Separating the Fan Theorem and Its Weakenings II

Robert S. Lubarsky
Dept. of Mathematical Sciences
Florida Atlantic University
Boca Raton, FL 33431
Robert.Lubarsky@alum.mit.edu

April 20, 2017

Abstract

Varieties of the Fan Theorem have recently been developed in reverse constructive mathematics, corresponding to different continuity principles. They form a natural implicational hierarchy. Earlier work showed all of these implications to be strict. Here we re-prove one of the strictness results, using very different arguments. The technique used is a mixture of realizability, forcing in the guise of Heyting-valued models, and Kripke models.

keywords: fan theorems, Kripke models, forcing, Heyting-valued models, formal topology, recursive realizability
AMS 2010 MSC: 03C90, 03E70, 03F50, 03F60, 03D80

1 Introduction

The Fan Theorem states that, in $2^{<\omega}$, every bar (i.e. set of nodes which contains a member of every (infinite) path) is uniform (i.e. contains a member of every (infinite) path some fixed level of $2^{<\omega}$). It has in recent years become an important principle in foundational studies of constructivism. In particular, various weakenings of it have been shown to be equivalent to some principles involving continuity and compactness [2, 4, 7]. These weakenings all involve strengthening the hypothesis, by restricting which bars they apply to. The strictest version, FAN_Δ or Decidable Fan, is to say that the bar B in question is **decidable**: every node is either in B or not. Another natural version, $\mathrm{FAN}_{\Pi^0_1}$ or Π^0_1 Fan, is to consider Π^0_1 **bars**: there is a decidable set $C \subseteq 2^{<\omega} \times \mathbb{N}$ such that $\sigma \in B$ iff, for all $n \in \mathbb{N}$, $(\sigma, n) \in C$. Nestled in between these two is FAN_c or c-Fan, which is based on the notion of a c**-bar**, which is a particular kind of Π^0_1 bar: for some decidable set $C \subseteq 2^{<\omega}$, $\sigma \in B$ iff every extension of σ is in C. It is easy to see that the implications

$$\mathrm{FAN}_{\mathrm{full}} \implies \mathrm{FAN}_{\Pi^0_1} \implies \mathrm{FAN}_c \implies \mathrm{FAN}_\Delta.$$

all hold over a weak base theory. What about the reverse implications? (We always include the implication of FAN_Δ from basic set theory when discussing the converses of the conditionals above.)

There had been several proofs that some of the converses did not hold [1,3,5]. These were piecemeal, in that each applied to only one converse, or even just a weak form of the converse, and used totally different techniques, so that there was no uniform view of the matter. This situation changed with [8], which provided a family of Kripke models showing the non-reversal of all the implications. It was asked there whether those models were in some sense the right, or canonical, models for this purpose; implicit was the question whether the other common modeling techniques, realizability and Heyting-valued models, could provide the same separations.

Here we do not answer those questions. We merely bring the discussion along, by providing a different kind of model. It should be pointed out early on that, at this point, the only separation provided is that FAN_Δ does not imply FAN_c, although we see no reason the arguments could not be extended to the other versions of Fan.

There are several ways that the model here differs from those of [8]. In the earlier paper, a tree with no simple paths was built over a model of classical ZFC via forcing, and the non-implications were shown by hiding that tree better or worse in various models of IZF. In particular, we showed there that FAN_Δ does not imply FAN_c by including that tree as the complement of a c-bar in a gentle enough way that no new decidable bars were introduced. Here, we start with a model of $\neg\mathrm{FAN}_\Delta$, and extend it by including paths that miss decidable (former) bars. If this is done to all decidable bars, FAN_Δ can be made to hold. If this is done gently enough, counter-examples to FAN_c will remain as counter-examples.

The other difference is in the techniques used. It is like a Kripke model built using Heyting-valued extensions of a realizability model. This is not the first time that some of these techniques have been combined (see [9] for references and discussion). This is the first time we are aware of that all three have been combined. Perhaps that in and of itself makes this work to be of some interest.

This work was started while the author was a fellow at the Isaac Newton Institute's fall 2015 program in the Higher Infinite. The author warmly thanks them for their support and hospitality during that time. Thanks are due also to Andrew Swan, a conversation with whom led to this work. Thanks go in addition to Francois Dorais and Noah Schweber for their input on Math Overflow about Francois's example of a c-bar which is not decidable.

2 Kripke structures of constructive models

A Kripke model is usually, if not always, built within a classical meta-theory from classical models. That is, working classically, given a partial order, to each node is associated a classical model (for the similarity type in question), along with a family of transition functions, and this determines a model of constructive logic. Our current setting is different.

A warning shot is given by the fact that the root of this model is, effectively, Kleene's recursive realizability model $\mathcal{M}_{K_1}$; to the degree we work within this model, we must work constructively, and not classically. This point could be finessed, though, by insisting that $\mathcal{M}_{K_1}$ was itself built within a classical theory.

More crucially, the structures at each node will be determined by Heyting-valued extensions of $\mathcal{M}_{K_1}$. These structures are no longer of the right type. A structure of the language of set theory would, among other things, determine, for objects a and b, whether $a \in b$. A Heyting-valued model, though, determines whether $\mathcal{H} \Vdash$ "$a \in b$", where $\mathcal{H}$ is a value in the Heyting algebra $\mathcal{T}$ in question.

To address these matters, a structure at a node will be, in addition to a Heyting-valued model, also a Heyting value $\mathcal{H}$ from the Heyting algebra, and we will let a formula be true there only if $\mathcal{H}$ forces it. To have forcing truth still be valid within this Kripke structure, we will allow an extension of a node to be determined in part by a strengthening of the Heyting value $\mathcal{H}$.

Before giving the formal definitions, we sketch the idea. We will need to iterate taking Heyting-valued extensions. By way of notation, $\mathcal{T}$ will be taken to be a typical complete Heyting algebra, to be consistent with the notation below. Some of these Heyting algebras $\mathcal{T}$ will show up only in some previously constructed Heyting-valued extension. So we assume we have a definable collection of Heyting algebras, say with definition $\phi(\mathcal{T})$, which IZF proves to be a set. To each node p will be associated a string $\langle(\mathcal{T}_0, \mathcal{O}_0), (\mathcal{T}_1, \mathcal{O}_1), ..., (\mathcal{T}_n, \mathcal{O}_n)\rangle$ such that each Heyting value $\mathcal{O}_i$ is not $\perp$ and forces that the next $\mathcal{T}_{i+1}$ is an allowable Heyting algebra, as given by ϕ. A child of p is determined by, optionally, extending some of the $\mathcal{O}_i$'s, and, optionally, including another pair $(\mathcal{T}_{n+1}, \mathcal{O}_{n+1})$ onto the string.

More formally, let $\phi(x)$ be a formula such that IZF proves "if $\phi(x)$ then x is a complete Heyting algebra, and ϕ is satisfied by only set-many objects." We will have occasion to consider ϕ as evaluated in a Heyting-valued extension, and so as applied to a term. Even if there are only set-many objects satisfying ϕ in this extension, there could still be class-many such terms. In order to keep this construction fully set-sized, we will implicitly allow only minimal terms $\mathcal{T}$ to be applied to ϕ. In more detail, suppose we assert in some context that $\phi(\mathcal{T})$. If this context is the ground model $\mathcal{M}_{K_1}$, then there are only set-many such $\mathcal{T}$'s within this model. Else the context will be some Heyting-valued extension, given by a Heyting value $\mathcal{H}$ – $\mathcal{H} \Vdash \phi(\mathcal{T})$ – within some other context. This latter context will then satisfy "for any other term t of rank less than that of $\mathcal{T}$, $\mathcal{H} \Vdash t \neq \mathcal{T}$". Similarly for the members of, or the Heyting values in, $\mathcal{T}$.

Definition 1. *Definition of the nodes, and their associated models, by induction on ω.*

The unique node of length 0 is the empty sequence $\langle\rangle$, with associated model $\mathcal{M}_{\langle\rangle} = \mathcal{M}_{K_1}$.

Inductively, given the set of nodes of length n, a node p of length $n+1$ will be a string of the form $\langle(\mathcal{T}_0, \mathcal{O}_0), (\mathcal{T}_1, \mathcal{O}_1), ..., (\mathcal{T}_n, \mathcal{O}_n)\rangle$ such that $p \restriction n$ is a node, and, in $\mathcal{M}_{K_1}$, $\mathcal{O}_0 \Vdash$ "$\mathcal{O}_1 \Vdash$ ..."$\mathcal{O}_{n-1} \Vdash$ "$\phi(\mathcal{T}_n)$ and $\mathcal{O}_n \neq \perp$ is a Heyting value in the Heyting algebra $\mathcal{T}_n$"" ...". The model $\mathcal{M}_p$ associated to p is the forcing extension by $\mathcal{T}_n$, with truth determined by $\mathcal{O}_n$, as evaluated within the model for $p \restriction n$.

We abbreviate the iterated forcing $\mathcal{O}_0 \Vdash$ "$\mathcal{O}_1 \Vdash$ ..."$\mathcal{O}_{n-1} \Vdash \psi$"..." as $\langle\mathcal{O}_0, ..., \mathcal{O}_{n-1}\rangle \Vdash_H \psi$. The reason for the subscript H is to emphasize that this notion of truth is given by iterated forcing. In contrast, for example, truth in $\mathcal{M}_{K_1}$ is given by realizers, and will be written as $e \Vdash_r \psi$. One important instance of that will be iterated forcing over $\mathcal{M}_{K_1}$. So truth in the model given by forcing with $\mathcal{T}$ over $\mathcal{M}_{K_1}$ would be written as $e \Vdash_r$ "$\mathcal{O} \Vdash_H \psi$".

By our various conventions, there are only set-many nodes.

By way of notation, we will typically suppress mention of the $\mathcal{T}_i$'s, as they are implicit in the choice of the $\mathcal{O}_i$'s. The opens of a node will (at least sometimes) be written as $\mathcal{O}^p_i$, so that p of length n will be $\langle \mathcal{O}^p_0, ..., \mathcal{O}^p_{n-1} \rangle$.

Definition 2. *The partial order on the set of nodes. For q to be an extension of p, written $q \geq p$, q has to be at least as long as p, and, for i less than the length of p, $q \restriction i \Vdash_H \mathcal{O}^q_i \leq \mathcal{O}^p_i$. (For this to make sense, implicitly $q \restriction i \Vdash_H \mathcal{T}^q_i = \mathcal{T}^p_i$.)*

We leave it to the reader to show that this is indeed a partial order. Notice that an extension of a node is indicated with the standard notation for partial orders, $\geq$, in contrast with the strengthening of a Heyting value, which is indicated with the standard notation for forcing, $\leq$.

This p.o. is in $\mathcal{M}_{K_1}$. Since a model embeds into any Heyting-valued extension, the p.o. is also in any of the models associated with a node. Furthermore, consider the p.o. restricted to a node (i.e. the extensions of any node, including itself). This restriction is definable in the node's model, uniformly from the node. That is, given any node as a parameter, the node's model can figure out the rest of the p.o.

We will need the notion of a node being covered by a set of nodes, akin to an open set in a topological space being covered by a collection of open sets, or, more generally, a member of a Heyting algebra being (less that) the join of a subset of the algebra.

Definition 3. *We define p of length n being covered by $P = \{p_j \mid j \in J\}$ by induction on n.*

- *For $n = 0, \langle\rangle$ is covered by only $\{\langle\rangle\}$.*
- *For $n = 1, p$ of the form $\langle(\mathcal{T}, \mathcal{O})\rangle$ is covered by P if each p_j also has length 1, and $p_j \geq p$ (so p_j is of the form $\langle(\mathcal{T}, \mathcal{O}_j)\rangle$, the point being that $\mathcal{T}$ is the same), and $\{\mathcal{O}_j \mid j \in J\}$ covers $\mathcal{O}$ in the sense of $\mathcal{T}$: $\mathcal{O} \leq \bigvee\{\mathcal{O}_j \mid j \in J\}$.*
- *For a length $n+1 > 1$, some conditions are immediate analogues: each p_j extends p in the Kripke order, and each p_j has length $n+1$. Furthermore, letting $P \restriction n$ be $\{p_j \restriction n \mid j \in J\}$, we have that $P \restriction n$ is to cover $p \restriction n$. Finally, we want to view P as a term for a set in the model associated with $p \restriction n$. Recall that a term for a Heyting-valued model is an arbitrary collection of pairs $\langle \mathcal{O}, \sigma\rangle$, where $\mathcal{O}$ is a member of the Heyting algebra and σ is (inductively) a term. If we are considering a two-step iteration, then σ is (a term for) a pair $\langle \hat{\mathcal{O}}, \tau\rangle$, where $\hat{\mathcal{O}}$ is a value from the second Heyting algebra. This can be abbreviated by $\langle(\mathcal{O}, \hat{\mathcal{O}}), \tau\rangle$. Whereas each p_j is of the form $\langle \mathcal{O}_0, \ldots, \mathcal{O}_n\rangle$, it induces a set $alt - p_j := \langle(\mathcal{O}_0, \ldots, \mathcal{O}_{n-1}), \mathcal{O}_n\rangle$, which is a term, in the language for an n-fold forcing iteration, with value (forced to be) an open set in $\mathcal{T}_n$. Of course, the n-fold iteration in question is just the model associated with $p \restriction n$. So, letting P_n be $\{alt - p_j \mid j \in J\}$, $p \restriction n \Vdash_H P_n$ is a collection of open sets of $\mathcal{T}_n$. Our final condition is that $p \restriction n \Vdash_H P_n$ covers $p(n)$.*

(At some point, we might need that the transitivity condition is satisfied: if P covers p, and each $p_j \in P$ is covered by P_j, then $\bigcup_{j \in J} P_j$ covers p.)

We are finally in a position to define the model. Working within $\mathcal{M}_{K_1}$, inductively on the ordinals α, we define the members $\mathcal{M}^p_\alpha$ of the model at node p of rank α (where we associate α with its canonical image in each of the associated models), along with the transition functions f_{pq} from $\mathcal{M}^p_\alpha$ to $\mathcal{M}^q_\alpha$. We will usually drop the subscripts and just write f as a polymorphic transition function. Similarly, we will not adorn f with any α, since the definition of f will be uniform in α. Do not confuse the associated models $\mathcal{M}_p$ from above with the $\mathcal{M}^p$ about to be defined.

Definition 4. *The universe $\mathcal{M}^p$ of the model at node p. First we define $\mathcal{M}^p_\alpha$ inductively on ordinals α. A member σ of $\mathcal{M}^p_\alpha$ is a function with domain the p.o. restricted to p (i.e. p and its extensions). Furthermore, $\sigma(q) \subseteq \bigcup_{\beta<\alpha} \mathcal{M}^q_\beta$. In order to fulfill the basic Kripke condition, if $\tau \in \sigma(q)$, and $r \geq q$, then $f(\tau) \in \sigma(r)$. If $q \geq p$, then $f(\sigma) = \sigma \upharpoonright \mathcal{P}^{\geq q}$. Let $\mathcal{M}^p$ be $\bigcup_\alpha \mathcal{M}^p_\alpha$.*

(If you're wondering whether there are any such members, or whether instead the definition is vacuous, see the proof of IZF below; the reader is invited to think through now why the empty set is a set within this formalism.)

Because this model has aspects of both a Kripke and a Heyting-valued model, it is in actuality neither. So to give the semantics, we cannot rely on any standard definition already extant in the literature. Rather, we have to give an independent, inductive definition of satisfaction. In this case we do not subscript $\Vdash$, because it is our main notion of truth; if we ever need to disambiguate, it will be written as $\Vdash_K$. Furthermore, we will refer to it as a Kripke model, because it is similar to one, and so we can distinguish it verbally from the various Heyting-valued models considered and from the realizability model in which this is all taking place.

Definition 5. *Implicitly in what follows, when we write "$p \Vdash \phi$", the parameters in ϕ are all in $\mathcal{M}^p$. Also implicit is the application of the transition function f, as need be. For future reference, $\Vdash$ and $\Vdash_K$ are the same thing.*

- *$p \Vdash \sigma \in \tau$ iff p is covered by some P, and for all $p_j \in P$ there is a $\sigma_j \in \tau(p_j)$ such that $p_j \Vdash \sigma = \sigma_j$.*
- *$p \Vdash \sigma = \tau$ iff for all $q \geq p$ and all $\rho \in \sigma(q)$, $q \Vdash \rho \in \tau$, and vice versa.*
- *$p \Vdash \phi \wedge \psi$ iff $p \Vdash \phi$ and $p \Vdash \psi$.*
- *$p \Vdash \phi \vee \psi$ iff p is covered by some Q, and for each $q \in Q$ either $q \Vdash \phi$ or $q \Vdash \psi$.*
- *$p \Vdash \phi \to \psi$ iff for all $q \geq p$ if $q \Vdash \phi$ then $q \Vdash \psi$.*
- *$p \Vdash \perp$ never.*
- *$p \Vdash \forall x\ \phi(x)$ iff for all $q \geq p$ and $\sigma \in \mathcal{M}^q$ $q \Vdash \phi(\sigma)$.*
- *$p \Vdash \exists x\ \phi(x)$ iff p is covered by some Q and for all $q \in Q$ there is some σ such that $q \Vdash \phi(\sigma)$.*

Lemma 1. *Each node satisfies the equality axioms.*

Lemma 2. *If Q covers p, and for each $q \in Q$ we have $q \Vdash \phi$, then $p \Vdash \phi$.*

Proof. By a straightforward induction on ϕ. □

Corollary 3. *Each node satisfies constructive logic.*

Proof. For the existential cases (membership, disjunction, existence). □

Theorem 4. *This structure models IZF.*

Proof. For many of the axioms, we will merely give the construction of the witness, and leave it to the reader to check that the witness given satisfies the conditions listed above to be in the universe.

Empty Set: In $\mathcal{M}_{K_1}$, the constant function with domain the entire tree always returning the empty set is an object of rank 0, and represents the empty set.

Infinity: First, we show that each $n \in \omega$ is canonically represented by a set in the model, inductively on n. The case $n = 0$ is done by the empty set, above. Suppose for $i \leq n$, i is represented in the model at the root by σ_i of rank i. Then $n+1$ will be represented by σ_{n+1}, where $\sigma_{n+1}(p)$ will be $\{f(\sigma_0), f(\sigma_1), ..., f(\sigma_n)\}$ (where once again f is the (polymorphic) transition function). Finally, ω will be represented by σ_ω, where $\sigma_\omega(p) = \{f(\sigma_n) \mid n \in \omega\}$.

Pair: If $\sigma, \tau \in \mathcal{M}^p$, then the pair is given by ρ, where $\rho(q) = \{f(\sigma), f(\tau)\}$.

Union: For $\sigma \in \mathcal{M}^p, \bigcup \sigma$ is given by $(\bigcup \sigma)(q) = \bigcup\{\tau(q) \mid \tau \in \sigma(q)\}$.

Extensionality: This follows fairly directly from the definition of forcing equality.

$\in$-Induction: All of the models are actually well-founded.

Power Set: Let $\sigma \in \mathcal{M}^p$. Then let $(\wp(\sigma))(q)$ be $\{\tau \in \mathcal{M}^q \mid \forall r \geq q \; \tau(r) \subseteq \sigma(r)\}$.

Separation: Given $\sigma \in \mathcal{M}^p$, and $\phi(x)$ with parameters from $\mathcal{M}^p$, let $(\mathrm{Sep}(\sigma, \phi))(q)$ be $\{\rho \in \sigma(q) \mid q \Vdash \phi(\rho)\}$.

Collection: Suppose $p \Vdash \forall x \in \sigma \; \exists \tau \; \phi(x, \tau)$. Working in $\mathcal{M}_{K_1}$, for each $q \geq p$ and $\rho \in \sigma(q)$, there is an α such that some τ in $\mathcal{M}^q_\alpha$ satisfies $q \Vdash \phi(\rho, \tau)$. By Collection in $\mathcal{M}_{K_1}$, there is a bounding set for all of the α's needed. This bounding set can be restricted to contain only ordinals, and then expanded to be an ordinal itself, say β. Let Σ be such that, for $q \geq p, \Sigma(q) = \mathcal{M}^q_\beta$, which suffices for a bounding set.

□

3 FAN$_\Delta$ does not imply FAN$_c$

For the moment, we will work simply under IZF.

Our primary task is now to define the right ϕ, the class of Heyting algebras we will use to build the nodes. They will be induced by the possible counterexamples B to FAN$_\Delta$: B is a decidable set of binary strings, but is not uniform. It is safe to assume that B is closed upwards. Mostly we're interested in when B is in addition a bar, there famously being such a creature in Kleene's recursive realizability model. The reason that we do not include being a bar in this definition is that would then be another condition to check before being able to use B. This is more than just a matter of convenience, or saving a little work. When we're working within a Heyting-valued extension of a realizability model, say, different conditions might decide whether B is a bar differently, and if B had to be a bar then we'd need to find an infinite path through those conditions along which B became a bar, meaning either there is such a path, or we'd have to find a non-uniform bar forcing such a path, and all of a sudden the thicket starts to look impenetrable. Although it seems unaesthetic to force paths that we really don't need, this is a small price to pay for having a theorem with a proof.

Let T be the complement of B. So T is a decidable, infinite tree. We will generically shoot a branch through T.

We will define a formal topology S from T. To help make this paper self-contained, we present a definition of a formal topology. Such definitions are not uniform in the literature. Here we will use the one from [6], sec. 2.1.

Definition 6. *A formal topology is a poset $(S, \leq)$ and a relation $\lhd$ between elements and subsets of S. (One should think of the elements of S as open sets, with $\leq$ as containment and $\lhd$ as covering.) The axioms are:*

- *if $a \in p$ then $a \lhd p$,*
- *if $a \leq b$ and $b \lhd p$ then $a \lhd p$,*
- *if $a \lhd p$ and $\forall x \in p\ x \lhd q$ then $a \lhd q$, and*
- *if $a \lhd p$ and $a \lhd q$ then $a \lhd \downarrow p \cap \downarrow q$,*

where $\downarrow p$ is the downward closure of p.

Definition 7. *The formal topology induced by B:*

Let B be a decidable, upwards-closed, non-uniform set of binary strings, and T its complement in $2^{<\omega}$. A member of S is a union of finitely many basic members of S. A basic member of S, $\mathcal{O}_\sigma$, is given by a node $\sigma \in T$, and is the set of all nodes in T compatible with σ, that is, all initial segments and extensions, when it is infinite. A witness that $\mathcal{O} \in S$, that is, a finite set Σ such that $\mathcal{O} = \bigcup_{\sigma \in \Sigma} \mathcal{O}_\sigma$, is called a base *for $\mathcal{O}$; note that bases are not unique. The partial order $\leq$ on S is just the subset relation $\subseteq$.*

A subset $\mathcal{U}$ of S covers $\mathcal{O} \in S$, $\mathcal{O} \lhd \mathcal{U}$, if it is not the case that there is no finite length n such that, for all $\sigma \in T$ of length n, either $\sigma \notin \mathcal{O}$ or, for some initial segment τ of σ and for some $\mathcal{O}_\mathcal{U} \in \mathcal{U}$, we have $\mathcal{O}_\tau \subseteq \mathcal{O}$ and $\mathcal{O}_\tau \subseteq \mathcal{O}_\mathcal{U}$. In symbols, $\mathcal{U}$ covers $\mathcal{O}$ iff

$$\neg\neg\exists n\ \forall \sigma \in T\ \mid \sigma \mid = n \rightarrow (\sigma \notin \mathcal{O} \vee \exists \tau \subseteq \sigma\ \exists \mathcal{O}_\mathcal{U} \in \mathcal{U}\ \mathcal{O}_\tau \subseteq (\mathcal{O} \cap \mathcal{O}_\mathcal{U})).$$

For any such n, we say that $\mathcal{U}$ covers $\mathcal{O}$ by length n.

Remarks: By choosing the base to be empty, $\emptyset \in S$.

If the set of nodes compatible with σ is finite, then σ does not determine an open set; alternatively, we could allow the induced set to be open, and it will be covered by the empty set.

For $\sigma \in T$ and $\mathcal{O} \in S$ it is decidable from a base for $\mathcal{O}$ whether $\sigma \in \mathcal{O}$.

Note that if $\mathcal{U}$ covers $\mathcal{O}$ by n then $\mathcal{U}$ covers $\mathcal{O}$ by any $k \geq n$. The reason for the double-negation in the definition of covering should become clear, when it is used, in Theorems 6 and 7.

Proposition 5. $(S, \leq, \lhd)$ *from above constitutes a formal topology.*

Proof. We will have need of the fact from propositional logic that if $(\bigwedge_i \phi_i) \to \neg\psi$ then $(\bigwedge_i \neg\neg\phi_i) \to \neg\psi$. To see this, from the first assertion, take the contrapositive twice, eliminating the double negation in front of $\neg\psi$. Then note that $\neg\neg(\bigwedge_i \phi_i)$ is equivalent with $\bigwedge_i(\neg\neg\phi_i)$.

1. Suppose $\mathcal{O} \in \mathcal{U}$; we need to show $\mathcal{U}$ covers $\mathcal{O}$. Let Σ be a base for $\mathcal{O}$. Let n be the length of the longest sequence in Σ. Then for all σ of length n, either there is an initial segment τ of σ in Σ, or there's not. In the latter case, $\sigma \notin \mathcal{O}$. In the former, $\mathcal{O}_\mathcal{U}$ can be chosen to be $\mathcal{O}$ itself.

2. Suppose $\mathcal{O}_1 \subseteq \mathcal{O}_0$ and $\mathcal{U}$ covers $\mathcal{O}_0$. We need to show $\mathcal{U}$ covers $\mathcal{O}_1$. We can assume that we have bases Σ_0 and Σ_1 for $\mathcal{O}_0$ and $\mathcal{O}_1$ respectively such that no $\sigma_0 \in \Sigma_0$ extends any $\sigma_1 \in \Sigma_1$. Assuming that $\mathcal{U}$ covers $\mathcal{O}_0$ by some length n, we will find a k such that $\mathcal{U}$ covers $\mathcal{O}_1$ by k, which suffices, by taking the double contrapositive. Let m be the length of the longest $\sigma \in \Sigma_1$. Let k be the larger of m and n. Consider any σ of length k. If $\sigma \notin \mathcal{O}_1$, then we are done. Else consider the initial segment ρ of σ which is in Σ_1. Also consider $\sigma \restriction n \in \mathcal{O}_1$; recalling that $\mathcal{O}_1 \subseteq \mathcal{O}_0$, we conclude that $\sigma \restriction n \in \mathcal{O}_0$. By the choice of n, let $\tau \subseteq \sigma \restriction n$ and $\mathcal{O}_\mathcal{U} \in \mathcal{U}$ be such that $\mathcal{O}_\tau \subseteq (\mathcal{O}_0 \cap \mathcal{O}_\mathcal{U})$. If ρ is an initial segment of τ, then $\mathcal{O}_\tau \subseteq \mathcal{O}_1$ and the same τ and $\mathcal{O}_\mathcal{U}$ suffice. Else τ is an initial segment of ρ, and $\mathcal{O}_\rho$ is a subset of both $\mathcal{O}_\tau$ and $\mathcal{O}_1$, so use $\mathcal{O}_\rho$ and $\mathcal{O}_\mathcal{U}$.

3. Suppose that $\mathcal{U}$ covers $\mathcal{O}$, and that every $\mathcal{O}_\mathcal{U} \in \mathcal{U}$ is covered by $\mathcal{V}$. We need to show that $\mathcal{V}$ covers $\mathcal{O}$.

Being very careful with the logic here, work under the assumption that every $\mathcal{O}_\mathcal{U} \in \mathcal{U}$ is covered by $\mathcal{V}$. In showing that $\mathcal{U}$ covering $\mathcal{O}$ implies that $\mathcal{V}$ covers $\mathcal{O}$, we can assume that $\mathcal{U}$ covers $\mathcal{O}$ by some fixed length n, by taking the contrapositive twice. We also take n to be at least as big as any string in some base Σ for $\mathcal{O}$. For each of the finitely many σ's of length n that are in $\mathcal{O}$ let τ_σ and $\mathcal{O}_{\mathcal{U}\sigma}$ be as given by the definition of covering. Each such $\mathcal{O}_{\mathcal{U}\sigma}$ is covered by $\mathcal{V}$, which means there is not not an n_σ as in the definition of covering. By the remarks at the beginning of this proof, in showing that $\mathcal{V}$ covers $\mathcal{O}$, we may assume that for each such σ there is such an n_σ.

Let m be the maximum of n and the n_σ's. We will show $\mathcal{V}$ covers $\mathcal{O}$ by m. Let ρ be of length m. If $\rho \in \mathcal{O}$, let $\sigma = \rho \restriction n$. Then $\tau_\sigma \subseteq \sigma$, $\mathcal{O}_{\mathcal{U}_\sigma} \in \mathcal{U}$, and $\mathcal{O}_{\tau_\sigma} \subseteq (\mathcal{O} \cap \mathcal{O}_{\mathcal{U}_\sigma})$. In particular, $\rho \in \mathcal{O}_{\mathcal{U}_\sigma}$. Since $\mathcal{V}$ covers $\mathcal{O}_{\mathcal{U}_\sigma}$ by n_σ, there is a $\mu \subseteq \rho \restriction n_\sigma$ and $\mathcal{O}_\mathcal{V} \in \mathcal{V}$ with $\mathcal{O}_\mu \subseteq (\mathcal{O}_{\mathcal{U}_\sigma} \cap \mathcal{O}_\mathcal{V})$. Letting ν be the longer of μ and τ_σ, $\mathcal{O}_\nu \subseteq (\mathcal{O} \cap \mathcal{O}_\mathcal{V})$, which suffices.

4. Suppose $\mathcal{O}$ is covered by both $\mathcal{U}$ and $\mathcal{V}$. We need to show that $\mathcal{O}$ is covered by $\{\mathcal{O}' \mid \exists \mathcal{O}_\mathcal{U} \in \mathcal{U}\ \mathcal{O}' \subseteq \mathcal{O}_\mathcal{U} \text{ and } \exists \mathcal{O}_\mathcal{V} \in \mathcal{V}\ \mathcal{O}' \subseteq \mathcal{O}_\mathcal{V}\}$.

We can assume that both $\mathcal{U}$ and $\mathcal{V}$ cover $\mathcal{O}$ by n. Let $\sigma \in \mathcal{O}$ have length n. Let τ and $\mathcal{O}_\mathcal{U}$ be as given by $\mathcal{U}$ covering $\mathcal{O}$, and ρ and $\mathcal{O}_\mathcal{V}$ be as given by $\mathcal{V}$ covering $\mathcal{O}$. Let ν be the longer of ρ and τ. Then ν and $\mathcal{O}_\nu$ are as desired. □

The reason for this formal topology is so that we can take the Heyting-valued model $\mathcal{M}_T$ over it.

We do not know whether the next theorem is true in general (meaning provable in IZF). So for the moment, we work in the recursive realizability model. That is, the model $\mathcal{M}_T$ is taken as being built within it.

Theorem 6. *Working within the recursive realizability model, in $\mathcal{M}_T$, the generic G is (identifiable with) an infinite branch through T.*

Proof. We can identify the generic G with $\{\langle \mathcal{O}_\sigma, \tau\rangle \mid \tau \subseteq \sigma, \mathcal{O}_\sigma$ a basic open set$\}$. We want to show that $\mathcal{O}_\emptyset \Vdash_H$ "for all k there is a unique σ of length k with $\sigma \in G$." Since the natural numbers in the sense of $\mathcal{M}_T$ can be identified with those of $\mathcal{M}_{K_1}$, which are themselves just those of V, it suffices to fix a k in the sense of V. It is easy to see that if $\mathcal{O}_\sigma$ is a basic open set with σ of length k then $\mathcal{O}_\sigma \Vdash_H$ "σ is the unique member of G of length k." Let $\mathcal{U}$ be $\{\mathcal{O}_\sigma \mid \sigma$ has length k and $\mathcal{O}_\sigma$ is a basic open set$\}$. It suffices to show that $\mathcal{U}$ covers $\mathcal{O}_\emptyset$.

Because of the double negation in the definition of covering, when showing that $\mathcal{U}$ covers $\mathcal{O}_\emptyset$ it is not necessary to get the n as a computable function of k; rather, any realizer will do. So it's just a matter of finding an n in the ground model V such that the rest (of the definition of covering) is easily seen to be forced. Toward this end, let n be large enough so that, whenever T beneath σ of length k is finite, T contains no descendants of σ of length n. In other words, go through level k of T, take all those nodes whose subtrees will eventually die, of which there are only finitely many, and then go out far enough that all of them have died already. Now given a node τ of T of length n, $\tau \restriction k$ and $\mathcal{O}_{\tau\restriction k}$ are the desired witnesses. □

The preceding lemma will help us with the analysis of the Kripke model at nodes of length 1. Of course, we need to consider longer nodes too. Hence we must prove the corresponding lemma for base models a finite iteration of these extensions.

Theorem 7. *Let p be some node of the Kripke model, with associated model $\mathcal{M}_p$. Suppose that, within $\mathcal{M}_p$, T is an infinite, decidable tree. Then, working within $\mathcal{M}_p$, in $\mathcal{M}_T$, the generic G is (identifiable with) an infinite branch through T.*

Proof. For ease of exposition, we take p to be $\langle (T_0, \mathcal{O}_{\langle\rangle})\rangle$; that is, to have length 1 and to have the open set be the entire tree. We leave the general case to the reader.

By way of notation, let $\mathcal{O}^0$ refer to open sets in the formal topology in $\mathcal{M}_{K_1}$ induced by T_0, and $\mathcal{O}^1$ refer to open sets in the formal topology in $\mathcal{M}_p$ for T. As in the previous theorem, the natural numbers of $\mathcal{M}_T$ can be identified with those of V. So let k be a natural number, and in $\mathcal{M}_p$ let $\mathcal{U}^1$ be $\{\mathcal{O}^1_\sigma \mid \sigma$ has length k and $\mathcal{O}^1_\sigma$ is a basic open set$\}$. As before, it suffices to show that $\mathcal{U}^1$ covers $\mathcal{O}^1_\emptyset$; actually, we must show that $\mathcal{O}^0_\emptyset \Vdash_H$ "$\mathcal{U}^1$ covers $\mathcal{O}^1_\emptyset$"; actually, what we really must show is that there is some e which realizes the above forcing assertion, uniformly in k.

Recall that "$\mathcal{U}^1$ covers $\mathcal{O}^1_\emptyset$" is an abbreviation of "not not there is an n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_\emptyset$ by n." So we must realize "for every $\mathcal{O}^0_\sigma$ extending $\mathcal{O}^0_\emptyset, \mathcal{O}^0_\sigma \not\Vdash_H$ there is no n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n." Toward that end, suppose $f \Vdash_r \mathcal{O}^0_\sigma$ is an open set of the space T_0. Then we must have chosen e so that $\{e\}(k, f) \Vdash_r$ "$\mathcal{O}^0_\sigma \not\Vdash_H$ there is no n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n." By the realizability semantics, everything realizes a negation, as long as nothing realizes the statement being negated. So we must show, in V, that nothing realizes "$\mathcal{O}^0_\sigma \Vdash_H$ there is no n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n."

Suppose, toward a contradiction, that g does realize that statement. Unpacking the semantics further, $g \Vdash_r$ "for every extension $\mathcal{O}^0_\tau$ of $\mathcal{O}^0_\sigma$, $\mathcal{O}^0_\tau \not\Vdash$ there is an n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n." Let $h \Vdash_r \mathcal{O}^0_\tau$ extends $\mathcal{O}^0_\sigma$. Then $\{g\}(h) \Vdash_r$ "$\mathcal{O}^0_\tau \not\Vdash$ there is an n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n." The only way to realize a negation is if nothing realizes the statement being negation. So nothing realizes "$\mathcal{O}^0_\tau \Vdash$ there is an n such that $\mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n." We will have our desired contradiction once we find a $\mathcal{O}^0_\tau$ extending $\mathcal{O}^0_\sigma$, an integer n and a realizer of "$\mathcal{O}^0_\tau \Vdash \mathcal{U}^1$ covers $\mathcal{O}^1_{\langle\rangle}$ by n."

We will construct a finite sequence of basic open sets of T_0, starting with $\mathcal{O}^0_\sigma$ and ending with the desired $\mathcal{O}^0_\tau$. The steps of this procedure are indexed by the binary sequences of length k. At each step, indexed by say ρ, extend the current open set if possible to force T beneath ρ to be finite, and then again to force a level n_ρ witnessing this finiteness (i.e. that ρ has no extension in T of length n_ρ); whenever this is not possible, the open set at hand already forces T beneath ρ to be infinite. Let n be the maximum of the n_ρ's. Working beneath $\mathcal{O}^0_\tau$, if π of length n is ever forced to be in T, then $\pi \restriction k$ and $\mathcal{O}^1_{\pi\restriction k}$ will be as desired.

□

Corollary 8. *For p a node in the Kripke partial order, with final entry $(T, \mathcal{O})$, and B the complement of T, $p \Vdash_K B$ is not a bar.*

Proof. Let G be the generic for forcing with T. The function (with domain the partial order from p onwards) with constant output G (more accurately, the canonical image of G in the input's associated model) witnesses that B is not a bar. □

The next two theorems finish this paper.

Theorem 9. $\mathcal{M} \models FAN_\Delta$.

Proof. The idea is simple enough. If, at a node, B is forced to be a decidable bar, then B must also be forced to be uniform, because, if not, the node would have an extension given by forcing with the complement of B, showing that B could not have been a bar. We need to check the details though, to guard against things like the use of classical logic and to make sure we're using the semantics of the model at hand. For better or worse, I know of no other way to do this than to unravel the statement to be shown, using the semantics given.

We need to show $\langle\rangle \Vdash \mathrm{FAN}_\Delta$, working within $\mathcal{M}_{K_1}$, meaning we must find a realizer e for the statement $\langle\rangle \Vdash \mathrm{FAN}_\Delta$. As a reminder, $\langle\rangle$ is the empty sequence, the bottom node in the partial order underlying the model. For reference, FAN_Δ is the assertion "for all B, if B is an upwards-closed decidable bar (in $2^{<\omega}$), then B is uniform, i.e. there is a natural number n such that all binary sequences of length n are in B."

The Hypothesis: Unpacking the meaning of $\Vdash$, we need to show that within $\mathcal{M}_{K_1}$, if $B \in \mathcal{M}^p$ then $p \Vdash$ "if B is such a bar then B is uniform." That means that if $t \Vdash_r p$ is a node and $B \in \mathcal{M}^p$ then $\{e\}(t) \Vdash_r$ "$p \Vdash$ "if B is such a bar then B is uniform"." To save on notation, we will suppress mention of t. This means that we must show $e \Vdash_r$ "for all $q \geq p$, if $q \Vdash B$ is such a bar then $q \Vdash B$ is uniform." Again suppressing the realizer that $q \geq p$, we must show $e \Vdash_r$ "if $q \Vdash B$ is such a bar then $q \Vdash B$ is uniform." So, suppose $f \Vdash_r$ "$q \Vdash B$ is such a bar;" we must have that $\{e\}(f) \Vdash_r$ "$q \Vdash B$ is uniform."

Unpacking some more, there is a realizer g, easily computable from f, with $g \Vdash_r$ "$q \Vdash B$ is decidable;" and that means $g \Vdash_r$ "$q \Vdash$ for all $\sigma \in 2^{<\omega}$, either $\sigma \in B$ or $\sigma \notin B$." Since $2^{<\omega}$ does not change from node to node, that means

$g \Vdash_r$ "for all $\sigma \in 2^{<\omega}, q \Vdash$ (either $\sigma \in B$ or $\sigma \notin B$)." Identifying a realizer that $\sigma \in 2^{<\omega}$ with σ itself, that becomes "for all $\sigma \in 2^{<\omega}$ $\{g\}(\sigma) \Vdash_r q \Vdash (\sigma \in B \vee \sigma \notin B)$." And that means that "for all $\sigma \in 2^{<\omega}$ there is a Q_σ such that $\{g\}(\sigma) \Vdash_r$ (Q_σ covers q and for each $r \in Q_\sigma$ either $r \Vdash \sigma \in B$ or $r \Vdash \sigma \notin B$)."

The Conclusion: Having just unpacked the hypothesis, we will now analyze the conclusion. Recall what we need to show: $\{e\}(f) \Vdash_r$ "$q \Vdash B$ is uniform;" i.e. $\{e\}(f) \Vdash_r$ "$q \Vdash$ there is a bound n witnessing that B is uniform;" which is $\{e\}(f) \Vdash_r$ "there is a cover R of q, and for all $r \in R$ there is some object n such that $r \Vdash n$ is a natural number witnessing the uniformity of B." That comes down to the existence of a set R such that $\{e\}(f) \Vdash_r$ "R covers q, and for all $r \in R$ there is some object n such that $r \Vdash n$ is a natural number witnessing the uniformity of B."

Since this is complicated enough, we're now going to build up somewhat slowly. We will examine several cases, based on the length of q. Since e has access to a realizer that q is a node, e has access to q's length, and so can make this case distinction.

Case I: For the first pass, suppose $q = \langle\rangle$. Every element in the cover Q_σ must have the same length as the thing covered, which in this case is 0, so $Q_\sigma = \{\langle\rangle\}$. So we have $\{g\}(\sigma) \Vdash_r$ (either $\langle\rangle \Vdash \sigma \in B$ or $\langle\rangle \Vdash \sigma \notin B$)." Similarly for R: we must have $\{e\}(f) \Vdash_r$ "there is some n such that $\langle\rangle \Vdash n$ is a natural number witnessing the uniformity of B." The obvious algorithm to find a uniform bound for B is to run through the various $\{g\}(\sigma)$'s until one finds such a bound. All there is left to do in this case is to show that this algorithm terminates. If not, then every K_1 realizer will realize that B is not uniform. So, in $\mathcal{M}_{K_1}$, letting T be the complement of B, $\langle(S_T, \mathcal{O}_{\langle\rangle})\rangle$ is a node (where S_T is the formal topology induced by T). By the corollary, $\langle(S_T, \mathcal{O}_{\langle\rangle})\rangle \Vdash$ "B is not a bar," contradicting the assumption that $f \Vdash_r$ "$\langle\rangle \Vdash B$ is (such) a bar."

Case II: For our second pass, suppose that q has length 1. So the notion that Q_σ as a set of nodes covers q as a node reduces to covering in the sense of the Heyting algebra S_T, where $q = \langle(S_T, \mathcal{O})\rangle$. Letting $\mathcal{U}_\sigma$ be the set of second components of the members of Q_σ (actually, a member of Q_σ being a sequence of length 1, we are identifying such a sequence with its sole entry), $\{g\}(\sigma)$ yields a realizer that $\mathcal{U}_\sigma$ covers $\mathcal{O}$ in the sense of S_T. Now recall that the definition of covering begins with a double negation; unpacking what it is to realize a double negation, since something realizes that $\mathcal{U}_\sigma$ covers $\mathcal{O}$, everything does; and there is a realizer that $\mathcal{U}_\sigma$ covers $\mathcal{O}$ by some witnessing length n, but those latter two items are not uniformly computable from $\{f\}(\sigma)$.

We need to define a set R with which we can work easily. Towards this end, say that a binary sequence ρ is sufficiently long if it is at least as long as each σ in the given fixed base Σ for $\mathcal{O}$. (That is, e computes a realizer that q is a node. Being a node includes that $\mathcal{O}$ is open in S_T, meaning that $\mathcal{O}$ has a finite base, which is then witnessed by the realizers at hand.) By way of notation, for $r \geq q$ of length 1, r will be $\langle(S_T, \mathcal{O}_r)\rangle$. In V, let R be $\{\langle h, r\rangle \mid h \Vdash_r$ "$r \geq q, \mathcal{O}_r = \mathcal{O}_\rho \leq \mathcal{O}, \rho$ is sufficiently long, and there is an $n \leq\mid \rho \mid$ such that $r \Vdash n$ is a natural number witnessing the uniformity of B"$\}$. (The reason to insist that $r \leq q$ is that, in order for r to force B to be uniform, r already has to force that B is a set of binary strings; recalling that our assumption is no more than that q forces B to be a bar, we remain on the safe side by working beneath q.)

From a realizer that r is in R, it is easy (and uniform) to compute an n and a realizer that r forces n to witness the uniformity of B. So there is a uniform realizer, not even depending on f, of the second part of what $\{e\}(f)$ is supposed to realize. Take such a realizer for the second part of $\{e\}(f)$. As for the first part of $\{e\}(f)$, it is supposed to realizes that R covers q. Working in $\mathcal{M}_{K_1}$, let $\mathcal{U}$ be $\{\mathcal{O}_r \mid r \in R\}$. We need to realize that $\mathcal{U}$ covers $\mathcal{O}$.

Recall that the notion of covering begins with a double negation. So if anything realizes that, then everything does. Hence we need only the existence of an n and a realizer that n witnesses that $\mathcal{U}$ covers $\mathcal{O}$ by length n; n and the realizer do not have to be computed.

In V, either there are such an n and a realizer, or there aren't. If there are, we are done (with this case of q having length 1). If not, then we would like to force with the complement of $\mathcal{U}$. Now, it does us little good to do such forcing over $\mathcal{M}_{K_1}$; rather, we need to work beneath q, or, more accurately, in $\mathcal{M}_T$. Toward this end, let B_R be the term in $\mathcal{M}_{K_1}$ for the $\mathcal{M}_T$-set which is (the canonical embedding of) $\{\rho \mid \mathcal{O}_\rho \in \mathcal{U} \vee \rho \notin \mathcal{O}\}$. We would like B_R to induce an extension of $\mathcal{M}_T$. That is, for the complement T_R of B_R, we would like that $q^\frown\langle(S_{T_R}, \mathcal{O}_{\langle\rangle})\rangle$ is an extension of q. (Please note that $T_R = \mathcal{O}\backslash\{\rho \mid \mathcal{O}_\rho \in \mathcal{U}\}$.) We will argue later why that essentially does it for this case; in the meantime, our intermediate goal is to show that $\mathcal{O} \Vdash_H$ "B_R is a decidable, upwards closed set of binary strings, which is not uniform."

Easily, B_R is forced to be an upwards closed set of binary strings.

Turning to B_R being decidable, it is not clear that it (rather, its $\mathcal{M}_{K_1}$ version) is so in $\mathcal{M}_{K_1}$. After all, the definition of B_R depends on R, which itself depends on B; even though p forces B to be decidable, it is not clear that that translates to an algorithm for deciding facts about forcing the uniformity of B. Serendipitously, we do not need decidability of B_R in $\mathcal{M}_{K_1}$, but only in $\mathcal{M}_T$ (which, recall, is q's associated model), as forced by $\mathcal{O}$. We will be able to leverage the difference between decidability in $\mathcal{M}_{K_1}$ and in $\mathcal{M}_T$ to get the latter.

Unpacking the assertion $\mathcal{O} \Vdash_H$ "B_R is decidable," we must show that for all $\rho \in 2^{<\omega}$, $\mathcal{O} \Vdash \rho \in B_R \vee \rho \notin B_R$. And that means that for all such ρ there is a cover R_ρ of $\mathcal{O}$ such that, for all $\mathcal{O}_s \in R_\rho$, either $\mathcal{O}_s \Vdash \rho \in B_R$ or $\mathcal{O}_s \Vdash \rho \notin B_R$. Given such a ρ, first determine whether $\rho \in \mathcal{O}$, using the decidability of T and the finite base of $\mathcal{O}$. If so, then determine whether ρ is sufficiently long. If so, consider all binary σ of the same length as ρ. Using the meet operation of the Heyting-algebraic structure of S_T, applied to the (finitely many) $\mathcal{U}_\sigma$'s, there is a cover Q of $\mathcal{O}$ each member of which decides B of level $\mid \rho \mid$ (i.e. decides membership in B of all strings of the same length as ρ). It is not hard to see that this Q suffices for our R_ρ.

With decidability out of the way, we now show that $\mathcal{O} \Vdash_H B_R$ is not uniform. That is, we are interested in evaluating whether B_R is uniform in $\mathcal{M}_T$. That means that for no $\hat{\mathcal{O}} \leq \mathcal{O}$ do we have that $\hat{\mathcal{O}} \Vdash B_R$ is uniform. Suppose to the contrary we had such an $\hat{\mathcal{O}}$. Then any witness to the uniformity of B_R is then a witness to the uniformity of $\mathcal{U}$ as a cover of $\mathcal{O}$, which is assumed not to exist. So $\mathcal{O} \Vdash_H B_R$ is not uniform.

So we have an extension node of q, achieved by forcing over T_R. That induces a generic path through T_R. From this path, we can interpret B as a non-uniform tree. B is already assumed to be forced to be upwards closed and decidable. So now B induces an extension in the Kripke partial order. This gets a path avoiding B. So B cannot be forced to be a bar. This contradiction finishes the proof of this case.

Case III: It is time to finish the proof of this theorem. So, let q be a node. If need be, we will work inductively on the length of q, so q can be taken to have length at least 1. To recall, we have a realizer f that $q \Vdash B$ is an upwards closed decidable bar. We are searching for a realizer that $q \Vdash B$ is uniform. As above, we will need a set R such that it is easily realizable that R covers q, and for all $r \in R$ there is an n such that $r \Vdash n$ witnesses the uniformity of B. By our notational conventions, we can omit mention of the spaces $\mathcal{T}_i$ in q, and consider q to be $\langle \mathcal{O}^q_0, \ldots, \mathcal{O}^q_m \rangle (m \geq 0)$; similarly for extensions of q. This means that nodes can also be taken to be iterated forcing conditions, thereby eliminating the need for notation to translate between nodes and their corresponding conditions.

For $r \leq q$ as a forcing condition, we say that r is in normal form if (for all i) $r \restriction i \Vdash_H \mathcal{O}^r_i = \mathcal{O}_{\rho_i}$, for some specific $\rho_i \in 2^{<\omega}$. Let the length of r be the maximum of m and the length of each ρ_i occurring in r. Let R be $\{\langle h, r \rangle \mid h \Vdash_r$ "r is in normal form, and for some n at most the length of r, $r \Vdash n$ witnesses the uniformity of B"$\}$. All that remains to do is to realize that R covers q, the rest of what needs showing being trivial. For that, we show by induction on n up to $m+1$, that $R \restriction n$ covers $q \restriction n$.

For $n = 0$, this is just that R is non-empty. This is a simple case, which can be handled by methods similar to those that follow, and so is left to the reader. For $n \neq 0$, we need to show that $q \restriction n \Vdash_H R_n$ covers $q(n)$, i.e. $q \restriction n \Vdash_H$ " not not there is a witness k that R_n covers $q(n)$ by length k." Toward that end, let r extend $q \restriction n$ (in the iterated forcing $\mathcal{T}_0 \times \ldots \times \mathcal{T}_{n-1}$). We need to show that $r \not\Vdash_H$ "there is no such witness k." Aiming toward a contradiction, suppose instead that $r \Vdash_H$ "there is no such witness k."

Now viewing r and q as nodes in the Kripke partial order, let $r \vee q$ be the least common extension of r and q in this p.o. We would like to show that $r \vee q \Vdash_K$ "there is no such k" (i.e. in the sense of the Kripke model). Let $s \geq_K r \vee q$. Assume toward a contradiction that $s \Vdash_K$ there is such a k. Since s is covered by a set forcing a particular such witness, we can assume that s itself forces that R_n covers $q(n)$ by length k. Now extend s in the Kripke p.o. (to, by abuse of notation, s) to decide on membership in $q(n)$ of each σ of length k, and, for each such σ forced to be in $q(n)$, witnesses $\tau \subseteq \sigma$ and $\mathcal{O} \in R_n$ to the covering property. Keeping in mind that $q(n)$ and R_n are sets in the model associated with $q \restriction n$, these same facts about R_n and $q(n)$ are therefore forced by $s \restriction n$ in the sense of iterated forcing. Since $s \restriction n \leq r$, this is the desired contradiction. Hence we conclude that $r \vee q \Vdash_K$ "there is no such k."

Let B_{R_n} be a term (in $r \vee q$'s associated model) for $\{\rho \mid \mathcal{O}_\rho \in R_n$ or $\rho \notin q(n)\}$. We would like to show that $r \vee q \Vdash_H$ "B_{R_n} is a decidable, upwards closed, non-uniform set of binary strings." Once we do that, B_{R_n} induces a Heyting-valued extension, in which B is not witnessed to be uniform. So then B induces a Heyting-valued extension in which it is not a bar. This is the ultimate contradiction toward which we were aiming.

It is easy to see that B_{R_n} is forced to be an upwards closed set of binary strings. For the non-uniformity, use the previous result that there is no witness k to R_n covering $q(n)$.

All we need left to do in this entire [1] proof is to show that B_{R_n} is forced to

[1] If the reader is feeling any frustration at the length and detail of this proof, it might

be decidable. As in the previous case (for q of length 1), for any finite binary string ρ, we must find a set S covering $r \vee q$ such that each $s \in S$ forces either $\rho \in B_{R_n}$ or $\rho \notin B_{R_n}$. As before, cover $r \vee q$ with a set, each member of which decides on a finite base for q, as well as on B up to the level of the length of ρ. This suffices.

□

Theorem 10. $\mathcal{M} \models \neg FAN_c$.

Proof. Recall that a c-fan is based on a decidable set of C, which can be taken to be a computable assignment of "in" and "out" to all the nodes. A node is in the bar if it and all of its successors are assigned "in", and out of the bar, or in the tree, if one of its successors is "out".

Consider the following c-fan, due to Francois Dorais. Let K be some complete c.e. set, with enumeration K_s (K at stage s). Let C be such that all nodes on level n are labeled "in" except for the unique node consistent with K_n (i.e. convert K_n into a characteristic function). It is easy to see that the characteristic function of K, if it exists, is the unique branch missing the induced c-set B, which is a bar in the realizability model, because, as is well known, K's characteristic function does not exist in that model. We must, and need only, show that B remains a bar in $\mathcal{M}$, the idea being that generically K is not added by forcing to the model.

Suppose toward a contradiction that $p \Vdash K$ witnesses that B is not a bar, in that K is an infinite path avoiding B. What that comes down to is that for every natural number n there is a covering Q_n of p such that every $q \in Q$ forces a unique binary string of length n to be in K and that all such strings in K cohere. It is easy to see that p could not be $\langle\rangle$, because $\langle\rangle$ is covered by only $\{\langle\rangle\}$, at which point K could be externalized from $\mathcal{M}^{\langle\rangle}$ to $\mathcal{M}_{K_1}$, which as already remarked is not the case.

More generally, when $q \Vdash \sigma \in K$, then σ really is an initial part of K's characteristic function. After all, if not, then at some stage s an element k is enumerated into K such that $\sigma(k) = 0$. So all descendants of σ of length $t \geq s$ are in C, and hence in the bar B. That would mean that q could not force an extension of σ of length s to be in K, contradicting K being forced to be an infinite path. Hence any such σ in an initial segment of K. That means that all such σ's among all such q's cohere. By taking their union, the characteristic function of K can be built in $\mathcal{M}_{K_1}$. □

References

[1] Michael Beeson, *Foundations of Constructive Mathematics*, Springer-Verlag, 1985.

[2] Josef Berger, *The logical strength of the uniform continuity theorem*, Logical Approaches to Computational Barriers, Proceedings of CiE 2006, LNCS 3988, 2006, pp. 35–39.

[3] ———, *A separation result for varieties of Brouwer's Fan Theorem*, Proceedings of the 10th Asian Logic Conference (ALC 10), Kobe University in Kobe, Hyogo, Japan, September 1-6, 2008, 2010, pp. 85–92.

[4] Hannes Diener and Iris Loeb, *Sequences of real functions on [0, 1] in constructive reverse mathematics*, Annals of Pure and Applied Logic **157(1)** (2009), pp. 50-61.

amuse them to know that in an earlier draft, the spot above contained an expletive.

[5] M. Fourman and J. Hyland, *Sheaf models for analysis*, Applications of Sheaves (Michael Fourman, Christopher Mulvey, and Dana Scott, eds.), Lecture Notes in Mathematics, vol. 753, Springer Berlin / Heidelberg, 1979, pp. 280-301.

[6] Nicola Gambino, *Heyting-valued interpretations for Constructive Set Theory*, Annals of Pure and Applied Logic **137** (2006), pp. 164-188.

[7] William Julian and Fred Richman, *A uniformly continuous function on [0,1] that is everywhere different from its infimum*, Pacific Journal of Mathematics **111(2)** (1984), pp. 333-340.

[8] Robert Lubarsky and Hannes Diener, *Separating the Fan Theorem and its weakenings*, Journal of Symbolic Logic **79** (2014), pp. 792-813.

[9] Jaap van Oosten, *Realizability: An Introduction to its Categorical Side*, Elsevier, 2008.