

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

January 28, 2016

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 an entire 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, 6]. 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, $\text{FAN}_{\Pi_1^0}$ or Π_1^0 Fan, is to consider Π_1^0 **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_1^0 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

$$\text{FAN}_{\text{full}} \implies \text{FAN}_{\Pi_1^0} \implies \text{FAN}_c \implies \text{FAN}_\Delta.$$

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 [7], 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 [7]. 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 $\neg\text{FAN}_\Delta$, 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 a Kripke model within a Heyting-valued extension of a realizability model. This is not the first time that some of these techniques have been combined (see [8] 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 draft is being prepared for the Isaac Newton Institute's pre-print series, as a result of their fall 2015 program in the Higher Infinite. The author warmly thanks them for their support and hospitality during this program, when this work was started. In an effort to have it available in a timely fashion, this write-up is in some parts only a sketch. Of course it is incumbent upon us to complete this in the near future. The hope is that it's already far enough along to be convincing, or, failing that, at least far enough along to be plausible.

Thanks are due to Andrew Swan, a conversation with whom led to this work. Thanks also go 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 FAN_Δ does not imply FAN_c

For the moment, we will work simply under IZF.

Suppose B is a counter-example to FAN_Δ : B is a decidable bar, but is not uniform. It is safe to assume that B is closed upwards. Let T be the complement of B . So T is a decidable, infinite tree with no infinite branch. We will generically shoot a branch through T .

Define a formal topology on T as follows. A basic open set \mathcal{O}_σ given by a node $\sigma \in T$ is the set of all nodes in T compatible with σ , that is, all initial

segments and extensions, when it is infinite (otherwise σ does not determine an open set, or, arguably, the bottom or empty set). An open set \mathcal{O} is a union of finitely many basic open sets. Note that this means it is decidable whether $\sigma \in \mathcal{O}$. A witness that \mathcal{O} is open, 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. A collection of open sets \mathcal{U} covers \mathcal{O} 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 . Note that if \mathcal{U} covers \mathcal{O} by n then \mathcal{U} covers \mathcal{O} by any $k \geq n$.

Proposition 1. *This constitutes a formal topology.*

Proof. We will have need of the fact from propositional logic that if $(\bigwedge_i \phi_i) \rightarrow \neg \psi$ then $(\bigwedge_i \neg \neg \phi_i) \rightarrow \neg \psi$. 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 $\sigma \upharpoonright n \in \mathcal{O}_1$; recalling that $\mathcal{O}_1 \subseteq \mathcal{O}_0$, we conclude that $\sigma \upharpoonright n \in \mathcal{O}_0$. By the choice of n , let $\tau \subseteq \sigma \upharpoonright n$ 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}_ρ is a subset of both \mathcal{O}_τ and \mathcal{O}_1 , so use \mathcal{O}_ρ 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 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 .

By the remark before this proposition, by increasing n and the n_σ 's as necessary, we can take them all to be equal. We will show \mathcal{V} covers \mathcal{O} by n . Let ρ be of length n . If $\rho \in \mathcal{O}$, then for some $\tau \subseteq \rho$ and $\mathcal{O}_\mathcal{U} \in \mathcal{U}$, $\mathcal{O}_\tau \subseteq (\mathcal{O} \cap \mathcal{O}_\mathcal{U})$. In particular, $\sigma \in \mathcal{O}_\mathcal{U}$. Since \mathcal{V} covers $\mathcal{O}_\mathcal{U}$ by n , there is a $\rho \subseteq \sigma$ and $\mathcal{O}_\mathcal{V} \in \mathcal{V}$ with $\mathcal{O}_\rho \subseteq (\mathcal{O}_\mathcal{U} \cap \mathcal{O}_\mathcal{V})$. Letting ν be the longer of ρ and τ , $\mathcal{O}_\nu \subseteq (\mathcal{O} \cap \mathcal{O}_\mathcal{V})$, 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 $\{\mathcal{O}' \mid \exists \mathcal{O}_\mathcal{U} \mathcal{O}' \subseteq \mathcal{O}_\mathcal{U} \text{ and } \exists \mathcal{O}_\mathcal{V} \mathcal{O}' \subseteq \mathcal{O}_\mathcal{V}\}$.

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}_ν are as desired. \square

The reason for this formal topology is so that we can take the Heyting-value 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 2. *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 $\{\langle \mathcal{O}_\sigma, \tau \rangle \mid \tau \subseteq \sigma, \mathcal{O}_\sigma \text{ a basic open set}\}$. We want to show that $\mathcal{O}_\emptyset \Vdash$ “for all k there is a unique σ of length k with $\sigma \in G$.” Fix a k . It is easy to see that if \mathcal{O}_σ is a basic open set with σ of length k then $\mathcal{O}_\sigma \Vdash$ “ σ is the unique member of G of length k .” Let \mathcal{U} be $\{\mathcal{O}_\sigma \mid \sigma \text{ has length } k \text{ and } \mathcal{O}_\sigma \text{ is a basic open set}\}$. 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, there are only finitely many such, and then go out far enough that all of them have died already. Now given a node τ of T of length n , $\tau \upharpoonright k$ and $\mathcal{O}_{\tau \upharpoonright k}$ are the desired witnesses. \square

So now we have seen how to kill any particular counter-example to FAN_Δ . How can we handle all possible counter-examples, to come up with a model of FAN_Δ ?

The idea is to iterate. More particularly, to allow for any finite iteration of such forcings.

Toward this end, working in the recursive realizability model, let \mathcal{P} be the partial order consisting of all finite sequences $\mu = \langle T_0, \dots, T_n \rangle$ of decidable, infinite trees with no infinite branches (with the natural order). To each such sequence we associate a base model. To the empty sequence, associate the recursive realizability model. For $\mu = \langle T_0, \dots, T_n \rangle$, the associated base model \mathcal{M}_μ is the formal topological model (as described above) for T_n built over the base model for $\langle T_0, \dots, T_{n-1} \rangle$. We will need the extension of the previous theorem to these base models.

Theorem 3. *Working within any of these base models, in \mathcal{M}_T the generic G is (identifiable with) an infinite branch through T .*

We would like to consider the induced Kripke model \mathcal{M} . At any node $\mu \in \mathcal{P}$, we build the set \mathcal{T}_α^μ of terms of rank at most α , inductively on \mathcal{M}_μ -ordinals α . In fact, for all $\nu \geq \mu$, the base model \mathcal{M}_ν extends \mathcal{M}_μ , so α is also an ordinal of \mathcal{M}_ν , and we can (and do) work with a stronger inductive hypothesis: for

all $\beta < \alpha$ and all $\nu \geq \mu$, at node ν , \mathcal{T}_β^ν is defined. We also assume we have a Kripke transition function from \mathcal{T}_β^μ to \mathcal{T}_β^ν , for which we polymorphically use the notation f . Given, \mathcal{T}_β^μ , let $\wp(\mathcal{T}_\beta^\mu)$ consist of all functions g with domain $\mathcal{P}^{\geq \mu}$ such that:

- $g(\nu) \subseteq \mathcal{T}_\beta^\nu$;
- if $s \in g(\nu)$ and $\xi > \nu$ then $f(s) \in g(\xi)$; and
- $g \upharpoonright \mathcal{P}^{\geq \nu} \in \mathcal{M}_\nu$.

Then $\mathcal{T}_\alpha^\mu = \bigcup_{\beta < \alpha} \wp(\mathcal{T}_\beta^\mu)$. The transition function f just restricts the domain.

Theorem 4. $\mathcal{M} \models IZF$.

Proof. Deferred to the full version of this paper. □

Theorem 5. $\mathcal{M} \models FAN_\Delta$.

Proof. If B is a decidable bar in the ground (realizability) model, then $\mathcal{M} \models$ “ B is not a bar,” because every node has an extension which has an infinite path avoiding B . It remains to show that no essentially new bars are introduced along the way. □

Theorem 6. $\mathcal{M} \models \neg FAN_c$.

Proof. Consider the following c -fan, due to Francois Dorais. 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”. 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 is the unique branch missing the induced c -set B , which is a bar in the realizability model, because K ’s characteristic function is not an infinite branch there. It can be shown that B remains a bar in \mathcal{M} , because generically K is not added by forcing to the model. □

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.
- [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] 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.
- [7] Robert Lubarsky and Hannes Diener, *Separating the Fan Theorem and its weakenings*, Journal of Symbolic Logic **79** (2014), pp. 792-813.
- [8] Jaap van Oosten, *Realizability: An Introduction to its Categorical Side*, Elsevier, 2008.