in Section 3.4.5, there is a bit of weirdness in the fact that the theory RCA_0^ω + $(\mathcal{A}_1)$ does not prove the existence of many topologies (or even many sets). For example, in the model $\mathcal{N}$ defined in Proposition 3.1, all type-$(\alpha + 1)$ functionals have finite support (so all sets are finite or cofinite), and all type-$(\alpha+1)$ functionals differ from the identity on finite support. So every topological space in $\mathcal{N}$ is compact, since either it is finite or it has only cofinite open sets.

So the strength of the statement that the product of two compact spaces is compact lies strictly between that of $(\mathcal{A}_1)$ and that of $(\mathcal{A}_2)$.

3.4.3 Compact T_2 spaces

It is a theorem of ZFC that every compact T_2 space is both T_3 and T_4. (It is a theorem of ZF that every space that is both T_2 and T_4 is also T_3, since in every T_2 space points are closed.) One standard proof that every compact T_2 space is T_3 does not require the Axiom of Choice and goes through in RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$:

Proposition 3.14 (RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$)**.** *Every compact T_2 space is T_3.*

Proof. Let $(X, \mathcal{T})$ be a compact T_2 space. Then for every $x \in X$ and $y \in X$, there are open sets U and V such that $x \in U$, $y \in V$, and $U \cap V =_{\alpha+1} \emptyset$.

Let F be an arbitrary closed set and fix $y \notin F$. Form the collection $\mathcal{F} \subseteq \mathcal{T}$ defined by:

$$U \in \mathcal{F} \iff (U \in \mathcal{T}) \wedge (\exists V^{\alpha+1}(V \in \mathcal{T} \wedge y \in V \wedge V \cap U =_{\alpha+1} \emptyset)).$$

Note that this definition uses only A_1, A_2, and $\mathcal{T}$, so the theory RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$ proves that $\mathcal{F}$ exists. Since the space is T_2, $\mathcal{F}$ covers F.

The collection $\{X \setminus F\} \cup \mathcal{F}$ exists and is an open cover of X. By compactness, it has a finite subcover $\langle U_n : n < N + 1\rangle$; without loss of generality, assume that $U_N =_{\alpha+1} X \setminus F$. Then, for each $n < N$, there is an open $V_n \ni y$ such that $U_n \cap V_n =_{\alpha+1} \emptyset$.

Our goal is to form a finite sequence of such V_n's and then to show that their intersection is an open set containing y. Then $\bigcap\langle V_n : n < N\rangle$ and $\bigcup\langle U_n : n < N\rangle$ would separate y from F.

The existence of such a sequence $\langle V_n : n < N\rangle$ follows from quantifier-free induction, using A_2. To see this, note that the empty sequence $\langle V_m : m < 0\rangle$ trivially exists, and that, given $\langle V_m : m < n\rangle$, where $n < N$, we can define:

$$\langle V_m : m < n+1\rangle := \begin{cases} V, & \text{if } m =_0 n \\ V_m, & \text{otherwise,} \end{cases}$$

where V is an open set such that $V \ni y$ and $U_n \cap V =_{\alpha+1} \emptyset$. So, given a finite sequence up to $n < N$, the finite sequence up to $n+1$ exists. So we have:

$$\exists\langle V_m : m < 0\rangle\, \forall m < 0\, \Phi(\langle V_m : m < 0\rangle)$$

and, for every n:

$$\exists\langle V_m : m < n\rangle\, \forall m < n\, \Phi(\langle V_m : m < n\rangle) \\ \to \exists\langle V_m : m < n+1\rangle\, \forall m < n+1\, \Phi(\langle V_m : m < n+1\rangle),$$

where $\Phi(\langle V_m : m < n\rangle)$ is the quantifier-free formula:

$$(y \in V_m) \wedge (V_m \cap U_m =_{\alpha+1} \emptyset),$$

which uses A_1 to determine type-$(\alpha+1)$ equality. The existence of a type-$(\alpha+1)$ finite sequence can be asserted, without quantifiers, by using A_2, so quantifier-free induction proves the existence of the finite sequence $\langle V_m : m < n\rangle$ for all $n \leq N$. In particular, it proves it for $n = N$.

Now $V := \bigcap\langle V_n : n < N\rangle$ exists—it can be defined using the primitive recursion operator $\mathcal{R}_0$. And quantifier-free induction proves that it is open, completing our proof. □

In another standard proof, which uses the Axiom of Choice, one defines the family $\mathcal{F}$ to have exactly one U_n corresponding to each V_n containing y; this proof goes through in RCA_0^ω

+ (ATOMS) + $(\mathcal{A}_1)$ + QF-AC$^{\alpha,\alpha+1}$ + QF-AC$^{0,\alpha}$. (The first choice principle, QF-AC$^{\alpha,\alpha+1}$, proves the existence of $\mathcal{F}$; the second choice principle, QF-AC$^{0,\alpha}$, proves the existence of $\langle V_n : n < N \rangle$.) So the fact that every compact T_2 space is T_3 can be proved from $(\mathcal{A}_2)$ or from choice principles.

We also have the following proposition, suggested by Kenneth Kunen:

Proposition 3.15 (RCA$_0^\omega$ + (ATOMS) + $(\mathcal{A}_1)$)**.** *If every compact T_2 space is T_3, then every compact T_2 space is T_4.*

So, over RCA$_0^\omega$ + (ATOMS) + $(\mathcal{A}_1)$, the statement that compact T_2 spaces are T_3 is equivalent to the statement that they are T_4. In the next section we show that the first statement is implied by $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$ but does not imply $(\mathcal{A}_1)$; Proposition 3.15 says that the same is true of the second statement.

Proof. Fix compact T_2 space $(X, \mathcal{T}_X)$; we will show that this space is T_4. Fix non-empty closed subset F of X; fix some point $x_0 \in F$. We define the quotient space $(Y, \mathcal{T}_Y)$ by collapsing the set F to the point x_0. The quotient space is compact and T_2, so by hypothesis it is T_3. So x_0 can be separated from closed sets by disjoint open sets; by pulling these open sets back to $(X, \mathcal{T}_X)$, we have open sets separating F from any disjoint closed subset of X. This proves that $(X, \mathcal{T}_X)$ is T_4.

Formally, the set Y is $(X \setminus F) \cup \{x_o\}$, defined by the term:

$$(\lambda x^\alpha.((A_0\, x\, x_0) \vee ((X\, x) \wedge \neg(F\, x)))),$$

where we use the logical symbols $\vee$, $\wedge$, and $\neg$ to denote the corresponding boolean operation on type-0 objects, as in the proof of Proposition 3.1. The theory RCA$_0^\omega$ + (ATOMS) proves that the set Y exists.

We define the family $\mathcal{T}_Y$ on subsets of Y by the rule:

$$U \in \mathcal{T}_Y \iff (U \subseteq_{\alpha+1} Y) \wedge ([(x_0 \notin U) \wedge (U \in \mathcal{T}_X)] \vee [(x_0 \in U) \wedge (U \cup F \in \mathcal{T}_X)]).$$

The theory RCA$_0^\omega$ + + (ATOMS) + $(\mathcal{A}_1)$ proves that $\mathcal{T}_Y$ exists. (We need A_1 for "$\subseteq_{\alpha+1}$.")

The family $\mathcal{T}_Y$ is a topology:

- $\emptyset \in \mathcal{T}_Y$ since $x_0 \notin \emptyset$ and $\emptyset \in \mathcal{T}_X$; $Y \in \mathcal{T}_Y$ since $x_0 \in Y$ and $Y \cup F = X \in \mathcal{T}_X$.
- If $U, V \in \mathcal{T}_Y$, then we have four cases, depending on whether each of the two sets contains x_0:
  - $x_0 \notin U$ and $x_0 \notin V$: Then $U, V \in \mathcal{T}_X$. So $U \cap V \in \mathcal{T}_X$. Also, $x_0 \notin U \cap V$, so $U \cap V \in \mathcal{T}_Y$.
  - $x_0 \notin U$ but $x_0 \in V$: Then $U \in \mathcal{T}_X$ and $V \cup F \in \mathcal{T}_X$. Note that $U \cap F = \emptyset$, since $Y \cap F = \{x_0\}$. So $U \cap V =_{\alpha+1} U \cap (V \cup F) \in T_X$, and $x_0 \notin U \cap V$. So $U \cap V \in \mathcal{T}_Y$.
  - $x_0 \in U$ but $x_0 \notin V$: similar to previous case.
  - $x_0 \in U$ and $x_0 \in V$: Then $U \cup F, V \cup F \in \mathcal{T}_X$, which means that $(U \cup F) \cap (V \cup F) =_{\alpha+1} (U \cap V) \cup F \in \mathcal{T}_X$. Now $x_0 \in \mathcal{U} \cap V$, so $U \cap V \in \mathcal{T}_Y$.
- Finally, suppose $\mathcal{F} \subseteq \mathcal{T}_Y$. Define family $\mathcal{G} \subseteq \mathcal{T}_X$ by the rule:

  $$U \in \mathcal{G} \iff ((U \cap F =_{\alpha+1} \emptyset) \wedge (U \in \mathcal{F})) \vee ((F \subseteq_{\alpha+1} U) \wedge ((U \setminus F) \cup \{x_0\} \in \mathcal{F})).$$

  The theory RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ proves that $\mathcal{G}$ exists; that $\mathcal{G} \subseteq \mathcal{T}_X$ is clear from the definition of $\mathcal{T}_Y$. Since $\mathcal{T}_X$ is a topology, $\bigcup \mathcal{G}$ exists and is open. Note that either $\bigcup \mathcal{F} =_{\alpha+1} \bigcup \mathcal{G}$ or $\bigcup \mathcal{F} =_{\alpha+1} (\bigcup \mathcal{G} \setminus F) \cup \{x_0\}$, depending on whether some $U \in \mathcal{F}$ contains x_0.

  Either way, $\bigcup \mathcal{F}$ exists and (by definition) is in $\mathcal{T}_Y$.

The topology $\mathcal{T}_Y$ is T_2—the fact that $\mathcal{T}_X$ is T_2 proves that points $x, y \neq_\alpha x_0$ can be separated by disjoint open sets, while the fact that $\mathcal{T}_X$ is T_3 proves that x_0 can be separated from any other element of Y. To see that $\mathcal{T}_Y$ is compact, let $\mathcal{F} \subseteq \mathcal{T}_Y$ be an open cover of Y. Then the family $\mathcal{G}$ defined, above, from $\mathcal{F}$ is an open cover of X. Since X is compact, $\mathcal{G}$ has a finite subcover $\langle V_n : n < N \rangle$.

Define finite subcover $\langle U_n : n < N \rangle$ of $\mathcal{F}$ by the rule:

$$U_n :=_{\alpha+1} \begin{cases} V_n, & \text{if } V_n \cap F =_{\alpha+1} \emptyset \\ (V_n \setminus F) \cup \{x_0\}, & \text{otherwise.} \end{cases}$$

The definition can be made using A_1, so $\langle U_n : n < N\rangle$ exists, proving compactness.

By hypothesis, the space $(Y, \mathcal{T}_Y)$ is T_3. Fix arbitrary closed set $K \subseteq X$, disjoint from F. Then $K \subseteq Y$ and $x_0 \notin K$. Since $(Y, \mathcal{T}_Y)$ is T_3, there are disjoint open set $U, V \in \mathcal{T}_Y$ such that $x_0 \in U$ and $K \subseteq V$. So $V, U \cup F \in \mathcal{T}_X$, and these two open sets are disjoint. Trivially, $F \subseteq U \cup F$, so the arbitrary disjoint closed sets F and K are separated by disjoint open sets, proving that $(X, \mathcal{T}_X)$ is T_4. □

3.4.4 RCA_0^ω + $(\mathcal{E}_1)$ + (ATOMS) + $(\mathcal{A}_1)$ does not imply that every compact T_2 space is T_3

We again construct a minimal term model $\mathcal{M}$ (which is an ω-model) of RCA_0^ω + (ATOMS) + $(\mathcal{E}_1)$ + $(\mathcal{A}_1)$. In addition to the constant symbol μ, we include the following constants:

- Uncountably many type-α constants $\mathbf{d}_i$, for $i \in 2^{\aleph_0}$, representing the elements of our topological space.

- One type-$(\alpha+1)$ constant $\mathbf{A}$ representing the characteristic function of an uncountable set. The constant $\mathbf{A}$ is intended to represent a copy of the closed unit square $[0,1]^2$. We can define $\mathbf{A}$ by the term $(\lambda x^\alpha.1)$—$\mathbf{A}$ is the set of all atoms—but it is nice to have a constant for this.

  Although $\mathbf{A}$ is intended to represent a copy of $[0,1]^2$, we will not include copies of the component spaces $[0,1]$ in $\mathcal{M}$, so $\mathbf{A}$ will not be a product space.

- For each open rectangle in the standard topology on $[0,1]^2$, a type-$(\alpha+1)$ constant for a copy $\subseteq \mathbf{A}$. These constants, together with $\mathbf{A} \setminus \mathbf{F}$, will be the basic open sets for the space $(\mathbf{A}, \mathcal{T})$.

- A type-$(\alpha+1)$ constant $\mathbf{F}$ representing the subset F of $[0,1]^2$ satisfying:

  - For all $x \in [0,1]$, there is exactly one $y \in [0,1]$ such that $\langle x, y\rangle \in F$; for all $y \in [0,1]$, there is exactly one $x \in [0,1]$ such that $\langle x, y\rangle \in F$. (Recall that $\mathbf{A}$ is not a product space: here x and y are elements of $[0,1]$ and $\langle x, y\rangle$ is an element of $[0,1]^2$.)

– The set F is dense in $[0,1]^2$. Further, we require that if R is an open rectangle (for which we have added a constant symbol $\mathbf{R}$) then $|R \cap F| = 2^{\aleph_0}$.

For example, we can apply the Continuum Hypothesis, so that $|[0,1]^2| = \aleph_1$, and then use transfinite induction to place uncountably many elements of F inside each rectangle with rational corners.

- A type-(α+2) constant $\mathcal{T}$, with the interpretation that for every type-(α+1) functional U in the model, $U \in \mathcal{T}$ if and only if U corresponds to an open set in the topology generated by the open rectangles and the set $\mathbf{A} \setminus \mathbf{F}$. (As in Section 3.4.2, we will show that the only sets in $\mathcal{M}$ are finite disjoint unions of components: open rectangles, open segments, singletons, and the intersections of $\mathbf{F}$ or $\mathbf{A} \setminus \mathbf{F}$ with open rectangles and open segments.)

The term model consists of (equivalence classes of) closed terms involving the various constants listed above.

Assuming (in $\mathcal{M}$) that the family $\mathcal{T}$ is a topology, it is clear that $\mathcal{T}$ is T_2, since any two distinct points may be separated by disjoint open rectangles. However, $\mathcal{T}$ is not T_3 since the set $\mathbf{F}$ is dense in the standard topology on $[0,1]^2$, so $\mathbf{F}$ cannot be separated from any $x_0 \notin \mathbf{F}$. (Suppose U is an open neighborhood of x_0; then $\mathrm{cl}(U) \cap \mathbf{F} \neq \emptyset$, so any open set covering $\mathbf{F}$ intersects U.)

It remains for us to show that the model $\mathcal{M} \models$ "$\mathcal{T}$ is a topology" and that $\mathcal{M} \models$ "$\mathcal{T}$ is compact.". The proofs are similar to those from Section 3.4.2.

As in Section 3.4.2, we have:

Lemma 3.16. *In the model $\mathcal{M}$, all type-$(\alpha+1)$ functionals $X \subseteq_{\alpha+1} \mathbf{A}$ are finite disjoint unions of singletons; open segments; open rectangles; the intersections of open segments with $\mathbf{F}$ or $\mathbf{A} \setminus \mathbf{F}$; and the intersections of open rectangles with $\mathbf{F}$ or $\mathbf{A} \setminus \mathbf{F}$.*

Proof. Let $(\lambda x.t[x])$ define an arbitrary set $X \subseteq \mathbf{A}$. As in the proof of Lemma 3.8, form a grid refining all constants (other than $\mathbf{F}$) occurring in t. Here the space $\mathbf{A}$ corresponds to the two-dimensional square $[0,1]^2$, so the grid will consist of open rectangles, open segments, and singletons. However, $\mathbf{A}$ is not a product space—it has no components spaces in $\mathcal{M}$—so we do not have the projections of these components.

As in the proof of Lemma 3.8, we may assume that t is written so that we need consider only subterms of t of the form:

$$(\mathcal{G}\,(\lambda y.t'[y;\overline{x}])) :\equiv \left(\mathcal{G}\,\left(\lambda y. \dot{\bigvee_{\sigma\in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y;\overline{x}]\right)\right),$$

where "$\mathcal{G}$" refers either to A_1 or $\mathcal{T}$. (Note that we have only two type-$(\alpha+2)$ constants here.) As in that proof, we will show that the type-$(\alpha+2)$ constant represented by $\mathcal{G}$ can be eliminated and the term $(\mathcal{G}\,(\lambda y.t'[y;\overline{x}]))$ replaced by an equivalent term in which the type-$(\alpha+2)$ constant represented by $\mathcal{G}$ does not occur. Let Y' be the set defined (for each value $\overline{\mathbf{x}}$) by the term $(\lambda y.t'[y;\overline{\mathbf{x}}/\overline{x}])$.

As in the proof of Lemma 3.8, we assume (by induction) that for each non-singleton component $\mathbf{A}_i$ of the grid, one of the following holds: either $Y' \cap \mathbf{A}_i$ or $Y' \cap (\mathbf{A}_i \setminus \mathbf{F})$ is finite, or $Y' \cap \mathbf{A}_i$ is cofinite in $\mathbf{A}_i$, or $Y' \cap (\mathbf{A}_i \setminus \mathbf{F})$ is cofinite in $\mathbf{A}_i \setminus \mathbf{F}$.

The A_1 case is the same as in the proof of Lemma 3.8—one can determine whether the intersection of Y' with some component is nonempty via a finite boolean expression not involving y.

For the $\mathcal{T}$ case, note that whether the set Y' defined by $(\lambda.t'[y;\overline{\mathbf{x}}/\overline{x}])$ is open depends on two things: whether the intersection of Y' with each of the grid's open rectangles is open; and whether Y' contains a neighborhood of each of its intersections with the grid's open segments and singletons.

Determining whether the intersection of Y' with one of the grid's open rectangles $\mathbf{A}_i$ is open is the same as in the proof of Lemma 3.8: the intersection is open if and only if either $Y' \cap \mathbf{A}_i$ or $Y' \cap (\mathbf{A}_i \setminus \mathbf{F})$ is cofinite in their respective sets. The second part is more complicated.

Suppose $\{\mathbf{a}_i\}$ is one of the grid's singletons, and suppose $\mathbf{a}_i \in Y'$. If $\mathbf{a}_i \in \mathbf{F}$ then, for Y' to be open, the intersection of Y' with each of the two, three, or four open segments adjacent to $\{\mathbf{a}_i\}$ must be cofinite. If $\mathbf{a}_i \notin \mathbf{F}$ then, for Y' to be open, for each open segment $\mathbf{A}_j$ adjacent to $\{\mathbf{a}_i\}$, either $Y' \cap \mathbf{A}_j$ must be cofinite in $\mathbf{A}_j$ or $Y' \cap (\mathbf{A}_j \setminus \mathbf{F})$ must be cofinite in $\mathbf{A}_j \setminus \mathbf{F}$. (If $\mathbf{a}_i \notin Y'$ then $\{\mathbf{a}_i\}$ puts no restriction on its adjacent open segments.)

Now suppose $\mathbf{A}_j$ is one of the grid's open segments, and suppose $Y' \cap \mathbf{A}_j$ is cofinite in $\mathbf{A}_j$. For Y' to be open, the intersection of Y' with the one or two open rectangles adjacent to $\mathbf{A}_j$, in the grid, must be cofinite in the respective sets.

If, instead, $Y' \cap \mathbf{A}_j \cap (\mathbf{A} \setminus \mathbf{F})$ is cofinite in $\mathbf{A}_j \setminus \mathbf{F}$, then, for Y' to be open, for each of the one or two adjacent open rectangles, either its intersection with Y' or its intersection with $Y' \cap (\mathbf{A} \setminus \mathbf{F})$ must be cofinite in its respective set.

Finally, if either $Y' \cap \mathbf{A}_i$ or $Y' \cap (\mathbf{A}_i \setminus \mathbf{F})$ is finite, then Y' is not open.

Note that all of these cases can be determined via a finite boolean expression not involving y, so $\mathcal{T}$ can be eliminated. So t is equivalent to a finite boolean expression involving grid components and $\mathbf{F}$, proving the lemma. □

And:

Lemma 3.17. *In $\mathcal{M}$, $\mathcal{T}$ is closed under arbitrary unions: if $\mathcal{H} \in \mathcal{M}$, $\mathcal{H} \subseteq \mathcal{T}$ then the set:*

$$\bigcup \mathcal{H} := \{x^\alpha : \exists X^{\alpha+1} \in \mathcal{H}\ (x \in X)\}$$

exists in $\mathcal{M}$.

Proof. Let $(\lambda X.t[X])$ define $\mathcal{H}$; let $X_0 \in \mathcal{H}$ be arbitrary. As in the proof of Lemma 3.11, we will define a class of permutations φ on the type-α elements of $\mathcal{M}$ that give rise to partial maps Φ on $\mathcal{M}$'s type-$(\alpha+1)$ elements, such that for all type-0 subterms $t'[X;\overline{x}]$ of t, and for all values $\overline{\mathbf{x}}$,

$$t'[X_0/X;\overline{\mathbf{x}}/\overline{x}] =_0 t'[\Phi(X_0)/X;\varphi(\overline{\mathbf{x}})/\overline{x}].$$

In particular, we will have $t[\Phi(X_0)] =_0 t[X_0]$, so $\Phi(X_0) \in \mathcal{H}$. The class of permutations φ will be large enough that if X_0 intersects $\mathbf{A}_i \setminus \mathbf{F}$, for some component $\mathbf{A}_i$ of the finite grid generated from the constants occurring in the term t, then the collection of all $\Phi(X_0)$'s covers $\mathbf{A}_i \setminus \mathbf{F}$; and similarly, if X_0 intersects $\mathbf{A}_i \cap \mathbf{F}$ then the collection of all $\Phi(X_0)$'s covers $\mathbf{A}_i$.

Form a grid of open rectangles, open segments, and singletons from t, as in the previous proof. Since $\mathbf{F}$ intersects each horizontal and vertical slice exactly once, by refining the grid we may assume without loss of generality that $\mathbf{F}$ does not intersect any of the grid's open segments. (If $\mathbf{F}$ intersects some open segment $\mathbf{A}_i$ at some point $\mathbf{c} \in \mathbf{A}_i$, then we add an orthogonal slice that intersects $\mathbf{A}_i$ at $\mathbf{c}$. Since $\mathbf{F}$ intersects each slice exactly once, $\mathbf{F}$ does not intersect any open segment of the two orthogonal slices.) So $\mathbf{F}$ intersects only the grid's open rectangles and (possibly) singletons.

Our permutations φ will be similar in spirit to the continuous deformations used in the second part of the proof of Lemma 3.11. However, we must now preserve membership in the set $\mathbf{F}$. Fix grid component $\mathbf{A}_i$—either an open rectangle or an open segment. Fix open set $X_0 \in \mathcal{H}$ intersecting $\mathbf{A}_i$; by Lemma 3.16, X_0 is a union of finitely many components: open rectangles and intersections of open rectangles with $\mathbf{A} \setminus \mathbf{F}$. We first describe the continuous deformation φ', and then use φ' to present the restrictions on φ.

As in the proof of Lemma 3.11, we require that φ' be the identity outside of the one or two (depending on whether $\mathbf{A}_i$ is an open segment or an open rectangle) open tubes intersecting $\mathbf{A}_i$. Inside the tube(s), we let φ' be a continuous deformation of X_0, as in the proof of Lemma 3.11, except with an additional requirement. The boundaries of the finitely many rectangles that make up X_0 consist of only finitely many open segments (and intersections of open segments with $\mathbf{A} \setminus \mathbf{F}$) and singletons. Our additional requirement is that φ' must take each such segment or singleton that intersects $\mathbf{F}$ to another segment or singleton that intersects $\mathbf{F}$, while taking each such segment or singleton that does not intersect $\mathbf{F}$ to another segment or singleton that does not intersect $\mathbf{F}$.

(The idea here is that the intersection of $\mathbf{F}$ with any open rectangle is infinite, but the intersection of $\mathbf{F}$ with a segment or singleton is finite (or empty). By refining our grid, we end up with a grid that intersects $\mathbf{F}$ only at singletons (possibly) and at rectangles (definitely).

When deforming the set X_0 to get $\Phi'(X_0)$, we want to preserve the intersections with $\mathbf{F}$ of the boundaries of X_0's components.)

Note that the class of permutations φ' suffices to cover $\mathbf{A}_i$ by sets $\Phi'(X_0)$, since $\mathbf{F}$ is dense in $\mathbf{A}$, while X_0's boundary consists of only finitely many open segments and singletons. However, φ' does not preserve all subterms of t; in particular, the primitive subterm $(\mathbf{F}\, x)$ need not equal $(\mathbf{F}\, \varphi(x))$.

The permutation φ agrees with φ' on the boundaries of X_0's component rectangles, and is the identity outside the open tube(s) intersecting $\mathbf{A}_i$. Everywhere else, the permutation φ differs from φ' by preserving membership in $\mathbf{F}$:

$$x \in \mathbf{F} \iff \varphi(x) \in \mathbf{F},$$

and we require that $\Phi(X_0) = \Phi'(X_0)$. (So φ preserves the interiors and exteriors of X_0.)

This last condition ensures that (in general) φ is not a continuous deformation of X_0—the interiors of the open rectangles that make up X_0 are jumbled to preserve membership in $\mathbf{F}$. (We can jumble the interiors because $F \cap R$ and $R \setminus R$ have the same cardinality for every open rectangle R.) The partial map Φ induced by φ is not, in general, a permutation on $\mathcal{M}$: if $Y \subseteq_{\alpha+1} X_0$ there is no requirement that $\Phi(Y)$ satisfy Lemma 3.16. Since $\Phi(X_0) = \Phi'(X_0)$, the class of all such permutations φ suffices to cover $\mathbf{A}_i$ by deformations of X_0.

As in the proof of Lemma 3.11, we prove that φ preserves type-0 subterms of t by induction on terms. As before, the base cases are:

- $(X_0 \mathbf{x}_j)$, which by definition equals $(\Phi(X_0)\, \varphi(\mathbf{x}_j))$;
- $(A_0\, \mathbf{x}_j\, \mathbf{x}_k)$, which trivially equals $(A_0\, \varphi(\mathbf{x}_j)\, \varphi(\mathbf{x}_k))$;
- $(\mathbf{F}\, \mathbf{x}_j)$, which equals $(\mathbf{F}\, \varphi(\mathbf{x}_j))$, since φ preserves membership in $\mathbf{F}$;
- $(\mathbf{A}_j\, \mathbf{x}_k)$, where $\mathbf{A}_j$ is some grid component, which (since φ is the identity outside of the open tube(s) intersecting $\mathbf{A}_i$) equals $(\mathbf{A}_j\, \varphi(\mathbf{x}_k))$; and
- $(A_0\, \mathbf{a}_j\, \mathbf{x}_k)$, where $\{\mathbf{a}_i\}$ is some grid component, which equals $(A_0\, \mathbf{a}_j\, \varphi(\mathbf{x}_k))$.

As before, the induction step is:

$$(\mathcal{G}\,(\lambda y.t'[y;X;\overline{x}])) :\equiv \left(\mathcal{G}\left(\lambda y. \dot{\bigvee_{\sigma\in 2^n \wedge \tau(\sigma)=1}} s_\sigma[y;X;\overline{x}]\right)\right),$$

where $\mathcal{G}$ is now one of the two type-$(\alpha+2)$ constants A_1 and $\mathcal{T}$. The case where $\mathcal{G} \equiv A_1$ is still trivial, since φ is a bijection.

The remaining case is where $\mathcal{G} \equiv \mathcal{T}$. For each tuple $\overline{\mathbf{x}}$ of parameters, let Y' be the set defined by $(\lambda y.t'[y;X_0/X;\overline{\mathbf{x}}/\overline{x}])$. As in the proof of Lemma 3.11, we want to show that:

$$Y' \in \mathcal{T} \iff \Phi(Y') \in \mathcal{T},$$

where $\Phi(Y') = (\lambda y.t'[y;\Phi(X_0)/X;\varphi(\overline{\mathbf{x}})/\overline{x}])$.

Note that Y' is some finite boolean combination of X_0, $\mathbf{F}$, singletons $\{\mathbf{x}_i\}$, and grid components for t; so $\Phi(Y')$ is the same finite boolean combination, but with $\Phi(X_0)$ in place of X_0 and $\{\varphi(\mathbf{x}_i)\}$ in place of $\{\mathbf{x}_i\}$. So $\Phi(Y') \in \mathcal{M}$, since $\Phi(X_0) \in \mathcal{M}$.

Now the set Y' is open if and only if:

1. $\mathbf{A}_j \cap Y'$ is open, for all open-rectangle grid components $\mathbf{A}_j$;
2. the projection of $\mathbf{A}_j \cap Y'$, where $\mathbf{A}_j$ is an open-segment component, onto the relevant coordinate axis is open;
3. if $\{\mathbf{a}_j\} \cap Y' \neq_{\alpha+1} \emptyset$, for some singleton grid component $\{\mathbf{a}_j\}$, then Y' contains two, three, or four, as relevant:
   (a) adjacent open segments $\mathbf{A}_i$, if $\mathbf{a}_j \in \mathbf{F}$; or
   (b) adjacent open segments $\mathbf{A}_i$ or $\mathbf{A}_i \setminus \mathbf{F}$, otherwise;
4. if $\mathbf{A}_j \cap Y' \neq_{\alpha+1} \emptyset$ for some open-segment grid component $\mathbf{A}_j$, then for each subsegment of $\mathbf{A}_j \cap Y'$, the set Y' contains one or two, as relevant:
   (a) adjacent open rectangles $\mathbf{A}_i$, if the subsegment includes $\mathbf{F}$; or
   (b) adjacent open rectangles $\mathbf{A}_i$ or $\mathbf{A}_i \setminus \mathbf{F}$, otherwise.

Note that all of these conditions are preserved by φ, so we are done. □

And finally:

Lemma 3.18. *In $\mathcal{M}$, $\mathcal{T}$ is compact.*

Proof. Let $\mathcal{H} \subseteq_{\alpha+2} \mathcal{T}$ be an open cover of $\mathbf{A}$; let $(\lambda X.t[X])$ be a closed term defining $\mathcal{H}$.

As in the proof of Lemma 3.17, construct a grid from the constants occurring in t and assume (without loss of generality) that $\mathbf{F}$ does not intersect any of the grid's open segments. (So $\mathcal{F}$ intersects only open rectangles and singletons from the grid.)

Since $\mathcal{H}$ is an open cover, for every $x \in \mathbf{A}$ there is an open set $U \in \mathcal{H}$ such that $x \in U$. Now the topology $\mathcal{T}$ is generated from open rectangles $\mathbf{R}$ (in the standard topology on $[0,1]^2$) and $\mathbf{R} \setminus \mathbf{F}$. If $x \in \mathbf{F}$ then the $U \in \mathcal{H}$ containing x must also contain an open-rectangle neighborhood $\mathbf{R}$ of x. Since $[0,1]^2$ is compact, the only problem is in covering elements $x \notin \mathbf{F}$.

Each horizontal or vertical slice of $\mathbf{A}$ is relatively compact: every open cover of the slice has a finite subcover. This is because each slice intersects $\mathbf{F}$ exactly once, and if $x \in \mathbf{F}$ then any set covering it must contain an open-rectangle neighborhood of x. Fix finitely many sets in $\mathcal{H}$ covering all open segments and all singletons in the grid.

Look at the portion of the grid consisting of one open rectangle $\mathbf{A}_i$, the four open segments forming the rectangle's four sides, and the four singletons forming the rectangle's corners. Each of these corners may be in $\mathbf{F}$, or not; and, by assumption, none of the four open segments intersects $\mathbf{F}$. If a corner is in $\mathbf{F}$ then there is an open set $U \in \mathcal{H}$ containing an open-rectangle neighborhood $\mathbf{R}$ of x. By deforming U as in the proof of Lemma 3.17 we can ensure that $\mathbf{R}$ covers arbitrarily much of $\mathbf{A}_i$.

If all four of $\mathbf{A}_i$'s corners are in $\mathbf{F}$ then there is no problem—applying the proof of Lemma 3.11, $\mathbf{A}_i$ can be covered by finitely many open sets in $\mathcal{H}$. The problem arises when one or more of $\mathbf{A}_i$'s corners are not in $\mathbf{F}$; in that case, it may be that no open set $U \in \mathcal{H}$ containing a particular corner $\{\mathbf{a}_j\}$ includes an open-rectangle neighborhood of $\mathbf{a}_j$. In this case, deforming any such U, using the permutations φ from Lemma 3.17, yields a set $\Phi(U)$ that fails to include an open-rectangle neighborhood of $\mathbf{a}_j$—thus failing to witness compactness.

As with Lemma 3.12, we need a larger class of permutations φ; we reuse the idea from that lemma and add to U an open rectangle in $\mathbf{A}_i$ adjacent to $\{\mathbf{a}_j\}$.

Let X_0 be any open set in $\mathcal{H}$ that intersects $\mathbf{A}_i$ and $\mathbf{F}$; then X_0 contains an open rectangle. If X_0 contains an open rectangle $\mathbf{R} \subseteq_{\alpha+1} \mathbf{A}_i$ adjacent to the corner $\{\mathbf{a}_j\}$ then we are done. Suppose not; let φ deform X_0 by adding to X_0 an open rectangle $\mathbf{R}$ adjacent to the corner such that:

- $\mathbf{R}$ does not contain any open rectangle $\mathbf{S} \subseteq_{\alpha+1} X_0$ or $(\mathbf{S} \setminus \mathbf{F}) \subseteq_{\alpha+1} X_0$;
- for each side $\mathbf{T}$ of $\mathbf{A}_i$, $\mathbf{R}$ is not adjacent to the whole $\mathbf{T}$; and
- the boundary of $\mathbf{R}$ does not intersect $\mathbf{F}$.

That the last condition can be satisfied follows from the facts: (1) $\mathbf{F}$ is dense in $\mathbf{A}$; (2) by choice of grid, $\mathbf{F}$ does not intersect any of the grid's open segments; and (3) by assumption, the corner $\{\mathbf{a}_j\}$ does not intersect $\mathbf{F}$.

The proof of Lemma 3.17 applies to the class of permutations φ defined here, so we are done. □

So RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ does not prove that every compact T_2 space is T_3. Also, as Proposition 3.19 shows, "every compact T_2 space is T_3" does not imply $(\mathcal{A}_2)$. So this theorem is strictly weaker than $(\mathcal{A}_2)$ and is not implied by $(\mathcal{A}_1)$.

3.4.5 Summary of reverse-mathematical results, and future work

We have shown that the theory RCA_0^ω + (ATOMS) + $(\mathcal{A}_2)$ + $(\mathcal{A}_1)$ implies the two basic topological theorems:

- the product of two compact spaces is compact;
- every compact T_2 space is T_3;

and that these theorems are not implied by the theory RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$. In fact, neither theorem implies $(\mathcal{A}_2)$ over RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$; this follows from an argument similar to that used in the proof of Proposition 2.17. So the two basic topological theorems lie strictly between $(\mathcal{A}_1)$ and $(\mathcal{A}_2)$.

Proposition 3.19 (RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$)**.** *If* Φ *is either of the two (third-order, over* α*) topological theorems mentioned above, then* Φ *does not imply* $(\mathcal{A}_2)$*.*

Proof. Let $\mathcal{M}$ be a term model of RCA_0^ω + (ATOMS) + $(\mathcal{A}_1)$ + Φ, with constant symbols only for A_1, the type-$(\alpha+2)$ families and type-$(\alpha+1)$ sets mentioned in Φ, and the basic functionals mentioned in the axioms RCA_0^ω + (ATOMS). Note that $\mathcal{M}$ contains no functionals of types higher than $(\alpha+2)$, except for the combinators Σ and Π.

Suppose, for a contradiction, that $\mathcal{M}$ satisfies $(\mathcal{A}_2)$. Let A_2 be defined by a closed term $(\lambda\mathcal{F}.t[\mathcal{F}])$. Without loss of generality, t is in normal form, so its subterms have free variables only of types 0 and α. As in the proof of Proposition 2.17, the type-$(\alpha+2)$ variable $\mathcal{F}$ occurs in subterms of t only on the left-hand side of an application; let $(\mathcal{F}\,S_1),\ldots,(\mathcal{F}\,S_k)$ list all applications in t involving $\mathcal{F}$. Let $\overline{m_i}$ and $\overline{x_i}$ list all type-0 and type-α variables, respectively, free in S_i. Note that S_i has type $(\alpha+1)$.

The rest of the proof is the same as in Proposition 2.17, except that the sequence of elements now depends both on $\overline{m_i}$ and on $\overline{x_i}$. In the end, we have a sequence $(n,\overline{x})\mapsto Y_{n,\overline{x}}$ such that if $\mathcal{F}$ and $\mathcal{G}$ agree on all $Y_{n,\overline{x}}$'s, then $(\mathcal{A}_2\,\mathcal{F}) =_0 (\mathcal{A}_2\,\mathcal{G})$.

Letting $\mathcal{F}$ be $\emptyset$, we have that for any family $\mathcal{G}$ that excludes all $Y_{n,\overline{x}}$'s, $(\mathcal{A}_2\,\mathcal{G}) =_0 0$. It remains only for us to show that such a $\mathcal{G}$ exists—for example, the family defined by:

$$\mathcal{G} :=_{\alpha+2} \left(\lambda X. \begin{cases} 1, & \text{if } \exists(x_1,\ldots,x_k)\,\exists n\,(X =_{\alpha+1} Y_{n,\overline{x}}) \\ 0, & \text{otherwise.} \end{cases} \right)$$

Since t is finite, it has only finitely many type-α variables x_i. The type-α quantification and the type-$(\alpha+1)$ equality can be handled by A_1, while the type-0 quantification can be handled by E_1. (Note that t contains type-0 variables only if we have included a constant for E_1 in our term model.) But this is a contradiction. $\square$

Concerning the lower-order consequences of these two topological theorems, one quirk that seems to arise is that a model $\mathcal{M}$ need not contain a varied selection of topologies. I suspect that one can build models of RCA_0^ω + (ATOMS) + ($\mathcal{A}_1$) in which one of the theorems holds and the other fails. It might be interesting to include additional third-order, over α, axioms asserting the existence of certain classes of topologies. It seems that Proposition 3.19 would still apply to such a theory, and so ($\mathcal{A}_2$) would not be implied.

LIST OF REFERENCES

[AF98] Jeremy Avigad and Solomon Feferman. Gödel's functional ("dialectica") interpretation. In Samuel R. Buss, editor, *Handbook of Proof Theory*, volume 137 of *Studies in Logic and the Foundations of Mathematics*, pages 337–406. Elsevier Science, 1998.

[Fef77] Solomon Feferman. Theories of finite type related to mathematical practice. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 913–972. North-Holland, Amsterdam, 1977.

[Har87] Victor Harnik. Set existence axioms for general (not necessarily countable) stability theory. *Annals of Pure and Applied Logic*, 34:231–243, 1987.

[Jec73] Thomas J. Jech. *The Axiom of Choice*, volume 75 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, 1973.

[Koh05] Ulrich Kohlenbach. Higher order reverse mathematics. In Stephen G. Simpson, editor, *Reverse Mathematics 2001*, volume 21 of *Lecture Notes in Logic*, pages 281–295. A.K. Peters, 2005.

[KW98] Iraj Kalantari and Lawrence Welch. Point-free topological spaces, functions and recursive points: Filter foundation for recursive analysis I. *Annals of Pure and Applied Logic*, 93(1–3):125–151, 1998.

[KW04] Iraj Kalantari and Larry Welch. Density and baire category in recursive topology. *Mathematical Logic Quarterly*, 50(4–5):381–391, 2004.

[Mum05] Carl Mummert. *On the Reverse Mathematics of General Topology*. PhD thesis, Pennsylvania State University, 2005.

[Mum07] Carl Mummert. Reverse mathematics of MF spaces. *Journal of Mathematical Logic*, 6(2):203–232, 2007.

[Mun75] James Raymond Munkres. *Topology: a First Course*. Prentice-Hall, Englewood Cliffs, New Jersey, 1975.

[Sim99] Stephen G. Simpson. *Subsystems of Second Order Arithmetic*. Springer-Verlag, Berlin, 1999.

[Tro73] A. S. Troelstra, editor. *Metamathematical Investigation of Intuitionistic Arithmetic and Analysis*, volume 344 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1973.