On the Cauchy Completeness of the Constructive Cauchy Reals

Robert S. Lubarsky [1,2]

Dept. of Mathematical Sciences
Florida Atlantic University
Boca Raton, FL 33431, USA

Abstract

Intuitionistic set theory without choice axioms does not prove that every Cauchy sequence of rationals has a modulus of convergence, or that the set of Cauchy sequences of rationals is Cauchy complete. Several other related non-provability results are also shown.

Keywords: Cauchy sequence, modulus of convergence, intuitionistic set theory, topological model, Kripke model

1 Introduction

Are the reals Cauchy complete? This is, for every Cauchy sequence of real numbers, is there a real number which is its limit?

This sounds as though the answer should be "of course". After all, the reals are defined pretty much to make this true. The reason to move from the rationals to the reals is exactly to "fill in the holes" that the rationals have. So however you define **R**, you'd think its Cauchy completeness would be immediate. At the very least, this property would be a litmus test for any proffered definition.

In fact, for the two most common notions of real number, Dedekind and Cauchy real, this is indeed the case, under classical logic. First off, classically Cauchy and Dedekind reals are equivalent anyway. Then, taking a real as an equivalence class of Cauchy sequences, given any Cauchy sequence of reals, a canonical representative can be chosen from each real, and a limit real can be built from them by a kind of diagonalization, all pretty easily.

Intuitionistically, though, this whole procedure breaks down. Starting even at the beginning, Cauchy and Dedekind reals are no longer equivalent notions (see [2]

[1] Thanks are due to Thomas Streicher for his careful reading of the first proof, and to Michael Rathjen for bringing the question that got this started to my attention and for his support during this work.

[2] Email: Robert.Lubarsky@alum.mit.edu

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

or [4]). While the Dedekind reals are complete, working with the Cauchy reals, it's not clear that a representative can be chosen from each equivalence class, or, even if you could, that a limit could be built from them by any means.

It is the purpose of this paper to show that, indeed, such constructions are not in general possible, answering a question of Martin Escardo and Alex Simpson ([1]).

While the original motivation of this work was to show the final theorem, that the Cauchy reals are not provably Cauchy complete, it is instructive to lay out the framework and examine the related questions so laid bare. A Cauchy real is understood as an equivalence class of Cauchy sequences of rationals. When working with a Cauchy sequence, one usually needs to know not only that the sequence converges, but also how fast. In classical set theory, this is definable from the sequence itself, and so is not problematic. The same cannot be said for other contexts. For instance, in recursion theory, the complexity of the convergence rate might be important. In our context, intuitionistic set theory, the standard way to define a modulus of convergence just doesn't work. Certainly given a Cauchy sequence X(n) and positive rational ϵ, there is an integer N such that for m, n $>$ N $|$ X(n) - X(m) $|< \epsilon$: that's the definition of a Cauchy sequence. A modulus of convergence is a function f such that for any such ϵ f on ϵ returns such an N. Classically this is easy: let f return the least such N. Intuitionistically that won't work. And there's no obvious alternative. So a real is taken to be an equivalence class of pairs $\langle X, f\rangle$, where X is a Cauchy sequence and f a modulus of convergence.

One immediate source of confusion here is identifying reals with sequence-modulus pairs. A real is an equivalence class of such pairs, and it is not obvious how a representative can be chosen constructively from each real; in fact, this cannot in general be done, as we shall see. This distinction has not always be made though. For instance, as observed by Fred Richman, in [7], the Cauchy completeness of the reals was stated as a theorem, but what was proved was the Cauchy completeness of sequence-modulus pairs. To be precise, what was shown was that, given a countable sequence, with its own modulus of convergence, of sequence-modulus pairs, then there is a limit sequence, with modulus. For that matter, it is not hard (and left to the reader) that, even if the given sequence does not come equipped with its own modulus, it still has a Cauchy sequence as a limit, although we will have to punt on the limit having a modulus. But neither of those two observations is the Cauchy completeness of the reals.

Nonetheless, these observations open up the topic about what kinds of behavior in the limit one can expect given certain input data. There are two, independent parameters. Does the outside Cauchy sequence have a modulus of convergence? And are its individual members sequence-modulus pairs, or merely naked sequences? Notice that, while the first question is yes-no, the second has a middle option: the sequences have moduli, but not uniformly. Perhaps each entry in the big sequence is simply a Cauchy sequence of rationals, and it is hypothesized to have a modulus somewhere, with no information about the modulus given. These possibilities are all summarized in the following table.

Entries in the outside Cauchy sequence are:

	seq-mod pairs	seqs that have mods somewhere	seqs that may not have mods anywhere
Outside seq has a mod			
Outside seq doesn't have a mod			

Not to put the cart before the horse, is it even possible to have a Cauchy sequence with no modulus of convergence, or with only non-uniform moduli? It has already been observed that the obvious classical definition of such does not work intuitionistically, but it still remains to be shown that no such definition is possible.

The goal of this paper is to prove the negative results as much as possible, that any given hypothesis does not show that there is a limit Cauchy sequence, or in certain cases that there is no limit with a modulus.

The positive results are all easy enough (even if the last one is a bit tricky) and so are left as exercises. They are:

(i) Every Cauchy sequence with modulus of sequence-modulus pairs has a limit sequence with modulus.

(ii) Every Cauchy sequence of sequence-modulus pairs has a limit sequence.

(iii) Every Cauchy sequence with modulus of Cauchy sequences has a limit sequence.

In tabular form, the positive results are:

Entries in the outside Cauchy sequence are:

	seq-mod pairs	seqs that have mods somewhere	seqs that may not have mods anywhere
Outside seq has a mod	There is a limit with mod.	There is a limit sequence.	There is a limit sequence.
Outside seq doesn't have a mod	There is a limit sequence.		

Regarding the first and last columns, the negative results are that these are the positive results cited above are the best possible. In detail:

Theorem 2.1 *IZF_{Ref} does not prove that every Cauchy sequence has a modulus of convergence. It follows that IZF_{Ref} does not prove that every Cauchy sequence* *of sequence-modulus pairs converges to a Cauchy sequence with a modulus of convergence.*

Theorem 4.1 *IZF_{Ref} does not prove that every Cauchy sequence of Cauchy sequences converges to a Cauchy sequence.*

Theorem 5.1 *IZF_{Ref} does not prove that every Cauchy sequence with a modulus of convergence of Cauchy sequences converges to a Cauchy sequence with a modulus of convergence.*

The middle column is discussed briefly in the questions in the last section of this paper.

In tabular form, these negative results are:

Entries in the outside Cauchy sequence are:

	seq-mod pairs	seqs that have mods somewhere	seqs that may not have mods anywhere
Outside seq has a mod		See questions, section 7.	Limit may not have a mod. thm 5.1
Outside seq doesn't have a mod	Limit may not have a mod. thm 2.1	See questions, section 7.	There may not even be a limit. thm 4.1

Then there is the major negative result, the original and ultimate motivation of this work:

Theorem 6.1 *IZF_{Ref} does not prove that every Cauchy sequence of reals has a limit.*

Recalling that a real is here taken as an equivalence class of sequence-modulus pairs, to prove this result it would suffice to construct a Cauchy sequence (perhaps itself without modulus) of reals, with no sequence-modulus pair as a limit. We will do a tad better, constructing a Cauchy sequence, with modulus, of reals, with no Cauchy sequence, even without modulus, as a limit.

At this point a word about the meta-theory is in order. The results here are stated as non-theorems of IZF_{Ref}, which is the variant of IZF in which the Collection schema is replaced by the Reflection schema. The point is that these independence results are not meant to be based on a weakness of the underlying set theory. Hence the set theory taken is the strongest version of the intuitionistic theories commonly considered. The results would remain valid if IZF were augmented by yet stronger hypotheses, such as large cardinals. Of course, these remarks do not apply if IZF is augmented by whatever choice principle would be enough to build the sequences and moduli here claimed not to exist. Clearly Dependent Choice is strong enough for everything at issue here: choosing representatives from equivalence classes, building moduli, constructing Cauchy sequences. The question what weaker choice principle/s would suffice is addressed in [5].

Regarding the methodology, counter-examples are constructed in each case. These examples could be presented as either topological or Kripke models. While certain relations among topological and Kripke models are known, it is not clear to the author that the natural models in the cases before us are really the same. While the family resemblance is unmistakable (compare sections 2 and 3 below, for instance), it can be deceiving. For instance, the first model of [4] seems to be merely the Kripke version of one of the models of [2] (also in in [6], VI.8), there presented as a sheaf model but easily enough recastable as topological. But there are differences. For instance, consider Bas Spitters' proof that in the topological model the reals are not uncountable. (His proof consists of a countable list of reals and a demonstration that every real is not apart from all reals on his list. Recalling that in the topological model a real is given by a continuous function, the countable list consists of the two families of functions $x + r$ and $-x + r$, where r ranges over the rationals.) Pretty much the same argument works in the Kripke model, but it can stand a simplification that would not work in the topological model. (In the Kripke model, it suffices to consider reals of only the form $x + r$.) Such a difference is certainly subtle, arguably small, but exists nonetheless, and is enough to say that the models are not the same. A better understanding of the relations among Kripke and topological models would be a worthwhile project for some other time. For now, we would like to present the reader with adequate information without being long-winded. Hence all of the constructions will be presented as topological models, since there is better technology for dealing with them. In particular, there is already a meta-theorem (see [3]) that the (full) model over any Heyting algebra models IZF (easily, IZF_{Ref} too). So we will never have to prove that our topological models satisfy IZF_{Ref}. In contrast, we know of no such meta-theorem that would apply to the Kripke models in question. In the simplest case, the first theorem, the Kripke model will also be given, so the reader can see what's going on there. But even a cursory glance at that argument should make it clear why the author does not want to repeat the proofs of IZF_{Ref} and all the auxiliary lemmas, and the reader likely does not want to read them, three more times.

One last word about notation/terminology. For p an open set in a topological space and ϕ a formula in set theory (possibly with parameters from the topological model), "p $\subseteq \|\phi\|$", "p $\Vdash \phi$", and "p forces ϕ" all mean the same thing. Also, "WLOG" stands for "without loss of generality".

2 Not every Cauchy sequence has a modulus of convergence

Theorem 2.1 *IZF_{Ref} does not prove that every Cauchy sequence has a modulus of convergence. It follows that IZF_{Ref} does not prove that every Cauchy sequence of sequence-modulus pairs converges to a Cauchy sequence with a modulus of convergence.*

The second assertion follows immediately from the first: given a Cauchy sequence

X(n), for each n let X_n be the constant sequence X(n) paired with some modulus of convergence independent of n. Sending X to the sequence $\langle X_n \mid n \in \mathbf{N}\rangle$ embeds the Cauchy reals into Cauchy sequences of sequence-modulus pairs. If provably every one of the latter had a modulus, so would each of the former.

To prove the first assertion, we will build a topological model with a specific Cauchy sequence Z(n) of rationals with no modulus of convergence.

The topological space T consists of all Cauchy sequences of rationals.

A basic open set is given by (p, I), where p is a finite sequence of rationals and I is an open interval. A Cauchy sequence X is in (the open set determined by) (p, I) if p $\subseteq$ X, rng(X\p) $\subseteq$ I, and lim(X) $\in$ I. (Notice that, under this representation, the whole space T is given by ($\emptyset$, $\mathbf{R}$), and the empty set is given by (p, $\emptyset$) for any p.)

For this to generate a topology, it suffices to show that the basic open sets are closed under intersection. Given (p, I) and (q, J), if p and q are not compatible (i.e. neither is an extension of the other), then (p, I) $\cap$ (q, J) = $\emptyset$. Otherwise WLOG let q $\supseteq$ p. If rng(q\p) $\not\subseteq$ I then again (p, I) $\cap$ (q, J) = $\emptyset$. Otherwise (p, I) $\cap$ (q, J) = (p$\cup$q, I$\cap$J) = (q, I$\cap$J).

Let M be the Heyting-valued models based on T, as describes in e.g. [3]. Briefly, a set in M is a collection of objects of the form $\langle\sigma$, (p, I) $\rangle$, where σ inductively is a set in M. It is shown in [3] that M $\models$ IZF_{Coll} (assuming IZF_{Coll} in the meta-theory). Similarly, assuming IZF_{Ref} in the meta-theory yields M $\models$ IZF_{Ref}.

We are interested in the term $\{\langle \bar{p}, (p, I)\rangle$ |(p, I) is an open set}. (Here $\bar{p}$ is the canonical name for p. Each set in V has a canonical name in M by choosing (p, I) to be ($\emptyset$, $\mathbf{R}$) hereditarily: $\bar{x} = \{\langle \bar{y}, (\emptyset, \mathbf{R})\rangle \mid y \in x\}$.) We will call this term Z.

Proposition 2.2 $\|$ *Z is a Cauchy sequence* $\|$ = *T.*

Proof. To see that $\|$Z is total$\|$ = T, let N be an integer. (Note that each integer in M can be identified locally with an integer in V. For notational ease, we will identify integers in M and V.) Let p be any sequence of rationals of length > N. Then (p, $\mathbf{R}$) $\subseteq$ $\|$ Z(N) = p(N) $\|$ $\subseteq$ $\|$ N $\in$ dom(Z) $\|$. As T is covered by the open sets of that form, T $\subseteq$ $\|$ N $\in$ dom(Z) $\|$. That Z is a function is similarly easy.

As for Z being Cauchy, again let N be an integer and X be in T. Since X is Cauchy, there is an integer M such that beyond M X stays within an interval I of size 1/(2N). Of course, X's limit might be an endpoint of I. So let J extend I on either side and still have length less that 1/N. Then X $\in$ (X$\restriction$M, J) $\subseteq$ $\|\forall$ m, n > M | Z(m) - Z(n) |$\leq$ 1/N $\|$, making Z "Cauchy for 1/N", to coin a phrase. □

In order to complete the theorem, we need only prove the following

Proposition 2.3 $\|$*Z has no modulus of convergence*$\|$ = *T.*

Proof. Suppose (p, I) $\subseteq$ $\|$f is a modulus of convergence for Z$\|$. WLOG I is a finite interval. Let n be such that 1/n is less than the length of I, and let ϵ be (length(I) - 1/n)/2.

Let (q, J) $\subseteq$ (p, I) force a value m for f(n). WLOG length(q) > m, as q could be so extended. If J=I, then (q, J) could be extended simply by extending q with two values a distance greater than 1/n apart, thereby forcing f not to be a modulus of convergence. So $J \subset I$. That means either inf $J >$ inf I or sup $J <$ sup I. WLOG assume the latter. Let mid J be the midpoint of J, q_0 be q extended by mid J, and J_0 be (mid J, sup J). Then $(q_0, J_0) \subseteq (q, J)$, and therefore $(q_0, J_0) \subseteq \|f(n) = m\|$.

What (q_1, J_1) is depends:

CASE I: There is an open set K containing $\sup(J_0)$ such that $(q_0, K) \subseteq \|f(n) = m\|$. Then let j_{max} be the sup of the right-hand endpoints (i.e. sups) of all such K's. Let q_1 be q_0 extended by sup J_0.

<u>Claim</u>: $(q_1, (\sup J_0, j_{max})) \subseteq \|f(n) = m\|$.

Proof. Let $j \in (\sup J_0, j_{max})$. By hypothesis, there is a K such that $(q_1, (\sup J_0, j)) \subseteq (q_1, K) \subseteq \|f(n) = m\|$. Since $(q_1, (\sup J_0, j_{max}))$ is the union of the various $(q_1, (\sup J_0, j))$'s over such j's, the claim follows. □

Of course, J_1 will be $(\sup J_0, j_{max})$.

CASE II: Not Case I. Then extend by the midpoint again. That is, q_1 is q_0 extended by mid J_0, and J_1 is (mid J_0, sup J_0). Also in this case, $(q_1, J_1) \Vdash f(n) = m$.

Clearly we would like to continue this construction. The only thing that might be a problem is if the right-hand endpoint of some J_k equals (or goes beyond!) sup I, as we need to stay beneath (p, I). In fact, as soon as $\sup(J_k) > \sup(I) - \epsilon$ (if ever), extend q_k by something within ϵ of sup I, and continue the construction with left and right reversed. That is, instead of going right, we now go left. This is called "turning around".

What happens next depends.

CASE A: We turned around, and after finitely many more steps, some J_k has its inf under $\inf(I) + \epsilon$. Then extend q_k by something within ϵ of inf I. This explicitly blows f being a modulus of convergence for Z.

CASE B: Not Case A. So past a certain point (either the stage at which we turned around, or, if none, from the beginning) we're marching monotonically toward one of I's endpoints, but will always stay at least ϵ away. WLOG suppose we didn't turn around. Then the construction will continue for infinitely many stages. The q_k's so produced will in the limit be a (monotonic and bounded, hence) Cauchy sequence X. Furthermore, lim X is the limit of the $\sup(J_k)$'s. Finally, $X \in (p, I)$. Hence there is an open set (q', K) with $X \in (q', K) \subseteq \|f(n) = m'\|$, for some m'. Let k be such that $\sup(J_k) \in K$, and $q_k \supseteq q'$. Consider (q_k, J_k). Note that $(q_k, J_k \cap K)$ extends both (q_k, J_k) and (q', K), hence forces both $f(n) = m$ and $f(n) = m'$, which means that m=m'.

Therefore, at this stage in the construction, we are in Case I. By the construction, $J_{k+1} = (\sup J_k, j_{max})$, where $j_{max} \geq \sup K > \lim X = \lim_k (\sup(J_k)) \geq \sup(J_{k+1}) = j_{max}$, a contradiction. □

3 Same theorem, Kripke model version

Theorem 2.1 *IZF_{Ref} does not prove that every Cauchy sequence has a modulus of convergence.*

3.1 Construction of the Model

Let $M_0 \prec M_1 \prec \ldots$ be an ω-sequence of models of ZF set theory and of elementary embeddings among them, as indicated, such that the sequence from M_n on is definable in M_n, and such that each thinks that the next has non-standard integers. Notice that this is easy to define (mod getting a model of ZF in the first place): an iterated ultrapower using any non-principal ultrafilter on ω will do. We will ambiguously use the symbol f to stand for any of the elementary embeddings inherent in the M_n-sequence.

The Kripke model M will have underlying partial order a non-rooted tree; the bottom node (level 0) will have continuum (in the sense of M_0) many nodes, and the branching at a node of level n will be of size continuum in the sense of M_{n+1}. (We will eventually name each node by associating a Cauchy sequence to it. Some motivation will be presented during this section, and the final association will be at the end of this section.) Satisfaction at a node will be indicated with the symbol $\models$. There is a ground Kripke model, which, at each node of level n, has a copy of M_n. The transition functions (from a node to a following node) are the elementary embeddings given with the original sequence of models (and therefore will be notated by f again). Note that by the elementarity of the extensions, this Kripke model is a model of classical ZF. More importantly, the model restricted to any node of level n is definable in M_n, because the original M-sequence was so definable.

The final model M will be an extension of the ground model that will be described like a forcing extension. That is, M will consist of (equivalence classes of) the terms from the ground model. The terms are defined at each node separately, inductively on the ordinals in that model. At any stage α, a term of stage α is a set σ of the form $\{\langle\sigma_j, (p_j, I_j)\rangle \mid j \in J\}$, where J is some index set, each σ_i is a term of stage $< \alpha$, each p_j is a finite function from $\mathbf{N}$ to $\mathbf{Q}$, and each I_j is an open rational interval on the real line. Note that all sets from the ground model have canonical names, by choosing each p_j to be the empty function and I_j to be the whole real line, hereditarily.

Notice also that the definition of the terms given above will be interpreted differently at each node of the ground Kripke model, as the $\mathbf{N}$ and $\mathbf{Q}$ change from node to node. However, any term at a node gets sent by the transition function f to a corresponding term at any given later node. The definitions given later, such as the forcing relation $\Vdash$, are all interpretable in each M_n, and coherently so, via the elementary embeddings.

As a condition, each finite function p is saying "the Cauchy sequence includes me", and each interval I is saying "future rationals in the Cauchy sequence have to come from me". For each node of level n there will be an associated Cauchy sequence r (in the sense of M_n) such that at that node the true p's and I's will be those compatible with r (or, perhaps, those with which r is compatible, as the reader will). You might reasonably think that compatibility means "$p \subset r$ and $\mathrm{rng}(r \backslash p) \subseteq I$": roughly, "r extends p, and anything in r beyond p comes from I". But that's not quite right. Consider the Cauchy sequence $r(n) = 1/n$ ($n \geq 1$). $\mathrm{rng}(r) \subseteq (0, 2)$, but in a non-standard extension, r's pattern could change at a non-standard integer; at that point, it would be too late for r to change by a standard amount, but it could change by an infinitesimal amount. So the range of r could include (infinitely small) negative numbers, which are outside of (0, 2). Hence we have the following

Definition 3.1 A condition (p, I) and a Cauchy sequence r are compatible if p $\subseteq$ r, rng(r\p) $\subseteq$ I, and lim(r) $\in$ I.

(p, I) is compatible with a finite function q if p $\subseteq$ q and rng(q\p) $\subseteq$ I.

Given this notion of compatibility, speaking intuitively here, a term σ can be thought of as being interpretable (with notation σ^r) inductively in M_n as $\{\sigma_j^r \mid \langle\sigma_j, p_j, I_j\rangle \in \sigma$ and r is compatible with $(p_j, I_j)\}$. (This notion is hidden in the more formal development below, where we define and then mod out by $=_M$.)

Our next medium-term goal is to define the primitive relations at each node, $=_M$ and $\in_M$ (the subscript being used to prevent confusion with equality and membership of the ambient models M_n). In order to do this, we need first to develop our space's topology.

Definition 3.2 (q, J) $\leq$ (p, I) ((q, J) extends (p, I)) if q $\supseteq$ p, J $\subseteq$ I, and rng(q\p) $\subset$ I.

$C = \{(p_j, I_j) \mid j \in J\}$ covers (p, I) if each (p_j, I_j) extends (p, I) and each Cauchy sequence r compatible with (p, I) is compatible with some (p_j, I_j).

$\leq$ induces a notion of compatibility of conditions (having a common extension). We say that a typical member $\langle\sigma$, (p, I) $\rangle$ of a term is compatible with (q, J) if (p, I) and (q, J) are compatible.

We need some basic facts about this p.o., starting with the fact that it is a p.o.

Lemma 3.3 (i) $\leq$ *is reflexive, transitive, and anti-symmetric.*

(ii) *If (p, I) and (q, J) are each compatible with a Cauchy sequence r, then they are compatible with each other.*

(iii) *If (p, I) and (q, J) are compatible, then their glb in the p.o. is (p$\cup$q, I$\cap$J).*

(iv) $\{$*(p, I)*$\}$ *covers (p, I).*

(v) *A cover of a cover is a cover. That is, if C covers (p, I), and, for each (p_j, I_j) $\in$ C, C_j covers (p_j, I_j), then $\bigcup_j C_j$ covers (p, I).*

(vi) *If C covers (p, I) and (q, J) $\leq$ (p, I), then (q, J) is covered by C $\wedge$ (q, J) $=_{def}$ $\{$($p_j$$\cup$q, $I_j$$\cap$J) $\mid$ (q, J) is compatible with (p_j, I_j) $\in$ C$\}$.*

Proof. Left to the reader. $\square$

Now we are in a position to define $=_M$ and $\in_M$. This will be done via a forcing relation $\Vdash$.

Definition 3.4 $(p, I) \Vdash \sigma =_M \tau$ and $(p, I) \Vdash \sigma \in_M \tau$ are defined inductively on σ and τ, simultaneously for all (p, I):

$(p, I) \Vdash \sigma =_M \tau$ iff for all $\langle\sigma_j, (p_j, I_j)\rangle \in \sigma$ compatible with (p, I) $(p \cup p_j, I \cap I_j) \Vdash \sigma_j \in_M \tau$ and vice versa, and

$(p, I) \Vdash \sigma \in_M \tau$ iff there is a cover C of (p, I) such that for all $(p_j, I_j) \in C$ there is a $\langle\tau_k, (p_k, I_k)\rangle \in \tau$ such that $(p_j, I_j) \leq (p_k, I_k)$ and $(p_j, I_j) \Vdash \sigma =_M \tau_k$.

(We will later extend this forcing relation to all formulas.)

Definition 3.5 At a node (with associated real r), for any two terms σ and τ, $\sigma =_M \tau$ iff, for some (p, I) compatible with r, (p, I) $\Vdash \sigma =_M \tau$.

Also, $\sigma \in_M \tau$ iff for some (p, I) compatible with r, (p, I) $\Vdash \sigma \in_M \tau$.

Thus we have a first-order structure at each node.

The transition functions are the same as before. That is, if σ is an object at a node, then it's a term, meaning in particular it's a set in some M_n. Any later node has for its universe the terms from some $\mathrm{M}_m, m \geq n$. With f the elementary embedding from M_n to M_m, f can also serve as the transition function between the given nodes. These transition functions satisfy the coherence conditions necessary for a Kripke model.

To have a Kripke model, f must also respect $=_M$ and $\in_M$, meaning that f must be an $=_M$- and $\in_M$-homomorphism (i.e. $\sigma =_M \tau \ \rightarrow\ \mathrm{f}(\sigma) =_M \mathrm{f}(\tau)$, and similarly for $\in_M$). In order for these to be true, we need an additional restriction on the model. By way of motivation, one requirement is, intuitively speaking, that the sets σ can't shrink as we go to later nodes. That is, once σ_j gets into σ at some node, it can't be thrown out at a later node. σ_j gets into σ because r is compatible with (p_j, I_j) (where $\langle\sigma_j, (p_j, I_j)\rangle \in \sigma$). So we need to guarantee that if r and (p, I) are compatible and r' is associated to any extending node then r' and (p, I) are compatible for any condition (p, I). This holds exactly when r' extends r and all of the entries in r'\r are infinitesimally close to lim(r). This happens, for instance, when r' = f(r). Other such examples would be f(r) truncated at some non-standard place and arbitrarily extended by any Cauchy sequence through the reals with standard part lim(r); in fact, all such r' have that form. We henceforth take this as an additional condition on the construction: once r is associated to a node, then for any r' associated to an extending node, rng(r'\r) must consist only of rationals infinitely close to lim(r).

Lemma 3.6 *f is an $=_M$ and $\in_M$-homomorphism.*

Proof. If $\sigma =_M \tau$ then let (p, I) compatible with r witness as much. At any later node, (p, I) = f((p, I)) = (f(p), f(I)) $\Vdash f(\sigma) =_M f(\tau)$. Also, the associated real r' would still be compatible with (p, I). So the same (p, I) would witness $f(\sigma) =_M f(\tau)$ at that node. Similarly for $\in_M$. □

We can now conclude that we have a Kripke model.

Lemma 3.7 *This Kripke model satisfies the equality axioms:*

(i) $\forall x\, x = x$

(ii) $\forall x, y\, x = y \rightarrow y = x$

(iii) $\forall x, y, z\, x = y \wedge y = z \rightarrow x = z$

(iv) $\forall x, y, z\, x = y \wedge x \in z \rightarrow y \in z$

(v) $\forall x, y, z\, x = y \wedge z \in x \rightarrow z \in y$.

Proof. 1: It is easy to show with a simultaneous induction that, for all (p, I) and σ, (p, I) $\Vdash \sigma =_M \sigma$, and, for all $\langle\sigma_j, (p_j, I_j)\rangle \in \sigma$ compatible with (p, I), (p $\cup$ p_j, I $\cap$ I_j) $\Vdash \sigma_i \in_M \sigma$.

2: Trivial because the definition of (p, I) $\Vdash \sigma =_M \tau$ is itself symmetric.

3: For this and the subsequent parts, we need some lemmas.

Lemma 3.8 *If (p', I') $\leq$ (p, I) $\Vdash \sigma =_M \tau$ then (p', I') $\Vdash \sigma =_M \tau$, and similarly for $\in_M$.*

Proof. By induction on σ and τ. □

Lemma 3.9 *If (p, I) $\Vdash \rho =_M \sigma$ and (p, I) $\Vdash \sigma =_M \tau$ then (p, I) $\Vdash \rho =_M \tau$.*

Proof. Again, by induction on terms. □

Returning to proving property 3, the hypothesis is that for some (p, I) and (q, J) each compatible with r, (p, I) $\Vdash \rho =_M \sigma$ and (q, J) $\Vdash \sigma =_M \tau$. By the first lemma, (p $\cup$ q, I $\cap$ J) $\Vdash \rho =_M \sigma, \sigma =_M \tau$, and so by the second, (p $\cup$ q, I $\cap$ J) $\Vdash \rho =_M \tau$. Also, (p $\cup$ q, I $\cap$ J) is compatible with r.

4: Let (p, I) $\Vdash \rho =_M \sigma$ and (q, J) $\Vdash \rho \in_M \tau$. We will show that (p $\cup$ q, I $\cap$ J) $\Vdash \sigma \in_M \tau$. Let C be a cover of (q, J) witnessing (q, J) $\Vdash \rho \in_M \tau$. We will show that (p $\cup$ q, I $\cap$ J) $\wedge$ C = (p, I) $\wedge$ C is a cover of (p $\cup$ q, I $\cap$ J) witnessing (p $\cup$ q, I $\cap$ J) $\Vdash \sigma \in_M \tau$. Let (q_i, J_i) $\in C$ and $\langle \tau_k, p_k, I_k \rangle$ be the corresponding member of τ. By the first lemma, (p $\cup$ q_i, I $\cap$ J_i) $\Vdash \rho =_M \sigma$, and so by the second, (p $\cup$ q_i, I $\cap$ J_i) $\Vdash \sigma =_M \tau_k$.

5: Similar, and left to the reader. □

With this lemma in hand, we can now mod out by $=_M$, so that the symbol "=" is interpreted as actual set-theoretic equality. We will henceforth drop the subscript $_M$ from = and $\in$, although we will not distinguish notationally between a term σ and the model element it represents, σ's equivalence class.

At this point, we need to finish specifying the model in detail. What remains to be done is to associate a Cauchy sequence to each node. At the bottom level, assign each Cauchy sequence from M_0 to exactly one node. Inductively, suppose we chose have the sequence r at a node with ground model M_n. There are continuum-in-the-sense-of-M_{n+1}-many immediate successor nodes. Associate each possible candidate r' in M_{n+1} with exactly one such node. (As a reminder, that means each member of rng(r'\r) is infinitely close to lim(r).)

By way of notation, a node will be named by its associated sequence. Hence "r $\models \phi$" means ϕ holds at the node with sequence r.

Note that, at any node of level n, the choice of r's from that node on is definable in M_n. This means that the evaluation of terms (at and beyond the given node) can be carried out over M_n, and so the Kripke model (from the given node on) can be defined over M_n, truth predicate and all.

3.2 The Forcing Relation

Which (p, I)'s count as true determines the interpretation of all terms, and hence of truth in the end model. We need to get a handle on this. As with forcing, we need a relation (p, I) $\Vdash \phi$ which supports a truth lemma. Note that, by elementarity, it doesn't matter in which classical model M_n or at what node in the ground Kripke model $\Vdash$ is being interpreted (as long as the parameters are in the interpreting model, of course).

Definition 3.10 (p, I) $\Vdash \phi$ is defined inductively on ϕ:

(p, I) $\Vdash \sigma =_M \tau$ iff for all $\langle \sigma_j, (p_j, I_j) \rangle \in \sigma$ compatible with (p, I) $(p \cup p_j, I \cap I_j) \Vdash \sigma_j \in_M \tau$ and vice versa

(p, I) $\Vdash \sigma \in_M \tau$ iff there is a cover C of (p, I) such that for all $(p_j, I_j) \in C$ there is a $\langle \tau_k, (p_k, I_k) \rangle \in \tau$ such that $(p_j, I_j) \leq (p_k, I_k)$ and $(p_j, I_j) \Vdash \sigma =_M \tau_k$.

(p, I) $\Vdash \phi \wedge \psi$ iff (p, I) $\Vdash \phi$ and (p, I) $\Vdash \psi$

(p, I) $\Vdash \phi \vee \psi$ iff there is a cover C of (p, I) such that, for each (p_j, I_j) $\in C$, (p_j, I_j) $\Vdash \phi$ or (p_j, I_j) $\Vdash \psi$

(p, I) $\Vdash \phi \to \psi$ iff for all (q, J) $\leq$ (p, I) if (q, J) $\Vdash \phi$ then (q, J) $\Vdash \psi$

(p, I) $\Vdash \exists x\ \phi(x)$ iff there is a cover C of (p, I) such that, for each (p_j, I_j) $\in C$, there is a σ such that (p_j, I_j) $\Vdash \phi(\sigma)$

(p, I) $\Vdash \forall x\ \phi(x)$ iff for all σ (p, I) $\Vdash \phi(\sigma)$

Lemma 3.11 (i) *If (q, J) $\leq$ (p, I) $\Vdash \phi$ then (q, J) $\Vdash \phi$.*

(ii) *If C covers (p, I), and (p_j, I_j) $\Vdash \phi$ for all (p_j, I_j) $\in$ C, then (p, I) $\Vdash \phi$.*

(iii) *(p, I) $\Vdash \phi$ iff for all r compatible with (p, I) there is a (q, J) compatible with r such that (q, J) $\Vdash \phi$.*

(iv) *Truth Lemma: For any node r, r $\models \phi$ iff (p, I) $\Vdash \phi$ for some (p, I) compatible with r.*

Proof. 1. A trivial induction, using of course the earlier lemmas about $\leq$ and covers.

2. Easy induction. The one case to watch out for is $\to$, where you need to invoke the previous part of this lemma.

3. Trivial, using 2.

4. By induction on ϕ, in detail for a change.

In all cases, the right-to-left direction ("forced implies true") is pretty easy, by induction. (Note that only the $\to$ case needs the left-to-right direction in this induction.) Hence in the following we show only left-to-right ("if true at a node then forced").

$=$: This is exactly the definition of $=$.

$\in$: This is exactly the definition of $\in$.

$\wedge$: If r $\models \phi \wedge \psi$, then r $\models \phi$ and r $\models \psi$. Inductively let (p, I) $\Vdash \phi$ and (q, J) $\Vdash \psi$, where (p, I) and (q, J) are each compatible with r. That means that (p, I) and (q, J) are compatible with each other, and (p $\cup$ q, I $\cap$ J) suffices.

$\vee$: If r $\models \phi \vee \psi$, then WLOG r $\models \phi$. Inductively let (p, I) $\Vdash \phi$, (p, I) compatible with r. {(p, I)} suffices.

$\to$: Suppose to the contrary r $\models \phi \to \psi$ but no (p, I) compatible with r forces such. Work in the node f(r). (Recall that f is the universal symbol for the various transition functions in sight. What we mean more specifically is that if r $\in$ M_n, i.e. if r is a node from level n, then f(r) is the image of r in M_{n+1}, i.e. in the Kripke structure on level n+1.) Let (p, I) be compatible with f(r) and p have non-standard (in the sense of M_n) length (equivalently, I has infinitesimal length). Since (p, I) $\nVdash \phi \to \psi$ there is a (q, J) $\leq$ (p, I) such that (q, J) $\Vdash \phi$ but (q, J) $\nVdash \psi$. By the previous part of this lemma, there is an r' compatible with (q, J) such that no condition compatible with r' forces ψ. At the node r', by induction, r' $\not\models \psi$, even though r' $\models \phi$ (since r' is compatible with (p, I) $\Vdash \phi$). This contradicts the assumption on r (i.e. that r $\models \phi \rightarrow \psi$), since r' extends r (as nodes).

$\exists$: If r $\models \exists x\ \phi(x)$ then let σ be such that r $\models \phi(\sigma)$. Inductively there is a (p, I) compatible with r such that (p, I) $\Vdash \phi(\sigma)$. $\{$(p, I)$\}$ suffices.

$\forall$: Suppose to the contrary r $\models \forall x\ \phi(x)$ but no (p, I) compatible with r forces such. As with $\rightarrow$, let (p, I) non-standard be compatible with f(r). Since (p, I) $\not\Vdash \forall x\ \phi(x)$ there is a σ such that (p, I) $\not\Vdash \phi(\sigma)$. By the previous part of this lemma, there is an r' compatible with (p, I) such that, for all (q, J) compatible with r', (q, J) $\not\Vdash \phi(\sigma)$. By induction, that means that r' $\not\models \phi(\sigma)$. This contradicts the assumption on r (i.e. that r $\models \forall x\ \phi(x)$), since r' extends r (as nodes).

$\square$

3.3 The Final Proofs

Using $\Vdash$, we can now prove

Theorem 3.12 *This Kripke model satisfies IZF$_{Ref}$.*

Proof. Note that, as a Kripke model, the axioms of intuitionistic logic are satisfied, by general theorems about Kripke models.

- Empty Set: The interpretation of the term $\emptyset$ will do.
- Infinity: The canonical name for ω will do. (Recall that the canonical name $\bar{x}$ of any set $x \in$ V is defined inductively as $\{\langle \bar{y}, (\emptyset, \mathbf{R})\rangle \mid y \in x\}$.)
- Pairing: Given σ and τ, $\{\langle \sigma, (\emptyset, \mathbf{R})\rangle, \langle \tau, (\emptyset, \mathbf{R})\rangle\}$ will do.
- Union: Given σ, $\{\langle \tau, J \cap J_i\rangle \mid$ for some σ_i, $\langle \tau, J\rangle \in \sigma_i$ and $\langle \sigma_i, J_i\rangle \in \sigma\}$ will do.
- Extensionality: We need to show that $\forall x\ \forall y\ [\forall z\ (z \in x \leftrightarrow z \in y) \rightarrow x = y]$. So let σ and τ be any terms at a node r such that r $\models$ "$\forall z\ (z \in \sigma \leftrightarrow z \in \tau)$". We must show that r $\models$ "$\sigma = \tau$". By the Truth Lemma, let r $\in$ J $\Vdash$ "$\forall z\ (z \in \sigma \leftrightarrow z \in \tau)$"; i.e. for all r' $\in$ J, ρ there is a J' containing r' such that J $\cap$ J' $\Vdash \rho \in \sigma \leftrightarrow \rho \in \tau$. We claim that J $\Vdash$ "$\sigma = \tau$", which again by the Truth Lemma suffices. To this end, let $\langle \sigma_i, J_i\rangle$ be in σ; we need to show that J $\cap$ J$_i$ $\Vdash \sigma_i \in \tau$. Let r' be an arbitrary member of J $\cap$ J$_i$ and ρ be σ_i. By the choice of J, let J' containing r' be such that J $\cap$ J' $\Vdash \sigma_i \in \sigma \leftrightarrow \sigma_i \in \tau$; in particular, J $\cap$ J' $\Vdash \sigma_i \in \sigma \rightarrow \sigma_i \in \tau$. It has already been observed in 3.7, part 1, that J $\cap$ J' $\cap$ J$_i$ $\Vdash \sigma_i \in \sigma$, so J $\cap$ J' $\cap$ J$_i$ $\Vdash \sigma_i \in \tau$. By going through each r' in J $\cap$ J$_i$ and using 3.11, part 3, we can conclude that J $\cap$ J$_i$ $\Vdash \sigma_i \in \tau$, as desired. The other direction ("$\tau \subseteq \sigma$") is analogous.
- Set Induction (Schema): Suppose r $\models$ "$\forall x\ ((\forall y \in x\ \phi(y)) \rightarrow \phi(x))$"; by the Truth Lemma, let J containing r force as much. We must show r $\models$ "$\forall x\ \phi(x)$". Suppose not. Using the definition of satisfaction in Kripke models, there is an r' extending (i.e. infinitesimally close to) r (hence in J) and a σ such that r' $\not\models \phi(\sigma)$. By elementarity, there is such an r' in M$_n$, where n is the level of r. Let σ be such a term of minimal V-rank among all r's $\in$ J. Fix such an r'. By the Truth Lemma (and the choice of J), r' $\models$ "$(\forall y \in \sigma\ \phi(y)) \rightarrow \phi(\sigma)$". We claim that r' $\models$ "$\forall y \in \sigma\ \phi(y)$". If not, then for some r" extending r' (hence in J) and

τ, r" $\models \tau \in \sigma$ and r" $\not\models \phi(\tau)$. Unraveling the interpretation of $\in$, this choice of τ can be substituted by a term τ of lower V-rank than σ. By elementarity, such a τ would exist in M_n, in violation of the choice of σ, which proves the claim. Hence r' $\models \phi(\sigma)$, again violating the choice of σ. This contradiction shows that r $\models$ "$\forall x\ \phi(x)$".

- Separation (Schema): Let $\phi(x)$ be a formula and σ a term. Then $\{\langle \sigma_i, J \cap J_i \rangle \mid \langle \sigma_i, J_i \rangle \in \sigma$ and $J \Vdash \phi(\sigma_i)\}$ will do.
- Power Set: A term $\hat{\sigma}$ is a canonical subset of σ if for all $\langle \sigma_i, \hat{J}_i \rangle \in \hat{\sigma}$ there is a $J_i \supseteq \hat{J}_i$ such that $\langle \sigma_i, J_i \rangle \in \sigma$. $\{\langle \hat{\sigma}, (\emptyset, \mathbf{R}) \rangle \mid \hat{\sigma}$ is a canonical subset of $\sigma\}$ is a set (in M_n), and will do.
- Reflection (Schema): Recall that the statement of Reflection is that for every formula $\phi(x)$ (with free variable x and unmentioned parameters) and set z there is a transitive set Z containing z such that Z reflects the truth of $\phi(x)$ in V for all $x \in$ Z. So to this end, let $\phi(x)$ be a formula and σ be a set at a node r of level n (in the tree which is this Kripke model's partial order). Let k be such that the truth of $\phi(x)$ at node r and beyond is Σ_k definable in M_n. In M_n, let X be a set containing σ, r, and ϕ's parameters such that X $\prec_k M_n$. Let τ be $\{\langle \rho, (\emptyset, \mathbf{R}) \rangle \mid \rho \in$ X is a term$\}$. τ will do.

□

We are interested in the canonical term $\{\langle \bar{p}, (p, I) \rangle \mid$ p is a finite function from $\mathbf{N}$ to $\mathbf{Q}$ and I is a non-empty, open interval from the reals with rational endpoints$\}$, where $\bar{p}$ is the canonical name for p. We will call this term Z. Note that at node r Z gets interpreted as r.

Proposition 3.13 *For all nodes r, r $\models$ "Z is a Cauchy sequence".*

Proof. To see that $\perp \models$ "Z is total", suppose r $\models$ "N is an integer". Then $(\langle N, r(N) \rangle, \mathbf{R})$ is compatible with r and forces "Z(N) = r(N)". That Z is a function is similarly easy.

As for Z being Cauchy, again let r $\models$ "N is an integer". Since r is Cauchy, there is an integer M such that beyond M r stays within an interval I of size 1/(2N). Of course, future nodes might be indexed by Cauchy sequences s extending r that go outside of I, but only by an infinitesimal amount. So let J extend I on either side and still have length less that 1/N. Then (r$\restriction$M, J) is compatible with r, and forces that Z beyond M stay in J, making Z "Cauchy for 1/N", to coin a phrase. □

In order to complete the theorem, we need only prove the following

Proposition 3.14 *For all nodes r, r $\models$ "Z has no modulus of convergence."*

Proof. Suppose r $\models$ "f is a modulus of convergence for Z." Let (p, I) compatible with r force as much. WLOG I is a finite interval. Let n be such that 1/n is less than the length of I, and let ϵ be (length(I) - 1/n)/2.

Let (q, J) $\leq$ (p, I) force a value m for f(n). WLOG length(q) $>$ m, as q could be so extended. If J=I, then (q, J) could be extended simply by extending q with two values a distance greater than 1/n apart, thereby forcing f not to be a modulus of convergence. So J $\subset$ I. That means either inf J $>$ inf I or sup J $<$ sup I. WLOG assume the latter. Let mid J be the midpoint of J, q_0 be q extended by mid J, and J_0 be (mid J, sup J). Then $(q_0, J_0) \leq (q, J)$, and therefore $(q_0, J_0) \Vdash f(n) = m$.

What (q_1, J_1) is depends:

CASE I: There is an open set K containing $\sup(J_0)$ such that $(q_0, K) \Vdash f(n) = m$. Then let j_{max} be the sup of the right-hand endpoints (i.e. sups) of all such K's. Let q_1 be q_0 extended by sup J_0.

Claim: $(q_1, (\sup J_0, j_{max})) \Vdash f(n) = m$.

Proof. Let r be any Cauchy sequence compatible with $(q_1, (\sup J_0, j_{max}))$. Since $\lim r < j_{max}$, r (that is, $\mathrm{rng}(r \backslash q_0)$) is actually bounded below j_{max}. By the definition of j_{max}, there is an open K containing sup J_0 such that r is bounded by sup K. As r is bounded below by sup J_0, r (again, $\mathrm{rng}(r \backslash q_0)$) is contained within K. As $(q_0, K) \Vdash f(n) = m$, $r \models f(n) = m$. □

Of course, J_1 will be $(\sup J_0, j_{max})$.

CASE II: Not Case I. Then extend by the midpoint again. That is, q_1 is q_0 extended by mid J_0, and J_1 is (mid J_0, sup J_0). Also in this case, $(q_1, J_1) \Vdash f(n) = m$.

Clearly we would like to continue this construction. The only thing that might be a problem is if the right-hand endpoint of some J_k equals (or goes beyond!) sup I, as we need to stay beneath (p, I). In fact, as soon as $\sup(J_k) > \sup(I) - \epsilon$ (if ever), extend q_k by something within ϵ of sup I, and continue the construction with left and right reversed. That is, instead of going right, we now go left. This is called "turning around".

What happens next depends.

CASE A: We turned around, and after finitely many more steps, some J_k has its inf within $\inf(I) + \epsilon$. Then extend q_k by something within ϵ of inf I. This explicitly blows f being a modulus of convergence for Z.

CASE B: Not Case A. So past a certain point (either the stage at which we turned around, or, if none, from the beginning) we're marching monotonically toward one of I's endpoints, but will always stay at least ϵ away. WLOG suppose we didn't turn around. Then the construction will continue for infinitely many stages. The q_k's so produced will in the limit be a (monotonic and bounded, hence) Cauchy sequence r. Furthermore, lim r is the limit of the $\sup(J_k)$'s. Finally, r is compatible with (p, I). Hence $r \models$ "f is total", and so $r \models f(n) = m'$, for some m'. Let some condition compatible with r force as much. This condition will have the form (q', K), where $\lim r \in K$. Let k be such that $\sup(J_k) \in K$, and $q_k \supseteq q'$. Consider (q_k, J_k). Note that $(q_k, J_k \cap K)$ extends both (q_k, J_k) and (q', K), hence forces both $f(n) = m$ and $f(n) = m'$, which means that m=m'.

Therefore, at this stage in the construction, we are in Case I. By the construction, $J_{k+1} = (\sup J_k, j_{max})$, where $j_{max} \geq \sup K > \lim r = \lim_k (\sup(J_k)) \geq \sup(J_{k+1}) = j_{max}$, a contradiction. □

4 Not every Cauchy sequence of Cauchy sequences converges

Theorem 4.1 *IZF_{Ref} does not prove that every Cauchy sequence of Cauchy sequences converges to a Cauchy sequence.*

The statement of the theorem itself needs some elaboration. The distance $d(x_{0n}, x_{1n})$ between two Cauchy sequences x_{0n} and x_{1n} is the sequence $\mid x_{0n} - x_{1n} \mid$. $x_{0n} < x_{1n}$ if there are $m, N \in \mathbf{N}$ such that for all $k > N$ $x_{0k} + 1/m < x_{1k}$. A rational number r can be identified with the constant Cauchy sequence $x_n =$ r. $x_n = 0$ if $\forall m \ \exists N \ \forall k > N \ \mid x_k \mid < 1/m$. x_{0n} and x_{1n} are equal (as reals, equivalent as Cauchy sequences if you will) if $d(x_{0n}, x_{1n}) = 0$. With these definitions in place, we can talk about Cauchy sequences of Cauchy sequences, and limits of such. The theorem is then that it is consistent with IZF$_{Ref}$ to have a convergent sequences of Cauchy sequences with no limit.

Note that we are not talking about reals! A real number would be an equivalence class of Cauchy sequences (omitting, for the moment, considerations of moduli of convergence). It would be weaker to claim that the sequence of reals represented by the constructed sequence of sequences has no limit. After all, given a sequence of reals, it's not clear that there is a way to choose a Cauchy sequence from each real. We are claiming here that even if your task is made easier by being handed a Cauchy sequence from each real, it may still not be possible to get a "diagonalizing", i.e. limit, Cauchy sequence.

4.1 The Topological Space and Model

Let T be the space of Cauchy sequences of Cauchy sequences. By way of notation, if X is a member of T, then X$_j$ will be the j^{th} Cauchy sequence in X; as a Cauchy sequence of rationals, X$_j$ will have values X$_j$(0), X$_j$(1), etc. Still notationally, if X$_n \in$ T, then the j^{th} sequence in X$_n$ is X$_{nj}$. In the classical meta-universe, the Cauchy sequence X$_j$ has a limit, lim(X$_j$); in addition, the sequence X has a limit, which will be written as lim(X).

A basic open set p is given by a finite sequence $\langle (p_j, I_j) \mid j < n_p \rangle$ of basic open sets from the space of the previous theorem (i.e. p_j is a finite sequence of rationals and I_j is an open interval), plus an open interval I$_p$. X $\in p$ if X$_j \in (p_j, I_j)$ for each j<n$_p$, if lim(X$_j$) $\in$ I$_p$ for each j$\geq$n$_p$, and lim(X) $\in$ I. Note that q $\subseteq$ p (q <u>extends</u> p) if n$_q \geq$ n$_p$, (q$_j$, K$_j$) $\subseteq (p_j, I_j)$ for j $<$ n$_p$, K$_j \subseteq$ I$_p$ for j $\geq$ n$_p$, and I$_q \subseteq$ I$_p$.

p and q are compatible (where WLOG $n_p \leq n_q$) if, for $j < n_p$ (p$_j$, I$_j$) and (q$_j$, K$_j$) are compatible, for $n_p \leq j < n_q$ K$_j \cap$ I$_p \neq \emptyset$, and I$_q \cap$ I$_p \neq \emptyset$. In this case, p $\cap$ q is not the basic open set you'd think it is, but rather a union of such. The problem is that for $n_p \leq j < n_q$ it would be too much to take the j^{th} component to be (q$_j$, K$_j \cap$ I$_p$), because that would leave out all extensions of q$_j$ with entries from K$_j \setminus$I$_p$ before they finally settle down to K$_j \cap$ I$_p$. So p $\cap$ q will instead be covered by basic open sets in which the j^{th} component will be (r$_j$, K$_j \cap$ I$_p$), where (r$_j$, K$_j$) $\subseteq$ (q$_j$, K$_j$). (So the given basic open sets form not a basis for the topology, but rather a sub-basis.)

As always, the sets in the induced Heyting-valued model M are of the form

$\{\langle\sigma_k, p_k\rangle \mid k \in K\}$, where K is some index set, each σ_k is a set inductively, and each p_k is an open set. Note that all sets from the ground model have canonical names, by choosing each p_k to be T (i.e. $n_p = 0$ and $I_p = \mathbf{R}$), hereditarily. M satisfies IZF_{Ref}.

4.2 The Extensions $\leq_j$ and $\leq_\infty$

In the final proof, we will need the following notions.

Definition 4.2 j-extension $\leq_j$: q $\leq_j$ p for some $j < n_p$ if q and p satisfy all of the clauses of q extending p except possibly for the condition on the j^{th} component: (q_j, K_j) need not be a subset of (p_j, I_j), although we will still insist that (q_j, I_j) be a subset of (p_j, I_j).

More concretely, q_j comes from p_j by extending with elements from I_j; it's just that we're no longer promising to keep to I_j in the future. Notice that $\leq_j$ is not transitive; the transitive closure of $\leq_j$ will be notated as $\leq_j^*$.

Definition 4.3 ∞-extension $\leq_\infty$: q $\leq_\infty$ p if q and p satisfy all of the clauses of q extending p except possibly for the last, meaning that I_q need not be a subset of I_p.

$\leq_\infty^*$ is the transitive closure of $\leq_\infty$.

Lemma 4.4 *Suppose $q \subseteq p$, $q \subseteq \|f(n)=m\|$ for some particular m and n, and $j < n_p$. Then for all $x \in I_j$ there is an $r \subseteq p$, $r \leq_j^* q$ such that $r \subseteq \|f(n)=m\|$ and $x \in L_j$, where (r_j, L_j) is r's j^{th} component.*

Proof. If x $\in K_j$, then we are done: let r be q. So assume WLOG that x $\geq$ sup(K_j). The inspiration for this construction is the construction of the previous theorem. The main difference is that not only do we have (q_j, K_j) to contend with, we also have all of q's other components around. Hence the notion of a j-extension: we will do the last theorem's construction on the j^{th} coordinates, and leave all the others alone.

First off, we would like to show that q has a j-extension q' $\subseteq$ p also forcing $f(n) = m$ such that sup(K_j) $\in$ K'$_j$. Toward this end, let X $\in$ T be a member of (the open set determined by) q except that lim(X_j) = sup(K_j). X is in p, so there is some r $\subseteq$ p such that X $\in$ r and r forces a value for f(n), say m'. q and r are compatible though: apart from the j^{th} component, X is in both, and the only thing happening in the j^{th} component is that, in r, sup(K_j) $\in L_j$, meaning that K_j and L_j overlap. So any common extension of both q and r would have to force f(n)=m and f(n)=m'; since p already forces that f is a function, m=m'. Using r, it is easy to construct the desired q': take the j^{th} component from r, and let each other component be the intersection of the corresponding components from r and q.

If there is such a q' such that x $\in$ K'$_j$, then we are done. Else we would like to mimic the last theorem's construction by having in our next condition the interval part of the j^{th} component be (sup(K_j), j_{max}) (for a suitably defined j_{max}). The problem is, q has all these other components around. For any real $y < j_{max}$ we could find a j-extension of q with (sup(K_j), y) in the j^{th} component, but not necessarily for $y = j_{max}$ itself.

To this end, consider all such q' as above. Each q' can be extended (to say q") by restricting the interval in the j^{th} component to (sup(K_j), sup(K'_j)). Let q_1 be such a q" where that interval is at least half as big as possible (i.e. among all such q", where of course sup(K'_j) has to be bounded by sup(I_j)).

Continue this construction so that q_n is defined from q_{n-1} just as q_1 was defined from q. WLOG dovetail this construction with extending all other components so that after infinitely many steps we would have produced an X $\in$ T. (This remark needs a word of justification about the j^{th} components. By the definition of j-extension alone, it is not clear that a sequence of j-extending conditions $q_0 \geq_j q_1 \geq_j$ ... converges to a point in T. In our case, though, by the construction itself, the various K_{nj}'s are monotonically increasing and bounded, hence the X_j so determined is Cauchy.)

If at some finite stage we have covered x, then we are done. If not, then sup(X_j) = $\sup_n$(sup(K_{nj})) $\leq$ x $\in I_j$, so that X $\in$ p. So there is some r $\subseteq$ p with X $\in$ r such that r forces a value for f(n), say m'. Let ϵ be sup(L_j) - sup(X_j). Eventually in the construction, K_{nj} will be contained within ϵ of sup(X_j). With r as the witness, at the next stage $K_{(n+1)j}$ would go beyond sup(X_j), which is a contradiction. Hence this case is not possible, and at some finite stage we must have covered x, as desired. □

We have a similar lemma for ∞-extensions.

Lemma 4.5 *Suppose $q \subseteq p$ and $q \subseteq \|f(n){=}m\|$ for some particular m and n. Then for all $x \in I_p$ there is an $r \subseteq p$, $r \leq^*_\infty q$ such that $r \subseteq \|f(n){=}m\|$ and $x \in I_r$.*

Proof. Similar to the above. □

Observe that the same arguments work for preserving finitely many values of f simultaneously.

4.3 The Final Proof

We are interested in the canonical term $\{\langle \bar{p}_j, p\rangle \mid$ p is an open set$\}$, where $\bar{p}_j$ is the canonical name for the sequence $\langle p_j \mid j < n_p\rangle$ from p. We will call this term Z. It should be clear that T = $\|$Z is a Cauchy sequence of Cauchy sequences$\|$. Hence we need only prove

Proposition 4.6 *$T = \|Z$ does not have a limit$\|$.*

Proof. Suppose p $\subseteq$ $\|$f is a Cauchy sequence$\|$. It suffices to show that for some q $\subseteq$ p, q $\subseteq$ $\|$f $\neq$ lim(Z)$\|$.

If p can ever be extended to force infinitely many values for f simultaneously, then do so, and further extend (it suffices here to extend merely the last component) to force Z away from f's limit. This suffices for the theorem.

If this is not possible, then the construction will be to build one or two sequences of open sets, p_k and possibly r_k, indexed by natural numbers k. It is to be understood even though not again mentioned that the construction below is to be dovetailed with a countable sequence of moves designed to produce a single member of T in the end (i.e. each individual component must shrink to a real as in the previous theorem, the n_{p_k}'s must be unbounded as k goes through $\mathbf{N}$, and the last components I_{p_k} must shrink to something of length 0).

First, let p_0 be built by extending p by cutting I_p to its bottom third, and let L be some point in I_p's top half. If p_0 can be extended (to p_1) so that f is forced to have an additional value (that is, beyond what has already been forced) in I_p's top half, then do so. Else proceed as follows. First extend p_0 to force an additional value for f, necessarily in I_p's bottom half. Then by the second lemma above, ∞-extend that latter condition, to q say, preserving the finitely many values of f already determined, and getting L into I_q. Typically $n_q > n_p$, so let $\bar{q}$ be such that $n_{\bar{q}} = n_q$, if $j < n_p$ then $\bar{q}$'s j^{th} component is the same as p's, if $n_p \leq j < n_q$ then $\bar{q}$'s j^{th} component is ($\emptyset$, I_p), and $I_{\bar{q}} = I_p$. Note that $q \subseteq \bar{q} \subseteq p$, so we can apply the first lemma above to q and $\bar{q}$. Starting from q, iteratively on j from n_p up to n_q, j-extend to get L into the interval part of the j^{th} component, while preserving the finitely many values of f already determined. Call the last condition so obtained r_0. Finally, ∞-extend r_0 to get the last component to be a subset of I_{p_0}, while still preserving f of course. Let this latter condition be p_1.

Stages $k > 0$ will be similar. To start, if possible, extend p_k to force an additional value for f in I_p's top half. Call this condition p_{k+1}.

If that is not possible, first extend p_k to force a new value for f, necessarily in I_p's bottom half. Then ∞-extend (to q say) to get L into the last component I_q. After that, j-extend for each j from n_{r_i} to n_q to get L in those components, where i is the greatest integer less than k such that r_i is defined. (It bears mentioning that r_h is defined if and only if at stage h we are in this case.) If need be, shrink those components to be subsets of I_{r_i}, for the purpose of getting $r_k \subseteq r_i$ (once we define r_k). That last condition will be r_k. Next, ∞-extend r_k to get the last component to be a subset I_{p_k}. This final condition is p_{k+1}.

This completes the construction.

If the second option happens only finitely often, let k be greater than the last stage where it happens. Then not only does p_k force lim(Z) to be in I_p's bottom third, as all p_i's do actually, but also p_k is respected in the rest of the construction: for $i > k$, $p_i \subseteq p_k$. Let $l \geq k$ be such that $6/l <$ length(I_p) (i.e. the distance between I_p's top half and bottom third is greater than 1/l). Recall that $p \subseteq \|$f is a Cauchy sequence$\|$; that is, $p \subseteq \|\forall \epsilon > 0\ \exists N\ \forall m, n \geq N\ |f(m) - f(n)| < \epsilon\|$. Since $1/l > 0$, $p \subseteq \|\exists N\ \forall m, n \geq N\ |f(m) - f(n)| < 1/l\|$. That means there is a cover C of p such that each $q \in C$ forces a particular value for N. Let S be $\bigcap_{j \geq k} p_j$, and let $q \in C$ contain S. Similarly, let $\hat{q}$ containing S force a value for f(N). $q \wedge \hat{q} \wedge p_k$ is non-empty because it contains S, and $q \wedge \hat{q} \wedge p_k$ forces by the construction that f(N) is in I_p's top half, by the choice of q that lim(f) is away from I_p's bottom third, and by choice of k that lim(Z) is in I_p's bottom third. In short, $q \wedge \hat{q} \wedge p_k \subseteq \|f \neq \lim(Z)\|$.

Otherwise the second option happens infinitely often. Then we have an infinite descending sequence of open sets r_k, and a similar argument works. Let S be $\bigcap_j r_j$, where the intersection is taken only over those j's for which r_j is defined. Let k be such that $r_k \subseteq \|$lim(Z) - midpoint(I_p) $< \epsilon\|$, for some fixed $\epsilon > 0$. Let q containing S be such that, for a fixed value of N, $q \subseteq \|\forall m, n \geq N\ |f(m) - f(n)| < \epsilon\|$. Let $\hat{q}$ force a particular value for f(N), necessarily in I_p's bottom half. Again, $q \wedge \hat{q} \wedge$

$p_k \subseteq \|f \neq \lim(Z)\|$.

□

5 The given Cauchy sequence has a modulus, but the limit doesn't

Theorem 5.1 *IZF_{Ref} does not prove that every Cauchy sequence with a modulus of convergence of Cauchy sequences converges to a Cauchy sequence with a modulus of convergence.*

Definition 5.2 c is a convergence function for a Cauchy sequence $\langle X_j \mid j \in \mathbf{N} \rangle$ if c is a decreasing sequence of positive rationals; for all n, if j, k $\geq$ n then $| X_j - X_k | \leq c(n)$; and $\lim(c(n)) = 0$.

Notice that convergence functions and moduli of convergence are easily convertible to each other: if c is the former, then d(n) := the least m such that $c(m) \leq 2^{-n}$ is the latter; and if d is the latter, then $c(n) := 2^{-m}$, where m is the greatest integer such that max(m, d(m)) $\leq$ n, is the former. Therefore the current construction will be of a Cauchy sequence $\langle X_j \mid j \in \mathbf{N} \rangle$ with a convergence function but no limit. Without loss of generality, the convergence function in question can be taken to be $c(n) = 2^{-n}$.

Let the topological space T be $\{\langle X_j \mid j \in \mathbf{N} \rangle \mid \langle X_j \mid j \in \mathbf{N} \rangle$ is a Cauchy sequence of Cauchy sequences with convergence function 2^{-n} $\}$. As in the previous section, for X $\in$ T, X_j will be the j^{th} Cauchy sequence in X's first component. The real number represented by X_j, i.e. X_j's limit, will be written as $\lim(X_j)$. In the classical meta-universe, the limit of the sequence $\langle X_j \mid j \in \mathbf{N} \rangle$ will be written as lim(X).

T is a subset of the space from the previous section, and the topology of T is to be the subspace topology. That is, a basic open set p is given again by a finite sequence $\langle (p_j, I_j) \mid j < n_p \rangle$ and an open interval I_p. X $\in$ p if, again, $X_j \in (p_j, I_j)$ for each $j<n_p$, $\lim(X_j) \in I_p$ for each $j \geq n_p$, and $\lim(X) \in I_p$. p and q are compatible under the same conditions as before, and p $\cap$ q is covered by basic open sets, just as in the last theorem; the convergence function causes no extra trouble.

Note that q $\subseteq$ p (q <u>extends</u> p) if all of the same conditions from the last section hold: $n_q \geq n_p$, $(q_j, K_j) \subseteq (p_j, I_j)$ for $j < n_p$, $K_j \subseteq I_p$ for $j \geq n_p$, and $I_q \subseteq I_p$.

In the following, we will need to deal with basic open sets in canonical form. The issue is the following. Suppose, in p, I_0 = (0, 1) and I_1 = (0, 10). Then X_1 could certainly contain elements from (0, 10). However, when it comes to taking limits, X_1 has 2 as an upper bound, because of I_0 and the convergence function 2^{-n}, but this is not reflected in I_1.

Definition 5.3 p is in canonical form if, for $j < k < n_p$, $|\sup(I_j) - \sup(I_k)| \leq 2^{-j}$, and also $|\sup(I_j) - \sup(I_p)| \leq 2^{-j}$.

The value of canonical form is that, if for $j < n_p$ $lim(X_j) = sup(I_j)$ and if $\lim(X) = \sup(I_p)$, then, although X $\notin$ p, X could still be in T.

Proposition 5.4 *Every open set is covered by open sets in canonical form.*

Proof. Let X $\in$ p open. If, in q $\subseteq$ p, J_k is an interval with midpoint $\lim(X_k)$ and radius independent of k, and I_q an interval with midpoint lim(X) and the same radius, then q will be canonical. We will construct such a q containing X.

By way of choosing the appropriate radius, as well as n_q, let δ be half the distance from lim(X) to the closer of I_p's endpoints. Let N $\geq n_p$ be such that for all k $\geq$ N $\lim(X_k)$ is within δ of lim(X). Let r $\leq \delta$ be such that for all k $<$ N ($\lim(X_k)$ - r, $\lim(X_k)$ + r) $\subseteq I_k$. Let $n_q \geq$ N be such that for all k $\geq n_q$ $\lim(X_k)$ is within r of lim(X). For k $< n_q$ let J_k be the neighborhood with center $\lim(X_k)$ and radius r, and let q_k be an initial segment of X_k long enough so that beyond it X_k stays within J_k. Let I_q be the neighborhood with center lim(X) and radius r. This q suffices. □

As always, the sets in the induced Heyting-valued model M are of the form $\{\langle\sigma_k, p_k\rangle \mid k \in K\}$, where K is some index set, each σ_k is a set inductively, and each p_k is an open set. Note that all sets from the ground model have canonical names, by choosing each p_k to be T (i.e. n_p = 0 and I_p = $\mathbf{R}$), hereditarily. M satisfies IZF_{Ref}.

We are interested in the canonical term $\{\langle\bar{p}_j, p\rangle \mid$ p is an open set$\}$, where $\bar{p}_j$ is the canonical name for the sequences $\langle p_j \mid j < n_p\rangle$ from p. We will call this term Z. It should be clear that T = $\|$Z is a Cauchy sequence of Cauchy sequences with convergence function $2^{-n}\|$. Hence we need only prove

Proposition 5.5 *$T = \|$No Cauchy sequence equal to lim(Z) has a modulus of convergence$\|$.*

Proof. Suppose p $\subseteq \|$f is a modulus of convergence for a Cauchy sequence g$\|$, p in canonical form. It suffices to show that for some q $\subseteq$ p, q $\subseteq \|$g $\neq$ lim(Z)$\|$.

Let $\epsilon <$ (length I_p)/2. Let q $\subseteq$ p in canonical form force "f(ϵ) = N"; WLOG $n_q >$ N. We can also assume (by extending again if necessary) that q forces a value for g(N); WLOG g(N) $\leq$ midpoint(I_p). Let X $\in$ p be on the boundary of q; that is, X_k extends q_k ($k < n_q$), X_k beyond length(q_k) is a sequence through J_k with limit sup(J_k), and X beyond n_q is a sequence through I_q with limit sup(I_q) (more precisely, $\langle\lim(X_k) \mid k \geq n_q\rangle$ is such a sequence).

(Technical aside: By the canonicity of q's form, X $\in$ T. But why should X be in p? This could fail only if sup(J_k) = sup(I_k) or if sup(I_q) = sup(I_p). The latter case would actually be good. The point of the current argument is to get a condition r (forcing the things q forces) such that I_r contains points greater than g(N) + ϵ, which would fall in our lap if sup(I_q) = sup(I_p). If sup(I_q) $<$ sup(I_p) and sup(J_k) = sup(I_k), then X_k must be chosen so that $\lim(X_k)$ is slightly less than this sup. Could this interfere with 2^{-n} being a convergence function for X? No, by the canonicity of p. If l is another index such that sup(J_l) = sup(I_l), then by letting $\lim(X_l)$ be shy of this sup by the same amount as for k the convergence function 2^{-n} is respected (for these two indices). If sup(J_l) $<$ sup(I_l), then what to do depends on whether sup(J_k) and sup(J_l) are strictly less than $2^{-min(k,l)}$ apart or exactly that far apart. In the former case, there's some wiggle room in the k^{th} slot for $\lim(X_k)$ to be less than sup(J_k). In the latter, sup(J_l) must be sup(J_k) either increased or decreased by $2^{-min(k,l)}$. The first option is not possible, by the canonicity of p, as sup(J_l) $<$ sup(I_l). In the second option, having $\lim(X_k)$ be less than sup(J_k) brings $\lim(X_k)$ and $\lim(X_l)$ even closer together. Similar considerations apply to comparing $\lim(X_k)$ and $\lim_k(X_k) = \sup(I_q)$.)

Let q_1 in canonical form containing X force values for $f(\epsilon)$ and $g(f(\epsilon))$. Since q_1 and q are compatible, they force the same such values. WLOG q_1 is such that $\sup(I_{q_1})$ is big (that is, $\sup(I_{q_1})$ - $\sup(I_q)$ is at least half as big as possible). Continuing inductively, define q_{n+1} from q_n as q_1 was defined from q. Continue until I_{q_n} contains points greater than $g(N) + \epsilon$. This is guaranteed to happen, because, if not, the infinite sequence q_n will converge to a point X in p. Some neighborhood r of X forcing values for $f(\epsilon)$ and $g(f(\epsilon))$ will contain some q_n, witnessing that q_{n+1} would have been chosen with larger last component than it was, as in the previous proofs.

Once the desired q_n is reached, shrink I_{q_n} to be strictly above $g(N) + \epsilon$. Call this new condition r. r forces "$\lim(Z) > g(N) + \epsilon$", and r also forces "$g \leq g(N) + \epsilon$". So r forces "$g \neq \lim(Z)$", as desired. □

6 The reals are not Cauchy complete

Theorem 6.1 *IZF_{Ref} does not prove that every Cauchy sequence of reals has a limit.*

As stated in the introduction, what we will actually prove will be what seems to be the hardest version: there is a Cauchy sequence, with its own modulus of convergence, of real numbers, with no Cauchy sequence as a limit, even without a modulus of convergence. Other versions are possible, such as changing what does and doesn't have a modulus. After all of the preceding proofs, and after the following one, it should not be too hard for the reader to achieve any desired tweaking of this version.

Let T consist of all Cauchy sequences of Cauchy sequences, all with a fixed convergence function of 2^{-n}. An open set p is given by a finite sequence $\langle (p_j, I_j) \mid j < n_p \rangle$ as well as an interval I_p, with the usual meaning to $X \in p$.

Recall from the previous section:

Definition 6.2 p is in canonical form if, for each $j < k < n_p$ $\mid \sup(I_j)$ - $\sup(I_k) \mid \leq 2^{-j}$. Also, $\mid \sup(I_j)$ - $\sup(I_p) \mid \leq 2^{-j}$.

Also from the last section:

Proposition 6.3 *Every open set is covered by sets in canonical form.*

Henceforth when choosing open sets we will always assume they are in canonical form.

Definition 6.4 p and q are similar, p∼q, if $n_p = n_q, I_p = I_q, I_k = J_k$, and $\text{length}(p_k) = \text{length}(q_k)$. So p and q have the same form, and can differ only and arbitrarily on the rationals chosen for their components.

If moreover $p_k = q_k$ for each $k \in J$ then we say that p and q are J-similar, $p \sim_J q$.

If p∼q, this induces a homeomorphism on the topological space T, and therefore on the term structure. (To put it informally, wherever you see p_k, or an initial segment or extension thereof, replace it (or the corresponding part) with q_k, and vice versa. This applies equally well to members of T, open sets, and (hereditarily) terms.)

Definition 6.5 If p, q, and r are open sets, σ is a term, and q and r are similar, then the image of p under the induced homeomorphism is notated by p_{qr} and that of σ by σ_{qr}.

Lemma 6.6 *$p \Vdash \phi(\vec{\sigma})$ iff $p_{qr} \Vdash \phi(\vec{\sigma}_{qr})$.*

Proof. A straightforward induction. $\square$

Definition 6.7 σ has support J if for all $p \sim_J q$ $\perp \Vdash \sigma = \sigma_{pq}$. σ has finite support if σ has support J for some finite set J.

The final model M is the collection of all terms with hereditarily finite support.

As always, let Z be the canonical term. Note that Z is not in the symmetric submodel! However, each individual member of Z, Z_j, is, with support $\{j\}$. Also, so is $\langle [Z_j] \mid j \in \mathbf{N} \rangle$, which we will call [Z], with support $\emptyset$. (Here, for Y a Cauchy sequence, [Y] is the equivalence class of Cauchy sequences with the same limit as Y, i.e. the real number of which Y is a representative.) That's because no finite change in Z_j affects $[Z_j]$. (Notice that even though each member of $[Z_j]$ has support $\{j\}$, $[Z_j]$'s support is still empty.) It will ultimately be this sequence [Z] that will interest us. But first:

Proposition 6.8 *$M \models IZF_{Ref}$.*

Proof. As far as the author is aware, symmetric submodels have been studied only in the context of classical set theory, not intuitionistic, and, moreover, the only topological models in the literature are full models, in which the terms of any given model are all possible terms built on the space in question, not submodels. Nonetheless, the same proof that the full model satisfies IZF (easily, IZF_{Ref}) applies almost unchanged to the case at hand. To keep the author honest without trying the patience of the reader, only the toughest axiom, Separation, will be sketched.

To this end, suppose the term σ and formula ϕ have (combined) support J (where the support of a formula is the support of its parameters, which are hidden in the notation used). The obvious candidate for a term for the appropriate subset of σ is $\{\langle \sigma_i, p \cap p_i \rangle \mid \langle \sigma_i, p_i \rangle \in \sigma \wedge p \Vdash \phi(\sigma_i)\}$, which will be called $\mathrm{Sep}_{\sigma,\phi}$. We will show that this term has support J.

To this end, let $q \sim_J r$. We need to show that $\perp \Vdash \mathrm{Sep}_{\sigma,\phi} = (\mathrm{Sep}_{\sigma,\phi})_{qr}$. In one direction, any member of $(\mathrm{Sep}_{\sigma,\phi})_{qr}$ is of the form $\langle \sigma_i, p \cap p_i \rangle_{qr}$, where $\langle \sigma_i, p \cap p_i \rangle \in \mathrm{Sep}_{\sigma,\phi}$, i.e. $\langle \sigma_i, p_i \rangle \in \sigma$ and $p \Vdash \phi(\sigma_i)$. We need to show that $(p \cap p_i)_{qr} \Vdash (\sigma_i)_{qr} \in \mathrm{Sep}_{\sigma,\phi}$. Since $\perp \Vdash \sigma = \sigma_{qr}$ and $(p_i)_{qr} \Vdash (\sigma_i)_{qr} \in \sigma_{qr}$, $(p_i)_{qr} \Vdash (\sigma_i)_{qr} \in \sigma$. In addition, by the lemma above, $p_{qr} \Vdash \phi_{qr}((\sigma_i)_{qr})$ (where ϕ_{qr} is the result of taking ϕ and applying the homeomorphism to its parameters). Since ϕ's parameters have support J, $\perp \Vdash \phi_{qr} = \phi$, and $p_{qr} \Vdash \phi((\sigma_i)_{qr})$. Summarizing, $(p \cap p_i)_{qr} \Vdash (\sigma_i)_{qr} \in \sigma \wedge \phi((\sigma_i)_{qr})$, so $(p \cap p_i)_{qr} \Vdash (\sigma_i)_{qr} \in \mathrm{Sep}_{\sigma,\phi}$, as was to be shown.

The other direction is similar. $\square$

So there was no harm in taking the symmetric submodel. The benefit of having done so is the following

Lemma 6.9 *Extension Lemma: Suppose q, r $\subseteq$ p, q $\subseteq \|f(n) = m\|$ and for $j \in J$ (q_j, K_j) = (r_j, L_j) (i.e. q and r agree on f's support). Then r has an extension forcing f(n) = m.*

Proof. Take a sequence of refinements of q converging to a point X on q's boundary, as follows. Consider $j < n_r, j \notin$ J. If $K_j \cap L_j$ is non-empty, then just work within the latter set. Else either $\sup(K_j) < \inf(L_j)$, in which case let $\lim(X_j) = \sup(K_j)$, or $\inf(K_j) > \sup(L_j)$, in which case let $\lim(X_j) = \inf(K_j)$. (In what follows, we will consider only the first of those two cases.) Similarly for I_q and I_r. As usual, since X $\in$ p, X has a neighborhood forcing a value for f(n); since X is on q's boundary, any such neighborhood has to force the same value for f(n) that q did. Let q_1 be such a neighborhood where, for j the smallest integer not in J, $\sup((K_1)_j)$ - $\sup(K_j)$ is at least half as big as possible.

To continue this construction, consider what would happen if sup $(K_1)_j < \inf(L_j)$. We would like to take another point X, this time on the boundary of q_1, with $\lim(X_j) = \sup(K_1)_j$. The only possible obstruction is that $(q_1)_j$ might have entries far enough away from $\sup(K_1)_j$ so that the constraint of the convergence function would prevent there being such an X. In this case, change $(q_1)_j$ so that this is no longer an obstruction. Since $j \notin$ J, the new condition is J-similar to the old, and so will still force the same value for f(n).

Repeat this construction, making such that each of the finitely many components $j < n_r$, $j \notin$ J and the final component get paid attention infinitely often (meaning $\sup((K_{n+1})_j)$ - $\sup((K_n)_j)$ is at least half as big as possible). This produces a sequence q_n. Eventually q_n will be compatible with r. If not, let X be the limit of the q_n's. If X $\notin$ r then for some component j $\lim(X_j) < \inf(L_j)$. X has a neighborhood, say q_∞, forcing f(n) = m. At some large enough stage at which j gets paid attention, the existence of q_∞ would have made the j^{th} component of the next q_n contain $\lim(X_j)$, a contradiction. □

With the Extension Lemma in hand, the rest of the proof is easy. It should be clear that [Z] has convergence function 2^{-n}. So it remains only to show

Proposition 6.10 *$\|$ [Z] has no limit $\|$ = T.*

Proof. Suppose p $\subseteq \|$ f is a Cauchy sequence $\|$. It suffices to find a q $\subseteq$ p such that q $\subseteq \|$ f $\neq$ lim([Z]) $\|$.

By the Extension Lemma, all of f's values are determined by f's finite support J. So f cannot be a limit for [Z], as any such limit has to be affected by infinitely many components. □

7 Questions

There is a variant of the questions considered nestled between the individual Cauchy sequences of the big Cauchy sequence being adorned with a modulus of convergence and not. It could be that each such sequence has a modulus of convergence, but the sequence is not paired with any modulus in the big sequence. Looked at differently, perhaps the big sequence is one of Cauchy sequences with moduli of convergence but not uniformly. Certainly this extra information would not weaken any of the positive results. Would it weaken any of the negative theorems though? Presumably not: knowing that each of the individual sequences has a modulus doesn't seem to help to build a limit sequence or a modulus for such, if there's no way you can get your hands on them. Still, in the course of trying to prove this some technical difficulties were encountered, so the questions remain open.

The negative results here open up other hierarchies. Starting with the rationals, one could consider equivalence classes of Cauchy sequence with moduli of convergence. By the last theorem, that may not be Cauchy complete. So equivalences classes can be taken of sequences of those. This process can be continued, presumably into the transfinite. Is there a useful structure theorem here? All of this can be viewed as taking place inside of the Dedekind reals, which are Cauchy complete. There is a smallest Cauchy complete set of reals, namely the intersection of all such sets. As pointed out to me by the referee, this could be a proper subset of the Dedekind reals, since that is the case in the topological model of [2]. Naturally enough, the same is also the case in the Kripke model of [4]. Is there any interesting structure between the Cauchy completion of the rationals and the Dedekind reals? What about the corresponding questions for other notions of reals, such as simply Cauchy sequences sans moduli?

As indicated in the introduction, the first two models, one topological and the other Kripke, are essentially, even if not substantially, different. What is the relation between the two?

In the presence of Countable Choice, all of the positive results you could want here are easily provable (e.g. every Cauchy sequence has a modulus of convergence, the reals are Cauchy complete, etc.). Countable Choice itself, though, is a stronger principle than necessary for this, since, as pointed out to me by Fred Richman, these positive results are true under classical logic, but classical logic does not imply Countable Choice. Are there extant, weaker choice principles that would suffice instead? Can the exact amount of choice necessary be specified? These questions will start to be addressed in the forthcoming [5], but there is certainly more that can be done than is even attempted there.

8 Bibliographical references

References

[1] M.H. Escardo and A.K. Simpson, A Universal Characterization of the Closed Euclidean Interval (Extended Abstract), Sixteenth Annual IEEE Symposium on Logic in Computer Science, pp.115-125, 2001

[2] M.P. Fourman and J.M.E. Hyland, Sheaf models for analysis, in M.P. Fourman, C.J. Mulvey, and D.S. Scott (eds.), Applications of Sheaves, Lecture Notes in Mathematics Vol. 753 (Springer-Verlag, Berlin Heidelberg New York, 1979), pp. 280-301

[3] R.J. Grayson, Heyting-valued models for intuitionistic set theory, in M.P. Fourman, C.J. Mulvey, and D.S. Scott (eds.), Applications of Sheaves, Lecture Notes in Mathematics Vol. 753 (Springer-Verlag, Berlin Heidelberg New York, 1979), pp. 402-414

[4] R. Lubarsky and M. Rathjen, On the constructive Dedekind reals, in preparation

[5] R. Lubarsky and F. Richman, in preparation

[6] S. MacLane and I. Moerdijk, Sheaves in Geometry and Logic (Springer-Verlag, New York, 1992)

[7] A.S. Troelstra and D. van Dalen, Constructivism in Mathematics, Vol. 1 (North Holland, Amsterdam New York Oxford Todyo, 1988)