SEPARATING FRAGMENTS OF WLEM, LPO, AND MP

MATT HENDTLASS AND ROBERT LUBARSKY

Abstract. We separate many of the basic fragments of classical logic which are used in reverse constructive mathematics. A group of related Kripke and topological models is used to show that various fragments of the Weak Law of the Excluded Middle, the Limited Principle of Omniscience, and Markov's Principle, including Weak Markov's Principle, do not imply each other.

§1. Introduction. At the beginning of the twentieth century, Brouwer identified a number of constructively dubious principles, which Bishop later, in his 1967 monograph [2], termed *omniscience principles*. In the intervening years many new (weaker) principles have been introduced and models have been given showing the nonderivability of these principles in, for example, intuitionistic ZF set theory (IZF) [7].

Omniscience principles are commonly used to show the independence of more subject specific theorems: if a (classical) result constructively implies an omniscience principle, then it cannot be proved using constructive techniques. By separating different omniscience principles over IZF we make this task easier: if under the assumption of a classical result together with an omniscience principle we can derive a stronger omniscience principle, then we can still conclude that the classical theorem is nonconstructive. More generally, implications among these principles and theorems of mainstream mathematics have been studied for a long time. Often this is the motivation for introducing these principles (some references being provided with the principles below), and often this study is done for foundational reasons after the principles are already established (as, for instance, in [26], which also includes variants that we do not examine here). A more general reason to do this work is that the purpose of the study of logic is to clarify foundational principles; certainly knowing when various combinations of such do and do not imply others is a part of this.

In this paper we present many models, often related to each other, that separate a large number of the omniscience principles defined in terms of binary sequences and related principles. The genesis of this work was the first author's question to the second of whether Richman's LLPO_n hierarchy [23] could be separated, a question about results. Since then, much interest has shifted to technique: could an argument

Received April 2, 2015.

2010 *Mathematics Subject Classification.* 03F50, 03C90, 03E70, 03H99.

Key words and phrases. omniscience principles, Markov's Principle, Weak Markov's Principle, weak LEM, topological models, Kripke models, nonstandard models.

© 2016, Association for Symbolic Logic
DOI:10.1017/jsl.2016.38

for a simple case be extended to a more complicated case? How could a model giving a weak separation (that is, one in which a certain principle is not true but also not false) be re-configured into one giving a strong separation (in which the principle in question is false)? What has to be done to make DC (Dependent Choice) true? We hope that some of the results are of wider interest, such as separations involving a central, traditional axiom like Weak König's Lemma, or a new axiom like Weak Markov's Principle. At the same time, the exposition has a strong orientation to technique, often providing several proofs of the same theorem, in the spirit of a saying we've heard and like, that it's better to prove the same thing in five different ways than to prove five different things the same way.

Turning now to the principles we will be studying, constructive mathematics is often crudely characterized as mathematics without the law of excluded middle

- **LEM**: For any proposition A, either A is true or A is false,

and IZF is essentially the result of expunging LEM from ZF. Our goal is always to end up with models satisfying IZF + Dependent Choice,

- **DC**: If a is a set, $x_0 \in a$, S is a subset of $a \times a$, and for each $x \in a$ there exists $y \in a$ such that $(x, y) \in S$, then there exists a sequence $(x_n)_{n\in\mathbf{N}}$ in X such that $x_0 = a$ and $(x_n, x_{n+1}) \in S$ for each n,

which is more than sufficient to formalise Bishop's constructive mathematics [2]. Since there is a double-negative translation of ZF into IZF, we lose nothing by adopting ZFC as our meta-theory, and since many of the models we present make use of ultrafilters this classical meta-theory is essential.

Many of the commonly occurring omniscience principles can be stated in terms of binary sequences. We denote the space of infinite binary sequences, Cantor space, by 2^ω and use decorations of α, β to represent elements of 2^ω. We reserve n, m to represent natural numbers, and use i, j, k to range over bounded sets of natural numbers.

Brouwer introduced the following three basic omniscience principles.

- **The Limited Principle of Omniscience (LPO)**: For any binary sequence α, either $\alpha(n) = 0$ for all n or there exists n such that $\alpha(n) = 1$.
- **The Weak Limited Principle of Omniscience (WLPO)**: For any binary sequence α, either $\alpha(n) = 0$ for all n or it is not the case that $\alpha(n) = 0$ for all n.
- **The Lesser Limited Principle of Omniscience (LLPO)**: For any binary sequence α with at most one nonzero term, either $\alpha(n) = 0$ for all even n or $\alpha(n) = 0$ for all odd n.

If we add countable choice to our system, then LLPO is equivalent to weak König's lemma:[1]

- **WKL**: Every infinite decidable binary tree has an infinite branch.

In [23] Richman defined a hierarchy of principles LLPO_ν ($\nu \in \omega + 1, \nu \geqslant 2$) related to LLPO:

[1] Ishihara has shown that weak König's lemma is equivalent to LLPO plus Π^0_1-$\mathrm{AC}_{\omega,2}$ over a weak constructive base system [11].

- **LLPO$_\nu$**: Let $(P_i)_{i<\nu}$ be a decidable partition of ω into blocks of size ω, and let α be a binary sequence with at most one nonzero term. Then there exists $k < \nu$ such that $\alpha(m) = 0$ for all $m \in P_k$.

By glueing together blocks from partitions, we see that LLPO$_\nu$ implies LLPO$_{\nu'}$ whenever $\nu < \nu'$. If $\nu = n \in \omega$, then without loss of generality we restrict our consideration to the natural n-partition of ω: $P_i = \{mn + i : m \in \omega\}$.

LLPO is just the restriction of De Morgan's law—for any propositions A, B, if $\neg(A \wedge B)$, then $\neg A$ or $\neg B$—to Σ^0_1-formulas. De Morgan's law is equivalent to LEM for negative formulas: the Weak Law of Excluded Middle

- **WLEM**: for any proposition A, either $\neg A$ or $\neg\neg A$.

The LLPO$_n$ hierarchy can be formed in a similar manner as the restrictions of a family of weakenings of WLEM:

- **WLEM$_n$**: $\quad \neg \bigvee_{i,j<n, i\neq j} A_i \wedge A_j \longrightarrow \bigvee_{i<n} \neg A_i$;

LLPO$_n$ is the restriction of WLEM$_n$ to Σ^0_1-formulas. We can define a similar principle WLEM$_\omega$ by quantifying over the natural numbers.

Another principle commonly considered in constructive reverse mathematics is Markov's Principle:

- **MP**: If it is impossible for all terms of α to be zero, then there exists an n such that $\alpha(n) = 1$.

Markov's Principle represents unbounded search and is accepted by some practitioners of constructive mathematics, notably those of the Russian school of recursive mathematics. There are two standard weakenings of Markov's Principle, Weak and Disjunctive MP.

- **WMP**:

$$\forall \alpha\, [\forall \beta\, (\neg\neg \exists n\, (\beta(n) = 1) \vee \neg\neg \exists n\, (\alpha(n) = 1 \wedge \beta(n) = 0)) \to \exists n\, \alpha(n) = 1].$$

- **MP$^\vee$**: If α has at most one nonzero term and it is impossible for all terms of α to be zero, then either all even terms are zero or all odd terms are zero.

Often MP and its variants are stated in terms of real numbers, instead of binary sequences. For instance, WMP in this form is

$$\forall a \in \mathbb{R}\, (\forall x \in \mathbb{R}\, (\neg\neg(x < a) \vee \neg\neg(0 < x)) \to a > 0).$$

The binary and real formulations are equivalent with a small amount of choice (weak countable choice suffices). The topological model over $\mathbb{R}$, with the standard topology, shows that some choice is necessary, since MP for binary sequences holds there but MP for reals does not.

Markov's Principle is equivalent to the conjunction of WMP and MP$^\vee$ [10, 17]. Finally, we shall also consider the generalisations MP$^\vee_n$ and MP$^\vee_\omega$ of MP$^\vee$ corresponding to LLPO$_n$ and LLPO$_\omega$.

It should be mentioned, even if only briefly, that many of these principles have been studied from a classical, computational point of view. For instance, the unsolvability of the halting problem shows the computational failure of a uniform version of LEM, and WKL, while studied here, is even more central in Reverse Mathematics (see for instance [25]). There have also been such studies from the point of view of Weihrauch reducibility, including LLPO and its hierarchy, our backbone [3,4,18,32].

1.1. Summary. The following diagram summarizes the known (to us) relationships, over IZF, among the principles we have introduced. The two-headed arrows are bi-implications, and the single arrows are strict, as either proven in this paper or already known.

We summarize briefly, often with well-known constructions, or give references to the reasons for the nonreversal of the black arrows, working from left-to-right in the diagram. Please note that we are not claiming that the proofs or arguments given are the best, or the simplest, or the first, or the canonical ones in some way. We merely want to demonstrate all these nonimplications, preferably in a not too difficult way. An example of a model of WLEM + ¬LEM is any (nontrivial) Kripke model over a linear order with no last element. The topological model over $\mathbb{R}$ satisfies LPO + ¬LEM. In fact, any (nontrivial) Kripke model over a linear order with no last element and which contains the standard natural numbers (such as the full model over the partial order) satisfies both WLEM and LPO and falsifies LEM. A quite different technique to separate LPO from LEM is realizability [21]. LLPO and WKL are separated in Lifschitz realizability [5]. Kleene's number realizability K_1 separates MP and LLPO_ω [23], giving a whole column of nonreversals; we give a very different model of the same separation below (Theorem 5.8).

The parallel black and red arrows are because the independence in question has already been shown, but over theories of a different nature from IZF. Kohlenbach [1, 12–14] achieves all of the nonreversals so indicated, and a lot more, using various realizabilities. The difference in all of these cases from our context is that the models are all of higher-type arithmetic, the higher types being function spaces. A clear difference from IZF is that there are no power sets and no transfinite iterations.

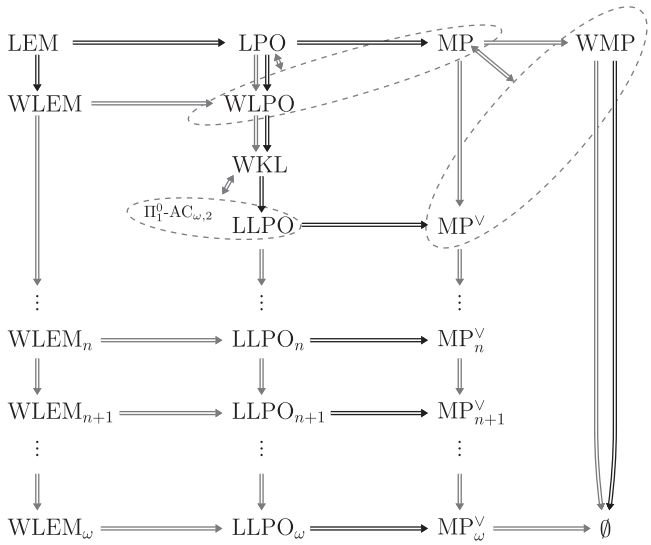

FIGURE 1. Fragments of LEM with all implications, over IZF, indicated. Red arrows correspond to separations given in this paper, black and blue arrows to those previously known.

A more subtle distinction is the role of Comprehension. In the setting of functionals, a subset is given by a 0-1-valued function. So Comprehension for a property gives Excluded Middle for that property. Hence the amount of Comprehension present in these models is of necessity limited (in the cases at hand, typically to negated formulas or $\exists$-free formulas). In contrast, IZF has full Comprehension, although its consequences are less weighty. Besides all that, the methods in this paper are so completely different from the earlier realizabilities that we'd want them to get an airing anyway.

For the blue arrows, the equivalence of LPO with WLPO + MP is trivial. It is easy to see that LLPO with Countable Choice is enough to yield WKL; Ishihara [11] has identified exactly the amount of choice needed and shown that to follow from WKL, effectively decomposing WKL into a choice and an omniscience principle. The equivalence of MP with WMP + $\text{MP}^\vee$ is in [17], notwithstanding the facts that all of those principles have different names there, and all are presented in their real number (as opposed to their binary) versions. With sufficient Choice, those versions are of course equivalent, but we would like to see the equivalence of those principles, as stated here (i.e., in terms of binary sequences), without using any Choice. This was done in [10], and with the same names as here to boot, albeit with $\text{MP}^\vee$ phrased differently. Because of that difference, and because the proof there is rather terse, we give the details here. So suppose that not all terms of α are 0. By WMP, we will be done if we can show α is pseudopositive (i.e., WMP's hypothesis). So let β be arbitrary. Define γ as follows. While generating the values of α and β, let γ continue taking the value 0, until the least n is reached (if ever) with either $\beta(n)$ or $\alpha(n)$ being 1. If $\beta(n) = 1$ then $\gamma(2n) = 1$, else $\gamma(2n + 1) = 1$, after which γ always has value 0. We would like to apply $\text{MP}^\vee$ to γ. By construction, γ has at most one nonzero term. If γ were always 0, then so would be α (because if $\alpha(n) = 1$ then γ takes on the value 1 at or before $2n + 1$). That contradicts the assumption on α, so γ is not always 0. By $\text{MP}^\vee$, either γ's even or odd terms are 0. Suppose the former. Working toward a contradiction, suppose there is no n with $\alpha(n) = 1$ and $\beta(n) = 0$. If there were an n with $\beta(n) = 1$ then for some $k \leq n$ we have $\gamma(2k) = 1$, which can't happen. So β is always 0. Hence there is no n with $\alpha(n) = 1$. But this contradicts the choice of α. So the hypothesis of WMP is satisfied in this case. Now suppose all of γ's odd terms are 0. Working toward a contradiction, assume β is always 0. But then all of γ's even terms would be 0 too, which we already saw cannot happen. So the hypothesis of WMP is satisfied in this case. This suffices.

For the nonreversal of the red arrows over IZF + DC, we refer the reader to Figure 2. By Theorem 4.1, none of the arrows from the WLEM column to the LLPO column reverse. By Theorems 5.1 and 5.6, WMP and $\text{MP}^\vee_\omega$ are unprovable; Theorem 5.1 also shows that $\text{MP}^\vee$ does not imply MP and that WLPO does not imply LPO. Theorem 5.7 shows that $\text{MP}^\vee_\omega$ does not imply $\exists n\ MP^\vee_n$, and Theorem 6.1 that WKL does not imply WLPO. That WMP does not show MP follows from Theorem 5.6, and all of the other arrows are shown not to reverse by Theorem 5.2.

By a nonreversal of the implication $A \to B$ over IZF + DC, we mean a model of IZF + DC in which B is true and A is false. Ultimately that is what is shown in each case. However, we also present many partial nonreversals along with way, just for the interest of the models themselves. For instance, in some models DC is false, and in others we just don't know. Also, some of the separations are *weak* as opposed to

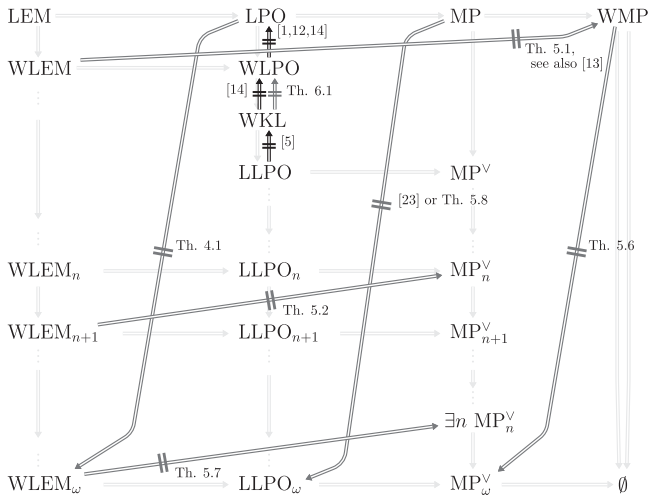

FIGURE 2. The nonderivabilities that give all the nonreversals among the fragments of LEM we study here.

strong separations. In a weak separation, A is not true, but also not false; rather, $\neg\neg A$ holds. In a strong separation, A is false, meaning $\neg A$ holds.

§2. Topological models. Many of the models we present are topological models. Although long known [9], they are not so widely understood, and so we summarize here the basics, as well as some particular results we will need.

Topological models are Heyting valued models where the complete Heyting algebra is the lattice of opens T_X of a topological space X. Meet and join in T_X are given by intersection and union, respectively, while the psuedo-complement $\to$ is defined by

$$U \to V \;\equiv\; (-U \cup V)^\circ,$$

where $-U$ denotes the complement of U in X and S° denotes the interior of S. The *full topological model* over X consists of the class of *names* or *terms*, defined inductively by

$$V_\alpha(X) = \mathcal{P}\left(\bigcup\{V_\nu(X) \times T_X : \nu \in \alpha\}\right),$$
$$V(X) = \bigcup_{\alpha\in\mathrm{ORD}} V_\alpha(X).$$

Given $\sigma \in V_\alpha(X)$, the meaning of $\langle\tau, \mathcal{O}\rangle \in \sigma$ is that $\mathcal{O}$ is the degree of truth, or truth-value, of τ being in σ. (Of course, the ultimate value of $\tau \in \sigma$ might be greater than $\mathcal{O}$, depending on what else is in σ.) The idea of the full model is to throw in absolutely everything you can. We will have occasion to look at submodels of the full model. An embedding $\check{}$ of the ground model V into $V(X)$ is defined inductively by

$$\check{a} = \{<\check{b}, X>: b \in a\}.$$

The truth value of any proposition A, with parameters from $V(X)$, is an open subset of X and is denoted by $[\![A]\!]$. To say that a proposition A is *true*, or *satisfied*, in a topological model M_X over X means $[\![A]\!] = X$, otherwise A is said to fail in M_X. Being false in M_X is a stronger property: A is said to be *false* in M_X if M_X satisfies $\neg A$, or equivalently $[\![A]\!] = \emptyset$. We freely switch between truth value notation $[\![\cdot]\!]$ for topological models and forcing notation: a point $x \in X$ *forces* a formula A, written $x \Vdash A$, if and only if $x \in [\![A]\!]$, and, for an open subset U of X, $U \Vdash A$ if and only if $U \subset [\![A]\!]$.

A particularly important object, the *generic*, in a topological model M_X is described by the name

$$G = \{< \check{U}, U >: U \in T_X\}.$$

The generic gives a new element of the topological space over which we are forcing, and is characterised by the equation

$$U \Vdash \check{U} \in G$$

for all $U \in T_X$.

See [9] for an introduction to topological models of a constructive formulation of second order arithmetic. Topological models preserve IZF; that is, IZF proves that the full topological model satisfies the axioms of IZF.

We will have need of two simple and well-known observations, summarized in the next lemma.

LEMMA 2.1. *Let X be a topological space, x a point of X, and U an open subset of X.*

1. *If $\{x\}$ is open, then $x \Vdash LEM$.*
2. *$U \Vdash \neg\neg A$ if and only if $\{x \in X : x \Vdash A\}$ is dense in U.*

PROOF. The second part of the lemma is just an unpacking of definitions. To see (1), note that $x \Vdash \neg A$ if and only if there is a neighbourhood U of x such that no point of U forces A which occurs precisely when $x \nVdash A$. ⊣

The next result is used to show that most of the topological models we present satisfy DC. A topological space X is *dim-zero-dimensional*, or of *covering dimension zero*, if for every open cover $(\mathcal{O}_i)_{i\in I}$ of X there exists an open partition $(\mathcal{O}'_j)_{j\in J}$ of X such that for each $j \in J$, $\mathcal{O}'_j \subset \mathcal{O}_i$ for some $i \in I$. We call $(\mathcal{O}'_j)_{j\in J}$ a *partition refinement* of $(\mathcal{O}_i)_{i\in I}$. We can always replace the partition refinement $(\mathcal{O}'_j)_{j\in J}$ of $(\mathcal{O}_i)_{i\in I}$ by a partition refinement indexed by $|T_X|$ simply by bulking out $(\mathcal{O}'_j)_{j\in J}$ with copies of the empty set.

PROPOSITION 2.2. *If X is dim-zero-dimensional, then $M_X \vDash DC$.*

PROOF. Suppose, without loss of generality, that

$$X \Vdash \forall s \in S\ \exists t \in S\ (\varphi(s,t)) \wedge s_0 \in S.$$

Then, whenever $X \Vdash s \in S$, $\{[\![\varphi(s,t)]\!] \cap \mathcal{O} : \langle t, \mathcal{O}\rangle \in S\}$ is an open cover of X. We inductively construct $\mathcal{O}_\sigma, t_\sigma$ $(\sigma \in S^{<\omega})$ such that for each $\sigma \in S^{<\omega}$ and each $n \in \omega$

(1) $\{\mathcal{O}_\sigma : |\sigma| = n\}$ is an open partition of X;
(2) $\mathcal{O}_{\sigma^\frown\langle t,\mathcal{O}\rangle} \subseteq [\![\varphi(t_\sigma, t_{\sigma^\frown\alpha})]\!]$.

We begin by setting $\mathcal{O}_{\langle\rangle} = X$ and $t_{\langle\rangle} = s_0$. Now suppose we have constructed $\mathcal{O}_\sigma$ and t_σ for all $\sigma \in S^n$ and fix a string σ in S^n. Let $(\mathcal{O}'_{\sigma,\langle t,\mathcal{O}\rangle})_{\langle t,\mathcal{O}\rangle\in S}$ be a partition refinement of $\{[\![\varphi(t_\sigma, t)]\!] \cap \mathcal{O} : \langle t, \mathcal{O}\rangle \in S\}$. For each $\langle t, \mathcal{O}\rangle \in S$ we set $\mathcal{O}_{\sigma^\frown\langle t,\mathcal{O}\rangle} = \mathcal{O}_\sigma \cap \mathcal{O}'_{\langle t,\mathcal{O}\rangle}$ and $t_{\sigma^\frown\langle t,\mathcal{O}\rangle} = t$. This completes the inductive construction.

Define a name

$$f = \{< pair(\check{n}, t_\sigma), \mathcal{O}_\sigma >: n \in \omega, \sigma \in S^n\},$$

where $pair(a, b)$ is a term for what is intended to be the ordered pair of a and b: $\{\{\langle a, X\rangle\}, X\rangle, \langle\{\langle a, X\rangle, \langle b, X\rangle\}, X\rangle\}$. By (1), X forces that f is a sequence in S, and by (2)

$$X \Vdash \forall n \in \omega\ \varphi(f(n), f(n+1)). \qquad \dashv$$

§3. Kripke models. Many of the remaining models are Kripke models. The constructions used bear a family resemblance. In an attempt to economize, we present here a general framework, which completely suffices for some of the models. Then we present a rather more particular construction, which suffices for most of the rest of the models.

Let $\mathcal{P}$ be a partial order. Suppose that to each node $\sigma \in \mathcal{P}$ there is an associated model M_σ of ZFC. Also suppose that whenever $\sigma \leq \tau$ there is an elementary embedding $f_{\sigma\tau} : M_\sigma \to M_\tau$. (Whenever $\sigma = \tau, f$ is the identity.) Moreover, the system of $f_{\sigma\tau}$'s coheres: $f_{\tau\rho} \circ f_{\sigma\tau} = f_{\sigma\rho}$. In the following, we will usually drop the subscripts to $f_{\sigma\tau}$ and let f act as a polymorphic embedding. We also assume that the system $\mathcal{P}^{\geq\sigma}$, with the associated assignment of M_τ to τ whenever $\tau \geq \sigma$, is in (well, more accurately, definable over) M_σ.

With this backdrop, we can now define the *full model* over this system. The objects at node σ are defined within M_σ inductively through the ordinals α of M_σ. Assume inductively that we have a set K^σ_β of objects of rank β ($\beta < \alpha$) at σ, along with transition functions $k : K^\sigma_\beta \to K^\tau_{f(\beta)}$. (It will be clear that the action of k does not depend on the choice of β, and so there is no need to distinguish various k's dependent on the choice of β. In some sense k does depend on σ, but again we will usually drop that from the notation and allow k to be polymorphic.) By elementarity, we also have the corresponding K^τ_β for any $\tau > \sigma$ and $\beta < f(\alpha)$. Then an object of K^σ_α will be a function g such that:

- $\mathrm{dom}(g) = \mathcal{P}^{\geq\sigma}$,
- $g \restriction \mathcal{P}^{\geq\tau} \in M_\tau$,
- $g(\tau) \subseteq \bigcup_{\beta<f(\alpha)} K^\tau_\beta$, and
- if $h \in g(\tau)$ and $\tau < \rho$ then $k_{\tau\rho}(h) \in g(\rho)$.

Let $k_{\sigma\tau}(g)$ be $g \restriction \mathcal{P}^{\geq\tau}$. This allows the inductive construction to continue through the ordinals.

Let the objects at node σ be the functions g that are in K^σ_α for any ordinal α from M_σ. Say $\sigma \models g \in h$ iff $g \in h(\sigma)$, and $\sigma \models g = h$ iff $g = h$. This gives a structure for the language of set theory at every node. Let the action of the transition function k be domain restriction, as with any of the partial k's from the construction.

THEOREM 3.1. *The full model satisfies IZF.*

PROOF. Empty Set: Let g be the function (with appropriate domain) that always returns the empty set.

Infinity: Within any M_σ, one can define inductively on $n \in \omega$ an element n_σ of the full model. To start, 0_σ is the set satisfying the Empty Set Axiom from above. Given m_σ for $m < n$, let $n_\sigma(\tau)$ be $\{k_{\sigma\tau}(m_\sigma) \mid m < n\}$. Notice that $k_{\sigma\tau}(m_\sigma) = m_\tau$. Let ω_σ be such that $\omega_\sigma(\tau) = \{m_\tau \mid m \in \omega^{M_\tau}\}$. Then ω_σ suffices.

Pairing: Given g and h at node σ, the desired pair is the function which, at node τ, yields the set $\{k(g), k(h)\}$.

Union: Given g at node σ, let h be such that $h(\tau) = \bigcup\{\hat{g}(\tau) \mid \hat{g} \in g(\tau)\}$.

Extensionality: Because equality is taken to be actual equality, it is easy to see that equal sets have the same members. On the other hand, if two sets are forced at a node to have the same members, then they are literally equal as functions from that node on. By Extensionality in the ambient universe, they are then equal.

$\in$-Induction: We must take some care here, because we did not assume that any M_σ is well-founded. So suppose we're dealing with a counter-example to $\in$-Induction: $\sigma \models \forall x(\forall y \in x\ \phi(y) \to \phi(x))$, yet $\sigma \not\models \forall x\ \phi(x)$. Then for some $\tau \geq \sigma$ there is a $g \in K^\tau$, $\tau \not\models \phi(g)$. By the inductive hypothesis, $\tau \not\models \forall y \in g\ \phi(y)$. Hence there is a $\rho \geq \tau$ and a term $h \in K^\rho$ with $\rho \models h \in g$, yet $\rho \not\models \phi(h)$.

Returning now to τ, the situation just described is a statement about M_τ : $M_\tau \models$ "There is a partial order $\mathcal{Q}$ with bottom element $\perp$, and a system of models with elementary embeddings, starting with $M_\perp = V$, such that the induced full model does not satisfy $\in$-Induction; moreover, there is such an example with a specific counter-example g already in $K^\perp$." Because M_τ is a model of ZFC, there is such a model in M_τ with such a g of least possible rank, say α, among all such models. Within that model, by the considerations above, there is a node $\rho \geq \perp$ and an $h \in K^\rho$ with $\rho \models h \in g$, yet $\rho \not\models \phi(h)$. So M_ρ models that h is also a specific counter-example to $\in$-Induction, and of rank strictly smaller than that of g. But by elementarity, within M_ρ, it is $f(\alpha)$ that is the least rank of such a counter-example. This is a contradiction. So there is no counter-example to $\in$-Induction.

Power Set: Let $g \in K^\sigma$. At any node $\tau \geq \sigma$, if $\tau \models h \subseteq g$, then h is (forced at τ to be) a function with domain $\mathcal{P}^{\geq\tau}$ such that $h(\rho) \subseteq g(\rho)$ (of course, also respecting the other condition on being in K^ρ). In fact, any such h at $\tau \geq \sigma$ can be taken to be an object at σ, by letting $h(\rho)$ be the empty set whenever ρ does not extend τ. Noticing there are only set many such h's, the set of all of them is the power set of g.

Separation: Given ϕ and g at σ, let $h \in K^\sigma$ be $h(\tau) = \{g' \in g(\tau) \mid \tau \models \phi(g')\}$.

Replacement \ Collection \ Reflection: While those three axiom schemas are equivalent classically, they are apparently of strictly increasing strength constructively. That is, Friedman–Scedrov [8] showed that Replacement does not imply Collection, and while a proof is wanting, presumably Collection does not imply Reflection. (The converse implications are all soft proofs.) Typically IZF is taken to be the theory with Collection. In fact, in our setting it is easy to see that Reflection holds true. Working in $V = M_\sigma$, by Reflection there, let $V_\alpha \prec_{\Sigma_n} V$. Cutting off the construction of the model at α, which is K^σ_α at σ and $K^\tau_{f(\alpha)}$ at any other τ, yields a Σ_n substructure of the entire model, in the sense that any property of K^σ which is Σ_n expressible in the ambient M_σ reflects. $\dashv$

For the second construction, start with a partial order $\mathcal{P}$ which is a tree of height ω. For notational convenience, we may as well take $\mathcal{P}$ to be a set of strings, closed under truncation (so $\langle\rangle$ is the root) and extension in $\mathcal{P}$ as a partial order is the same as string extension. Assume that, as before, there is an assignment of a ZFC-model M_σ to each node σ, with (polymorphic) elementary embeddings f. For convenience, we take $M_{\langle\rangle}$ to be V. This time, the tree $\mathcal{P}$ and the assignment are moreover assumed to be *uniform*, as follows. To define what it is for $\mathcal{P}$ to be uniform, let the extensions of σ, ext_σ, be the strings ρ of length 1 such that $\sigma^\frown\rho \in \mathcal{P}$ (which are exactly what yield σ's children). Then for $\sigma \leq \tau$, $ext_\tau = f(ext_\sigma)$. For instance, $\mathcal{P}$ could be the full binary tree, and then each ext_σ would be 2, or $\{0,1\}$. Another example more akin to Baire space would be $ext_\sigma = \omega^{M_\sigma}$. For the assignment of M_σ to each σ to be uniform means that, for $\sigma \leq \tau$, the definition of the assignment $\rho \mapsto M_{\tau^\frown\rho}$ for $\rho \in ext_\tau$ is the same definition as for $\rho \mapsto M_{\sigma^\frown\rho}$ for $\rho \in ext_\sigma$ (where, of course, any parameters are translated via f). For instance, if M_0 is an ultrapower of $M_{\langle\rangle}$ via the ultrafilter $\mathcal{U}$, then $M_{\tau^\frown 0}$ is an ultrapower of M_τ via $f(\mathcal{U})$. Also, if M_n (n a natural number) is the n-fold iteration of the ultrapower construction applied to $M_{\langle\rangle}$ based on the ultrafilter $\mathcal{U}$, then the same holds for $M_{\tau^\frown n}$ relative to M_τ and $f(\mathcal{U})$, even for n nonstandard, where the non-standard-length iteration is as defined in M_τ.

With this set-up, we can define the *immediate settling model* over this system. The idea is that a set can change arbitrarily from a node to an immediate successor, just as above, but then it can't change anymore, instead needing to settle down. Of course, the notion of not changing is mediated by the embeddings f.

More precisely, the universe K^σ at node σ will consist of functions g such that:

- the domain of g, $\mathrm{dom}(g)$, consists of $\langle\rangle$ and all of $\langle\rangle$'s immediate successors, the strings in $\mathcal{P}$ of length 1 (for which we use the variable τ),
- $g \in M_\sigma$ and $g(\tau) \in M_{\sigma^\frown\tau}$,
- inductively, $g(\langle\rangle) \subseteq K^\sigma$ and $g(\tau) \subseteq K^{\sigma^\frown\tau}$,
- if $h \in g(\langle\rangle)$ then $k(h) \in g(\tau)$, where k is the inductively defined transition function from K^σ to $K^{\sigma^\frown\tau}$, and
- if $h \in g(\tau)$ then $k(h) \in f(g(\tau))$, where k is the transition function from $K^{\sigma^\frown\tau}$ to $K^{\sigma^\frown\tau^\frown\rho}$ and f is the elementary embedding from $M_{\sigma^\frown\tau}$ to $M_{\sigma^\frown\tau^\frown\rho}$.

To define the transition function k from K^σ to some extension of σ, it suffices to define those k's from K^σ to an immediate extension $K^{\sigma^\frown\tau}$, since transition functions to nonimmediate extensions can be taken to be compositions of functions to immediate extensions. So for $g \in K^\sigma$ and τ of length 1, let $k(g)(\langle\rangle)$ be $g(\tau)$, and for ρ of length 1 let $k(g)(\rho)$ be $f(g(\tau))$, where f is the elementary embedding from $M_{\sigma^\frown\tau}$ to $M_{\sigma^\frown\tau^\frown\rho}$. We leave it to the reader to show that $k(g) \in K^\tau$, as well as the following.

LEMMA 3.2. *Suppose that $g \in K^\sigma$ is such that $g(\tau) = f(g(\langle\rangle))$. (This is the case, for instance, when g is in the image of k.) Then $k(g) = f(g)$.*

Let $\sigma \models g \in h$ iff $g \in h(\langle\rangle)$, and $\sigma \models g = h$ iff $g = h$.

THEOREM 3.3. *The immediate settling model satisfies IZF.*

PROOF. Empty Set: Let g be the function (with appropriate domain) that always returns the empty set.

Infinity: Within $V = M_{\langle\rangle}$, one can define inductively on $n \in \omega$ an element $n_{\langle\rangle}$ of this model, which will also be the n^{th} successor of 0 at any other node too. To start,

$0_{\langle\rangle}$ is the set satisfying the Empty Set Axiom from above. Given $m_{\langle\rangle}$ for $m < n$, let $n_{\langle\rangle}(\langle\rangle) = n_{\langle\rangle}(\tau)$ be $\{m_{\langle\rangle} \mid m < n\}$. Let $\omega_{\langle\rangle}$ be such that $\omega_{\langle\rangle}(\langle\rangle) = \{n_{\langle\rangle} \mid n \in \omega\}$, and $\omega_{\langle\rangle}(\tau) = f(\omega_{\langle\rangle}(\langle\rangle))$. Then $\omega_{\langle\rangle}$ suffices at the node $\langle\rangle$.

Pairing: Given g and h at node σ, the desired pair is the function which, at node τ, yields the set $\{k(g), k(h)\}$. To show that this set is in K^σ, most of the properties are simple; for the last, use the lemma before this theorem.

Union: Given g at node σ, let h be such that $h(\tau) = \bigcup\{\hat{g}(\langle\rangle) \mid \hat{g} \in g(\tau)\}$.

Extensionality: Because equality is taken to be actual equality, it is easy to see that equal sets have the same members. On the other hand, if two sets are forced at a node to have the same members, then they are literally equal as functions. By Extensionality in the ambient universe, they are then equal.

$\in$-Induction: Because of the uniformities and simplicities in the current set-up, this proof can proceed by consideration of only this model, as opposed to last time. So suppose $\in$-Induction fails: $\sigma \models \forall x(\forall y \in x\ \phi(y) \to \phi(x))$, yet $\sigma \not\models \forall x\ \phi(x)$. Therefore, for some $\tau \geq \sigma$ there is a $g \in K^\tau$, $\tau \not\models \phi(g)$.

First we're going to consider the simple case when there is such a τ which strictly extends σ. The advantage here is that, at τ, "ϕ" is really $k(\phi)$, meaning that the parameters are actually of the form $k(g)$, for g an original parameter of ϕ at σ. That means, by the earlier lemma, that with reference to any extension ρ of τ, $k(\phi) = f(\phi)$. Working in M_τ, let g be a set of least rank α such that $\tau \not\models \phi(g)$. Anything forced at τ to be a member of g has smaller rank, and so is forced at τ to satisfy ϕ. Also, anything forced at any $\rho > \tau$ to be a member of $k(g)$ has rank less than $f(\alpha)$. By elementarity, in M_ρ, if h has rank less than $f(\alpha)$ then $\rho \models f(\phi)(h)$. As remarked above, $f(\phi) = k(\phi)$, so if $\rho \models h \in k(g)$, then $\rho \models \phi(h)$. So $\tau \models \forall y \in g\ \phi(y)$, hence by the inductive hypothesis $\tau \models \phi(g)$, contrary to the choice of g.

So there is no such τ; hence, if $\tau > \sigma$, then $\tau \models \forall h\ \phi(h)$. Let g be of minimal rank such that $\sigma \not\models \phi(g)$. If $\sigma \models h \in g$ then h has smaller rank than g and so $\sigma \models \phi(g)$. If $\tau > \sigma$ and $h \in K^\tau$ then as already observed $\tau \models \phi(h)$. So $\sigma \models \forall y \in g\ \phi(y)$, hence by the inductive hypothesis $\sigma \models \phi(g)$, contrary to the choice of g.

Power Set: This is likely the most interesting axiom to check. After all, any alleged power set must settle down after one step through $\mathcal{P}$, yet a possible subset might show up later which then has its own step in which to settle down. Here we present a general proof; it might help the reader to look Model 5, where we work through the most basic example of the Power Set Axiom in detail.

Let $g \in K^\sigma$. Let $\wp(g)(\langle\rangle)$ consist of all functions $h \in K^\sigma$ such that $h(\langle\rangle) \subseteq g(\langle\rangle)$ and $h(\tau) \subseteq g(\tau)$. Also, $\wp(g)(\tau)$ consists of all $h \in K^\tau$ such that $h(\langle\rangle) \subseteq g(\tau)$ and $h(\rho) \subseteq f(g(\tau))$. Note there are only set-many such functions. It is an unenlightening technical exercise to show that $\wp(g) \in K^\sigma$. Then $\wp(g)$ is as desired: if $\sigma \models h \subseteq g$, then $h \in \wp(g)(\langle\rangle)$ by construction, and similarly for one-step extensions $\sigma^\frown\tau$; for longer $\rho > \sigma$, it essentially comes down to elementarity.

Separation: Given ϕ and g at σ, let $h \in K^\sigma$ be $h(\langle\rangle) = \{g' \in g(\langle\rangle) \mid \langle\rangle \models \phi(g')\}$ and $h(\tau) = \{g' \in g(\tau) \mid \tau \models \phi(g')\}$.

Replacement \ Collection \ Reflection: As in the previous theorem, it is easy to see that Reflection holds true. Working in $V = M_\sigma$, by Reflection there, let $V_\alpha \prec_{\Sigma_n} V$.

Cut off the construction of the model at α. That is, let $K(\langle\rangle)$ be K^{σ}_{α} and $K(\tau)$ be $K^{\sigma^\frown\tau}_{f(\alpha)}$. This yields a Σ_n substructure of the entire model, in the sense that any property of K^{σ} which is Σ_n expressible in the ambient M_σ reflects. ⊣

§4. LPO and WLEM.

THEOREM 4.1. *Over IZF + DC, LPO does not imply WLEM$_\omega$.*

PROOF. We give several constructions.

Model 1: This first model has the benefit of being well-known. It has the drawback of falsifying DC.

Consider the full topological model $M_{\mathbb{R}^2}$ over $\mathbb{R}^2$. If x forces that σ is an infinite binary sequence, then some connected neighbourhood $\mathcal{O}$ of x forces the same. Then for each n, $\mathcal{O}$ is the disjoint union of $[\![\sigma(n) = 0]\!] \cap \mathcal{O}$ and $[\![\sigma(n) = 1]\!] \cap \mathcal{O}$, and so either $\mathcal{O} \subset [\![\sigma(n) = 0]\!]$ or $\mathcal{O} \subset [\![\sigma(n) = 1]\!]$. Thus $\mathcal{O}$ forces σ to be a ground model sequence and hence, applying LPO at the meta-level, either $x \Vdash \forall n\ \sigma(n) = 0$ or $x \Vdash \exists n\ \sigma(n) = 1$.

We show that $(0,0) \nVdash \mathrm{WLEM}_\omega$. Since there is nothing special about $(0,0)$, it follows that $M_{\mathbb{R}^2} \vDash \neg\, \mathrm{WLEM}_\omega$. For each $n \in \omega$ set

$$S_n = \left\{ \mathbf{x} \in \mathbb{R}^2 : \mathbf{x} \neq \mathbf{0} \wedge \frac{1}{\pi} \cos^{-1}\left(\frac{\mathbf{x} \cdot (0,1)}{||\mathbf{x}||} \right) \in \left(1 - 1/2^n, 1 - 1/2^{n+1}\right) \right\};$$

that is, take a countable sequence of disjoint wedges in the upper half-plane with vertex the origin. Let A_n be the statement $\check{0} \in \{< \check{0}, S_n >\}$; viewing S_n as a truth value, A_n means "S_n is true." Since the S_n are mutually disjoint $[\![\exists n, m\ (n \neq m \wedge A_n \wedge A_m)]\!] = \emptyset$, so $\mathbb{R}^2 \Vdash \neg\exists n, m\ (n \neq m \wedge A_n \wedge A_m)$. However, since $(0,0)$ is in the closure of each S_n, $(0,0)$ cannot force any A_n to be false, and hence does not force WLEM_ω.

Model 2: The most basic models separating two principles are models in which the weaker principle is true while the stronger principle is not satisfied, as opposed to being false—we call these *weak separations*. This model is such a weak separation. It has the benefit of satisfying DC.

Let $\mathbb{N}^+$ be $\mathbb{N} \cup \{*\}$, with the discrete topology on $\mathbb{N}$ and the only neighborhood of $\{^*\}$ being the whole space. Then the full topological model of $\mathbb{N}^+$ is the same thing as the full Kripke model on the system with the partial order having * as the bottom node and its immediate successors n $(n \in \mathbb{N})$, and V at each node with elementary functions the identity. If $* \models \alpha$ is a binary sequence, then using LPO in the meta-theory either $* \models \alpha(n) = 1$ for some particular n, or $* \models \alpha(n) = 0$ for all n, which fact persists to all the successor nodes. LPO holds at the other nodes because they are terminal nodes and so classical logic holds there. DC holds at all nodes using DC in the meta-theory. WLEM_ω fails: letting $0_n \subseteq 1$ be such that $k \models 0 \in 0_n$ iff $k = n$, the statements "$0 \in 0_n$" are pairwise incompatible yet none are false at *.

Model 3: Of course, while WLEM_ω is not true in the model above, it is also not false, because full classical logic holds at the terminal nodes. In order to get a *strong separation*, one in which WLEM_ω is false, we must iterate the previous construction.

Consider the partial order $\mathbb{N}^{<\mathbb{N}}$, ordered by end-extension, with V at each node and the identity for the embeddings. Take the full Kripke model over that p.o. Then DC and LPO hold at each node as above, and WLEM_ω fails at each node as above. ⊣

§5. **Markov's principle.** The next batch of models has to do with falsifying MP and its fragments, sometimes while retaining other parts of it.

Since MP has been so prominent for so long, it's no surprise that there have already been models developed that make it false. Such models have been of all sorts: topological [15, 19], Beth [6], realizability [20, 27], a mix [31]. Much of the earlier interest came from the incompatibility of MP with Kripke's Schema. That is, MP + KS proves full Excluded Middle ([29], p. 237), which was enough to stop Brouwer right there, and with any amount of continuity proves an outright contradiction. So some of the earlier falsifications of MP were almost accidental by-products of modeling KS plus continuity. Perhaps surprisingly, the most prominent model of Brouwer's intuitionism, Kleene's functional realizability K_2, actually satisfies MP (see [28], p. 428).

Each model falsifying MP of necessity provides one of the separations in Figure 2. The thing is, we don't know which. These fragments of MP have been less studied, and so the original authors did not concern themselves with which of them held in their models. It would be interesting to see what holds where. In what follows, we present instead what believe are new models, in part to have some more models in the literature, in part because this is how we view these separations, and in part because it's easier this way.

5.1. Weak Markov's principle.

THEOREM 5.1. *Over IZF + DC, WLEM does not imply WMP.*

PROOF. **Model 4:** This is only a weak separation, and may not even satisfy DC. Let $f : V \to M$ be an elementary embedding of the universe of sets V into a model M with nonstandard integers. For instance, M could be an ultrapower of V using a nonprincipal ultrafilter $\mathcal{U}$ on ω. Let $\mathcal{P}$ be the two-node partial order with bottom node $\bot$ and top node $\top$. Let $M_\bot$ be V and $M_\top$ be M. Consider the full Kripke model over that system.

WLEM is true because the partial order is linear. WMP fails by considering at $\bot$ a binary sequence α which always takes on the value 0 at all standard places and has a 1 in a nonstandard place. The hypothesis of WMP holds: given any β, either $\top \models \exists n\ \beta(n) = 1$, thereby satisfying the first disjunct, or it doesn't, forcing β to be the 0 sequence and thereby satisfying the second disjunct. But the conclusion of WMP fails at $\bot$.

Model 5: To get a strong separation, still potentially without DC, we can simply iterate the construction from above. So let $\mathcal{P}$ be ω. Let M_0 be V, M_1 be M, and let M_{n+1} be defined from M_n the way M was defined from V. For instance, if M is the ultrapower of V via the ultrafilter $\mathcal{U}$, then M_{n+1} is the ultrapower of M_n built within M_n via $f(\mathcal{U})$, where f is the (polymorphic) elementary embedding. Take the immediate settling model over this system.

WLEM and WMP stand and fall for the same reasons as above.

As promised (during the proof of IZF for the immediate settling model), we provide a central example of Power Set here, namely of 1, a.k.a. $\{0\}$. At any node,

1 viewed externally has three subsets: $0 = \emptyset$, 1 itself, and the set that now looks like 0 but at the next node will be 1, which we will call $1_\top$. So $\mathcal{P}(1)$ is given by $\mathcal{P}(1)(\bot) = \{0, 1, 1_\top\}$, and $\mathcal{P}(1)(\top) = f(\mathcal{P}(1)(\bot))$. So while the "three" members at $\bot$ collapse to two at $\top$, a new third member re-appears.

Model 6: This is a weak separation in which DC holds. The key idea of many of the following separation theorems is contained in this construction, which is a full topological model.

The points of the topological space are the natural numbers with an extra point at infinity: $X = \omega \cup \{*\}$. To define the topology on X take a nonprincipal ultrafilter $\mathcal{U}$ on ω. Define a topology T_X on X by taking

- $\emptyset \in T_X$
- $\{k\} \in T_X$ for each k in ω,
- $u \cup \{*\} \in T_X$ for each $u \in \mathcal{U}$,

and closing under union; for singleton opens, we write k in place of $\{k\}$.

The intuition is as follows. Interpret the elements of X as binary sequences with at most one nonzero term: $k \in \omega$ corresponds to the sequence α_k with a unique 1 in the k^{th} position, and $*$ corresponds to the constant zero sequence α_*. So the generic is a binary sequence α_G, with at most one nonzero term. Then $k \Vdash \alpha_G = \check{\alpha}_k$ and $u \cup \{*\} \Vdash \alpha_G(n) = 0$ for each $n \notin u$, since no extension of $u \cup \{*\}$ forces $\alpha_G(n) = 1$. Note that for each $n \in \omega$ since $\mathcal{U}$ is a nonprincipal ultrafilter there is a u in $\mathcal{U}$ such that $n \notin u$ and so $u \cup \{*\} \Vdash \alpha_G(n) = 0$. Hence, in M_X, α_G is indeed a binary sequence with at most one nonzero term.

We claim that the full topological model M_X over (X, T_X) satisfies WLEM and DC, and does not satisfy WMP.

For WLEM, let A be an arbitrary formula. Since isolated points behave classically, we need only give a neighbourhood of $*$ which forces $\neg A \vee \neg\neg A$. Let

$$X_0 = \{k \in \omega : k \Vdash \neg A\} \text{ and}$$
$$X_1 = \{k \in \omega : k \Vdash A\}.$$

Since the isolated nodes behave classically, $\omega = X_0 \cup X_1$ and so either $X_0 \in \mathcal{U}$ or $X_1 \in \mathcal{U}$. We show that in the former case, $X_0 \cup \{*\} \Vdash \neg A$ and in the latter $X_1 \cup \{*\} \Vdash \neg\neg A$. In the first case, every extension $\mathcal{O}$ of $X_0 \cup \{*\}$ can be extended to k for some $k \in X_0$. Since such a k forces $\neg A$, no extension of $X_0 \cup \{*\}$ can force A; thus $X_0 \cup \{*\} \Vdash \neg A$. The second case is similar: any extension of $X_1 \cup \{*\}$ can be further extended to some k forcing A, so no such extension can force $\neg A$.

As for dependent choice, since $\{k\}$ is open for each $k \in \omega$, X is dim-zero dimensional and hence, by Lemma 2.2, M_X satisfies DC.

While the failure of Weak Markov's Principle is stronger than that of simple Markov's Principle, the latter is easier to understand and more common, so it might be instructive to see first why MP fails for the generic α_G. Given an open set in X we can always extend to a singleton and so, since $k \Vdash \alpha_G(\check{k}) = 1$ for each singleton open k, no open forces that $\alpha_G(n) = 0$ for all n; that is $M_X \vDash \neg\forall n \in \omega\ \alpha_G(n) = 0$. Consider an open neighbourhood $\mathcal{O}$ of $*$ and fix $n \in \omega$. Then for any $k \in \mathcal{O} \setminus \{n, *\}$, k extends $\mathcal{O}$ and k forces $\alpha_G(n) = 0$. Hence $*$ does not force 'there exists n such that $\alpha_G(n) = 1$', so

$$M_X \nVdash \neg\forall n \in \omega\ \alpha_G(n) = 0 \to \exists n \in \omega\ \alpha_G(n) = 1.$$

Regarding WMP, since we have already seen that α_G is not forced to be 1 somewhere, it suffices to show that α_G is forced to be pseudo-positive. So let β be arbitrary. On the open set $\{k\}$ this is clear: $\alpha(k) = 1$, and $\beta(k)$ is either 0 or 1. We need only find an appropriate neighborhood of $\{*\}$. Let X_0 be $\{k \mid k \Vdash \exists n\ \beta(n) = 1\}$, and X_1 be $\{k \mid k \Vdash \forall n\ \beta(n) = 0\}$. If $X_0 \in \mathcal{U}$ then $X_0 \cup \{*\} \Vdash \neg\neg\exists n\ \beta(n) = 1$, else $X_1 \cup \{*\} \Vdash \neg\neg(\exists n\ \alpha(n) = 1 \wedge \beta(n) = 0)$.

Model 7: This is a strong separation in which DC holds, and can be viewed as an iteration of the previous construction. Let T be $\omega^{<\omega}$, the nodes of the countably branching tree. Recall that $\mathcal{U}$ is a nonprincipal ultrafilter on the natural numbers. A basic open set $\mathcal{O}$ contains a unique shortest node $\sigma_{\mathcal{O}}$, called the root, and, for all $\sigma \in \mathcal{O}$, $\{n \mid \sigma^\frown n \in \mathcal{O}\} \in \mathcal{U}$. It is easy to see that these are closed under intersection, and so an arbitrary open set is a union of these basic opens. It is also not hard to see that each basic open set is clopen. It's fair to call T an iteration of the earlier space X, because X naturally embeds into any basic open set $\mathcal{O}$, in such a way that $\mathcal{O}$ looks like a refinement of X: $\{^*\}$ gets sent to $\sigma_{\mathcal{O}}$, and ω is bijected with the immediate successors of $\sigma_{\mathcal{O}}$ in $\mathcal{O}$, where we identify the open set $\{n\}$ with the entire subset of $\mathcal{O}$ beneath (that is, the extensions of) the image of n.

Take the full topological model over T.

DC holds by Lemma 2.2, which applies because each basic open is clopen.

For WLEM, we have to show that each node σ is in some open set $\mathcal{O}$ either forcing $\neg A$ or forcing $\neg\neg A$. We say that σ is *determined* if some neighborhood of σ decides $\neg A$ (i.e. forces $\neg A$ or $\neg\neg A$); we can assume that such a neighborhood is a basic open set with σ as its root. Notice that if $\{n \mid \sigma^\frown n \text{ is determined}\} \in \mathcal{U}$ then σ is determined: letting $\mathcal{O}_n$ (with root $\sigma^\frown n$) decide $\neg A$ for each such n, one of $\{\sigma\} \cup \bigcup\{\mathcal{O}_n \mid \mathcal{O}_n \Vdash \neg A\}$ and $\{\sigma\} \cup \bigcup\{\mathcal{O}_n \mid \mathcal{O}_n \Vdash \neg\neg A\}$ is open.

We need to show that every node is determined. Suppose σ is not. We build a set of nodes $\mathcal{O}$ inductively, all of which are undetermined. We start with $\sigma \in \mathcal{O}$. Notice that if τ is undetermined then $\mathrm{UND}_\tau = \{n \mid \tau^\frown n \text{ is undetermined}\}$ is in $\mathcal{U}$. So for each $\tau \in \mathcal{O}$, include UND_τ in $\mathcal{O}$. By construction, $\mathcal{O}$ is open. If some open subset (without loss of generality, basic open) of $\mathcal{O}$ forces A, then its root is determined, contradicting its membership in $\mathcal{O}$. That means that $\mathcal{O}$ itself forces $\neg A$, which means that σ is determined, also a contradiction.

Regarding WMP, we show that no basic open set $\mathcal{O}$ with arbitrary root σ forces WMP, which suffices. Let α be a term such that, for k with $\sigma^\frown k \in \mathcal{O}$, we have $\mathcal{O}_k = \mathcal{O} \cap \{\tau \mid \tau \supseteq \sigma^\frown k\} \Vdash$ "$\alpha(k) = 1 \wedge \forall n \neq k\ \alpha(n) = 0$". Then $\mathcal{O} \nVdash \exists n\ \alpha(n) = 1$. It remains only to show that $\mathcal{O}$ forces the hypothesis of WMP.

Toward that end, let β be an arbitrary term for a binary sequence. Since $\mathcal{O}_k \Vdash \alpha(n) = 1 \leftrightarrow n = k$, beneath $\mathcal{O}_k$ the WMP hypothesis reduces to $\neg\neg\exists n\ \beta(n) = 1 \vee \beta(k) = 0$. By WLEM, we have $\neg\exists n\ \beta(n) = 1 \vee \neg\neg\exists n\ \beta(n) = 1$, which implies what we want for any $\mathcal{O}_k$. It remains only to find a neighborhood of σ forcing what we want.

Let X_0 be $\{n \mid$ some neighborhood of $\sigma^\frown n$ forces $\neg\neg\exists n\ \beta(n) = 1\}$ and X_1 be $\{n \mid$ some neighborhood of $\sigma^\frown n$ forces $\neg\exists n\ \beta(n) = 1\}$. If $X_0 \in \mathcal{U}$ then some neighborhood of σ forces the first disjunct in WMP; if $X_1 \in \mathcal{U}$ then some neighborhood of σ forces the second. ⊣

There is a different way of looking at the previous construction, which we sketch briefly (as opposed to calling it a fifth construction ;-)).

Rather than having the entire tree $\omega^{<\omega}$ be present all at once, it can be rolled out level-by-level within a Kripke model. Toward this end, let $\mathcal{K}$ be the full Kripke model over the partial order ω. Let the topological space T consist at node n of $\omega^{\leq n}$. So, for instance, at node 0, T contains as a subset $\omega^0 = \omega^{\emptyset} = \{\emptyset\}$, so T looks like a single-point space, $\{^*\}$. A basic open set $\mathcal{O}$ at node n contains a unique shortest sequence σ, and if, at n, $\tau \in \mathcal{O}$ has length less than n, then $\{i \mid n \models \tau^\frown i \in \mathcal{O}\} \in \mathcal{U}$. Finally, we need to explain what $f_n(\mathcal{O})$ looks like, where f_n is the transition function from node n. For $\mathcal{O}$ basic open at Kripke node n, $n+1 \models \tau \in f_n(\mathcal{O})$ if $\tau \in \mathcal{O}$ or if $\tau = \sigma^\frown m$ has length $n+1$ and $\sigma \in \mathcal{O}$. So if a basic open set at a node contains what looks like a terminal sequence of T, that seemingly isolated point actually stands for the entire space extending it.

Notice that we have chosen a restricted notion of open set. For instance, at node 0 the only open set is the entire space. Other subsets of T exist in $\mathcal{K}$, which could legitimately have been taken to be open. For instance, the set which at 0 contains $\emptyset$ and at 1 contains $\mathcal{U}$-many but not all extensions. This restricted notion demands a corresponding restriction in the sets allowed in the model. At node n, no name for a set in the topological model may grow at future nodes. That is, suppose t is a name for a set in the model at node n, so that t consists of pairs of the form $\langle s, \mathcal{O}\rangle$, where s is also such a name and $\mathcal{O}$ an open set at n. Then at any future node $f_n(t)$ consists entirely of pairs $\langle f_n(s), f_n(\mathcal{O})\rangle$. In particular, no new open sets may appear within $f_n(t)$.

The proof that a full topological model satisfies IZF is valid constructively, and so holds within $\mathcal{K}$. That our restricted topological model satisfies IZF would have to be checked in detail. DC holds by re-doing the proof of Proposition 2.2, using DC in $\mathcal{K}$ and the countability of each level of T. WMP fails for the same reason as in the previous construction. The validation of WLEM uses the fact that the truth value of a statement at node n is an open set, of the kind described, at node n.

For a sixth construction, start as above, while allowing all possible open sets, and take the full model.

5.2. Separating the hierarchies. In this section we separate the WLEM_n, LLPO_n, and $\mathrm{MP}_n^\vee$ hierarchies in the strongest ways possible by Figure 1, as indicated by Figure 2.

THEOREM 5.2. *Over IZF + DC, $WLEM_{n+1}$ does not imply $MP_n^\vee$ for each n.*

Note that Theorem 5.2 separates all three of the hierarchies at once, for if WLEM_{n+1} does not imply $\mathrm{MP}_n^\vee$, then it does not imply the stronger principles WLEM_n or LLPO_n, and hence these latter principles cannot be proved from LLPO_{n+1} or $\mathrm{MP}_{n+1}^\vee$ which are weaker than WLEM_{n+1}.

PROOF. **Model 8:** (A weak separation without DC.) Corresponding to Model 4, consider the Kripke partial order with root $\perp$ and n-many immediate successors 0, $\ldots$, $n-1$. At the successor nodes let M_i be an ultrapower of V with nonstandard integers. Let $M_\perp$ be V. Take the full model over that system. WLEM_{n+1} holds because, given $n+1$-many mutually incompatible assertions, at each successor node at most one is true, so there's at least one false in all of them, and that one is false at $\perp$. $\mathrm{MP}_n^\vee$ fails, by taking the sequence which is all 0's at $\perp$, and at node i has a 1 in a (necessarily nonstandard) slot in the i^{th} slice of ω.

Model 9: (A strong separation without DC.) As in Model 5, let the Kripke partial order be the n-branching tree. The base model for a node of length k is M_k. Take the immediate settling model.

$\mathrm{MP}_n^\vee$ fails as above. To see why WLEM_{n+1} holds, consider a statement A at node σ. Because of A's parameters, the truth of A at some successor $\sigma^\frown i$ might be different from at $\sigma^\frown j$. But by elementarity, the truth at $\sigma^\frown i$ is the same as at all of its successors.

Model 10: (A weak separation with DC.) This is a topological model corresponding to Model 6. Since the notation gets sufficiently ghastly to obscure the simple idea behind these models, we give first the case $n = 2$ before sketching the general construction.

Let $X = \omega \cup \{*\}$ as before and let $\mathcal{U}$ be a nonprincipal ultrafilter on ω. Let f_0, f_1 be the functions on ω sending n to $2n$ and $2n+1$, respectively, and set $\mathcal{U}_i = \{f_i{}"u : u \in \mathcal{U}\}$ and $\omega_i = f_i{}"\omega$ $(i = 0, 1)$. We give X the topology generated by taking

- the emptyset $\emptyset$,
- $\{k\}$ for each $k \in \omega$, and
- $u_0 \cup u_1 \cup \{*\}$ where $u_i \in \mathcal{U}_i$ $(i = 0, 1)$,

and closing under unions. Let M_X be the full model over X. As in the previous section, we can associate to the generic a binary sequence α_G forced by X to have at most one nonzero term and such that $\neg\forall n \in \omega\ \alpha_G(n) = 0$. Recall, in particular, that $k \Vdash \alpha_G = \check{\alpha}_k$ for each $k \in \omega$. Since any open $\mathcal{O}$ containing $*$ can be extended to a singleton $k_0 \in \omega_0$ and a singleton $k_1 \in \omega_1$, $*$ does not force $\forall n \in \omega\ \alpha_G(2n) = 0 \vee \forall n \in \omega\ \alpha_G(2n+1) = 0$. Hence $M_X \not\Vdash \mathbf{MP}^\vee$.

Suppose, without loss of generality, that X forces $\neg\bigvee_{i,j=0,i\neq j}^2 A_i \wedge A_j$ for some formulas A_0, A_1, A_2 with parameters from $V(X)$. For $i = 0, 1$ we define

$$X_0^i = \{k \in \omega_i : k \Vdash \neg A_1 \wedge \neg A_2\};$$
$$X_1^i = \{k \in \omega_i : k \Vdash \neg A_0 \wedge \neg A_2\};$$
$$X_2^i = \{k \in \omega_i : k \Vdash \neg A_0 \wedge \neg A_1\}.$$

Since $\omega_i = X_0^i \cup X_1^i \cup X_2^i$ $(i = 0, 1)$, one of X_0^0, X_1^0, X_2^0 is in $\mathcal{U}_0$ and one of X_0^1, X_1^1, X_2^1 is in $\mathcal{U}_1$. Taking for illustration the case where $X_0^0 \in \mathcal{U}_0$ and $X_1^1 \in \mathcal{U}_1$ we have that $\mathcal{O} = X_0^0 \cup X_1^1 \cup \{*\}$ is an open subset of X. Since any extension of $\mathcal{O}$ can be further extended to some $k \in X_0^0 \cup X_1^1$ and each such k forces that $\neg A_2$, $\mathcal{O}$ forces $\neg\neg\neg A_2$ or equivalently $\neg A_2$. The other cases are essentially the same. Hence $M_X \vDash WLEM_3$. DC holds by Lemma 2.2.

For arbitrary n, let $X, \mathcal{U}$ be as before. For $i = 0, \ldots, n-1$ define a function f_i on ω by $f_i : m \mapsto nm + i$; let $\mathcal{U}_i$ be the image of $\mathcal{U}$ under f_i and ω_i be the image of ω under f_i. A base for the topology on X is given by taking

- the emptyset $\emptyset$,
- $\{k\}$ for each $k \in \omega$, and
- $u_0 \cup \cdots \cup u_{n-1} \cup \{*\}$ where $u_i \in \mathcal{U}_i$ $(0 \leqslant i \leqslant n-1)$.

The generic demonstrates that $*$ does not satisfy $\mathrm{MP}_n^\vee$.

For any formulas $A_0, \ldots, A_n$, with parameters from $V(X)$, such that X forces $\neg \bigvee_{i,j=0, i\neq j}^{n} A_i \wedge A_j$, we form the sets

$$X_j^i = \{k \in \omega_i : k \Vdash \bigwedge_{l \neq j} \neg A_l\} \qquad (0 \leqslant i \leqslant n-1, 0 \leqslant j \leqslant n).$$

Since $\mathcal{U}$ is an ultrafilter, for each i there exists j_i such that $X_{j_i}^i \in \mathcal{U}_i$. Then

$$\mathcal{O} = X_{j_0}^0 \cup \cdots \cup X_{j_{n-1}}^{n-1} \cup \{*\}$$

is an open set and by the pigeon-hole principle there exists $j \in \{0, \ldots, n\}$ such that $\mathcal{O}$ forces $\neg A_j$.
DC holds, by Lemma 2.2.

The isolated points are dense in X, so Excluded Middle holds densely, hence this is a weak separation.

Model 11: We would like a strong separation with DC. The most straightforward way you'd think to do this is a start, but does not work entirely. That would be to adapt Model 7. Namely (in the case n=2, say), let the points be $\omega^{<\omega}$; consider two copies of $\mathcal{U}$, one of which concentrates on the even naturals and the other on the odds; and say that a set $\mathcal{O}$ is open if, whenever $\sigma \in \mathcal{O}$, each of the set of even and the set of odd children of σ in $\mathcal{O}$ is in its copy of $\mathcal{U}$. That doesn't do what we want, because by considering the grandchildren of σ there's enough room for four statements to be true.

The additional component needed is like immediate settling, taking the submodel consisting of those sets that can change from a node to any of the node's children, and then have to stop changing. In order for the model not to be trivial at those children, new sets have to be introduced, which then themselves have one generation during which they are allowed to change. What this involves is partial existence: there are some sets at nodes that are not the images of any sets from earlier nodes. This phenomenon is not new. It is perhaps best known via Kripke models with nonconstant domains, and has also been studied in more general contexts, first in [24] and later exposited in [9] and [30] (section 13.6), where the logic of partial existence was worked out. The model here has aspects of a Kripke model, in that it is based on a partial order, $\omega^{<\omega}$, and depends on partial existence. But it's not a Kripke model, instead making essential use of the topology, because a statement being true at a node does not mean that it's true at all children, rather merely on an open set of children. It could be called the immediate settling topological model. This situation imposes an additional burden: since no extant meta-theorem applies, IZF has to be checked by hand. (Possibly this is a full topological model within the Kripke immediate settling model, but at the very least that would have to be checked. We prefer just to check IZF directly).

The guiding intuition is still the space from Model 7. That would be to work over $\omega^{<\omega}$, and to require of an open set $\mathcal{O}$ (in the case of $n = 2$, for notational simplicity) that for $\sigma \in \mathcal{O}$ both $\{n \mid \sigma^\frown 2n \in \mathcal{O}\}$ and $\{n \mid \sigma^\frown 2n+1 \in \mathcal{O}\}$ be in $\mathcal{U}$. The fact is, though, that the model is homogeneous: every node looks like every other node. So it suffices to work with only the bottom level, $\{\perp\} \cup \omega$.

The objects at all nodes are defined inductively, as well as the transition functions f_i from $\perp$ to i. The sets built by stage α will be the same at all nodes, and so

the notation $\mathcal{T}_\alpha$ suffices. The construction will be a lot like the previous theorem's third construction. A big difference is that here the base models are always V. This obviates of course all concerns about where these objects exist. Another difference is that there it was easier to formalize the notion of a set not changing anymore, because there was only one successor node, whereas here we have infinitely many successors. Hence there is a need for the notion of *canonical sets* $\mathcal{CT}_\alpha$, which we also define inductively, as those functions g such that

- $\mathrm{dom}(g) = \{\bot\} \cup \omega$,
- $g(\bot), g(i) \subseteq \bigcup_{\beta<\alpha} \mathcal{T}_\beta$,
- $g(\bot) = g(i)$, and
- if $h \in g(\bot)$ then $f_i(h) \in g(\bot)$.

As for $\mathcal{T}_\alpha$,

- $\mathrm{dom}(g) = \{\bot\} \cup \omega$,
- $g(\bot), g(i) \subseteq \bigcup_{\beta<\alpha} \mathcal{T}_\beta$,
- if $h \in g(i)$ then $f_j(h) \in g(i)$, and
- if $h \in g(\bot)$ then for all $i < n$, $\{k \mid f_{kn+i}(h) \in g(kn+i)\} \in \mathcal{U}$.

In addition, we must extend f_i to these new sets: $f_i(g)(\bot) = g(i) = f_i(g)(j)$. It is easy to check that $\mathcal{CT}_\alpha \subseteq \mathcal{T}_\alpha$, that $f_i(g) \in \mathcal{CT}_\alpha$ and, for $g \in \mathcal{CT}_\alpha$, $f_i(g) = g$.

Now we need to define truth in this model. Inductively on formulas ϕ, we define a set $[\![\phi]\!] \subseteq \{\bot\} \cup \omega$. By way of notation, $f_i(\phi)$ refers to the result of applying f_i to each of ϕ's parameters, and $f_\bot(g) = g$. Also, if $A \subseteq \omega$, then A_i is the i^{th} slice of A: $A_i = \{k \mid kn + i \in A\}$. If $A \subseteq \omega \cup \{\bot\}$ then by A_i we mean $(A \cap \omega)_i$.

- $[\![g \in h]\!] = \{q \mid \exists f \in h(q)\ \bot \in [\![f_q(g) = f]\!]\}$
- $[\![g = h]\!] = \{q \in \omega \mid \text{for all } f \in g(q)\ \bot \in [\![f \in f_q(h)]\!], \text{ and vice versa}\} \cup$
  $\{\bot \mid \forall f \in g(\bot)\ \bot \in [\![f \in h]\!] \text{ and vice versa, and } \forall i < n\ [\![g = h]\!]_i \in \mathcal{U}\}$
- $[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$
- $[\![\phi \vee \psi]\!] = [\![\phi]\!] \cup [\![\psi]\!]$
- $[\![\phi \to \psi]\!] = [\![\psi]\!] \cup (\omega \backslash [\![\phi]\!]) \cup$
  $\{\bot \mid \bot \notin [\![\phi]\!] \wedge \forall i < n\ [[\![\psi]\!] \cup (\omega \backslash [\![\phi]\!])]_i \in \mathcal{U}\}$
- $[\![\exists x\ \phi(x)]\!] = \{q \mid \exists h\ q \in [\![\phi(h)]\!]\}$
- $[\![\forall x\ \phi(x)]\!] = \{q \in \omega \mid \text{for all } h\ \bot \in [\![f_q(\phi)(h)]\!]\} \cup$
  $\{\bot \mid \text{for all } h\ \bot \in [\![\phi(h)]\!], \text{ and } \forall i < n\ [\![\forall x\ \phi(x)]\!]_i \in \mathcal{U}\}$

LEMMA 5.3. a) *$i \in [\![\phi]\!]$ iff $\bot \in [\![f_i(\phi)]\!]$.*
b) *If $\bot \in [\![\phi]\!]$ then for $i < n$ $[\![\phi]\!]_i \in \mathcal{U}$.*

PROOF. Both parts are proved together via a simultaneous induction.
Regarding a), for $i = \bot$ this is trivial.

Otherwise, $i \in [\![g \in h]\!]$ iff for some $f \in h(i)$ we have $\bot \in [\![f_i(g) = f]\!]$. On the other hand, $f_i(g = h)$ is $f_i(g) = f_i(h)$, so $\bot \in [\![f_i(g = h)]\!]$ iff for some $f \in h(i)$ we have $\bot \in [\![f_\bot(f_i(g)) = f]\!]$. But $f_\bot$ is the identity function, so those are the same.

For $=$, $i \in [\![g = h]\!]$ iff for all $f \in g(i)$ we have $\bot \in [\![f \in f_i(h)]\!]$ and vice versa. Notice this is exactly the first clause of $\bot \in f_i(g = h)$, so the right-to-left direction is proven. The second clause of $\bot \in f_i(g = h)$ is $[\![f_i(g) = f_i(h)]\!]_j \in \mathcal{U}$. If $i \in [\![g = h]\!]$, then it is straightforward to check that $[\![f_i(g) = f_i(h)]\!]$ contains all of ω, and so each slice is in the ultrafilter.

The cases $[\![\phi \wedge \psi]\!]$ and $[\![\phi \vee \psi]\!]$ are trivial to check via induction.

For $\rightarrow$, we have $i \in [\![\phi \rightarrow \psi]\!]$ iff $i \in [\![\psi]\!]$ or $i \notin [\![\phi]\!]$. The first case inductively is equivalent to $\perp \in [\![f_i(\psi)]\!]$, which is one of the two ways that $\perp$ gets into $[\![f_i(\phi \rightarrow \psi)]\!]$. The other way $\perp$ gets in is if $\perp \notin [\![f_i(\phi)]\!]$ and for each $j < n$ we have $[[\![f_i(\psi)]\!] \cup (\omega \backslash [\![f_i(\phi)]\!])]_j \in \mathcal{U}\}$. The first of those two clauses is equivalent inductively to the second case. Furthermore, if $\perp \notin [\![f_i(\phi)]\!]$, then, inductively, $\perp \notin [\![f_\perp(f_i(\phi))]\!] = [\![f_j(f_i(\phi))]\!]$, so $j \notin [\![f_i(\phi)]\!]$, which is why $[\![f_i(\phi)]\!]$ is empty.

For the quantifiers, consider $\exists$: $i \in [\![\exists x\ \phi(x)]\!]$ iff for some h we have $i \in [\![\phi(h)]\!]$ iff, inductively, for some h $\perp \in [\![f_i(\phi(h))]\!] = [\![f_i(\phi)(f_i(h))]\!]$. If that's true, then, for some h, namely $f_i(h)$, we have $\perp \in [\![f_i(\phi)(h)]\!]$, so $\perp \in [\![\exists x\ f_i(\phi)(x)]\!] = [\![f_i(\exists x\ \phi(x))]\!]$. In the other direction, suppose $\perp \in [\![f_i(\exists x\ \phi(x))]\!]$, i.e. for some h we have $\perp \in [\![f_i(\phi)(h)]\!]$. Using part b) inductively, there is a $j \in \omega$ with $j \in [\![f_i(\phi)(h)]\!]$. Using a) inductively, $\perp \in [\![f_j(f_i(\phi)(h))]\!] = [\![f_j(f_i(\phi))(f_j(h))]\!] = [\![f_i(\phi)(f_j(h))]\!]$. But $f_j(h) = f_i(h \circ \sigma)$, where σ is the permutation that interchanges i and j and leaves everything else alone. So for some h, namely $h \circ \sigma$, we have $\perp \in [\![f_i(\phi)(f_i(h))]\!]$, which we have already seen is equivalent to $i \in [\![\exists x\ \phi(x)]\!]$.

Finally, $i \in [\![\forall x\ \phi(x)]\!]$ iff for each h we have $\perp \in [\![f_i(\phi)(h)]\!]$. Notice that condition is exactly the first clause of $\perp \in [\![f_i(\forall x\ \phi(x))]\!]$. The second clause is $\forall j < n\ [\![\forall x\ f_i(\phi)(x)]\!]_j \in \mathcal{U}$. So suppose that for each h we have $\perp \in [\![f_i(\phi)(h)]\!]$. We want to see when $k \in [\![\forall x\ f_i(\phi)(x)]\!]$, that is, whether for all h $\perp \in [\![f_k(f_i(\phi))(h)]\!]$. But $[\![f_k(f_i(\phi))(h)]\!] = [\![f_i(\phi)(h)]\!]$, so this is exactly our supposition.

Now to show part b). There are some easy cases: if ϕ is an equality or a universal statement, this is built right into the definition of $[\![\phi]\!]$. If ϕ is a disjunction or conjunction, this follows easily from $\mathcal{U}$ being an ultrafilter. For an implication, if $\perp \in [\![\psi]\!]$, inductively each $[\![\psi]\!]_i \in \mathcal{U}$, and those slices are subsets of $[[\![\psi]\!] \cup (\omega \backslash [\![\phi]\!])]_i$, which are then also in $\mathcal{U}$. If $\perp \notin [\![\psi]\!]$ then there is only one other way $\perp$ can get into $[\![\phi \rightarrow \psi]\!]$, and that other way has part b) built in.

For the existential case, suppose for some h we have $\perp \in [\![\phi(h)]\!]$. Inductively, for each $i < n$, $[\![\phi(h)]\!]_i \in \mathcal{U}$. Suppose $j \in [\![\phi(h)]\!]_i$, i.e. $jn + i \in [\![\phi(h)]\!]$. So for that h it holds that $jn + i \in [\![\phi(h)]\!]$, which means $jn + i \in [\![\exists x\ \phi(x)]\!]$. In short, $[\![\phi(h)]\!]_i \subseteq [\![\exists x\ \phi(x)]\!]_i$, and so the latter set is also in $\mathcal{U}$.

Finally, consider membership. If $\perp \in [\![g \in h]\!]$, then there is an $f \in h(\perp)$ with $\perp \in [\![g = f]\!]$. We therefore have each $[\![g = f]\!]_i$ in $\mathcal{U}$. If j is in that set, by part a), $\perp \in [\![f_{jn+i}(g = f)]\!] = [\![f_{jn+i}(g) = f_{jn+i}(f)]\!]$. Separately from that, by the last clause of the definition of a set, $\{k \mid f_k(f) \in h(k)\}_i$ is in $\mathcal{U}$. Now if j is in that latter set, then $f_{jn+i}(f) \in h(jn + i)$. If j is in both sets, then $j \in [\![g \in h]\!]_i$, and of course ultrafilters are closed under intersection. $\dashv$

Corollary 5.4. *If all of ϕ's parameters are canonical, then $[\![\phi]\!]$ is either empty or $\omega \cup \{\perp\}$.*

Proof. Suppose $[\![\phi]\!]$ is nonempty. If $\perp \in [\![\phi]\!] = [\![f_i(\phi)]\!]$ then by a) $i \in [\![\phi]\!]$. If $i \in [\![\phi]\!] \cap \omega$ then again by a) $\perp \in [\![f_i(\phi)]\!] = [\![f_j(\phi)]\!]$, so $j \in [\![\phi]\!]$. $\dashv$

Lemma 5.5. *The equality axioms (reflexivity, symmetry, transitivity) all get full value.*

Now we check that that the axioms of IZF + DC get full value. First note that the ground model embeds into this model as the hereditarily constant functions. So Emptyset and Infinity hold, as witnessed by the images of $\emptyset$ and ω. Pair holds, for,

given f and g, let h be such that $h(\bot) = \{f, g\}$ and $h(i) = \{f_i(f), f_i(g)\}$. For Union, given g, let $(\bigcup g)(\bot)$ be $\{f \mid f \in h(\bot) \text{ for some } h \in g(\bot)\}$, and $(\bigcup g)(i)$ be $\{f \mid f \in h(\bot) \text{ for some } h \in g(i)\}$. We leave the verification that these four constructions satisfy their axioms to the reader.

About the two structural axioms, $\in$-Induction is simple, because the sets were defined inductively on the ordinals.

Regarding Extensionality, we need to show that $[\![f = g \leftrightarrow \forall h\ (h \in f \leftrightarrow h \in g)]\!] = \{\bot\} \cup \omega$. That is equivalent to $[\![f = g]\!] = [\![\forall h\ (h \in f \leftrightarrow h \in g)]\!]$. Consider first $i \in \omega$. Then $i \in [\![f = g]\!]$ iff for all $h \in f(i)$ it holds that $\bot \in [\![h \in f_i(g)]\!]$ and vice versa. Also, $i \in [\![\forall h\ (h \in f \leftrightarrow h \in g)]\!]$ iff for all h we have $\bot \in [\![h \in f_i(f) \leftrightarrow h \in f_i(g)]\!]$, which means $\bot$ is in both $[\![h \in f_i(f) \rightarrow h \in f_i(g)]\!]$ and $[\![h \in f_i(g) \rightarrow h \in f_i(f)]\!]$. We will show for all $h \in f(i)$ $(\bot \in [\![h \in f_i(g)]\!])$ if and only if for all h $(\bot \in [\![h \in f_i(f) \rightarrow h \in f_i(g)]\!])$, the reasoning for the other directions being analogous.

Considering the left-to-right direction first. Let h be arbitrary. The RHS holds when either $\bot \in [\![h \in f_i(g)]\!]$, or both $\bot \notin [\![h \in f_i(f)]\!]$ and for each $j < n$ the j^{th} slice of $[\![h \in f_i(g)]\!] \cup (\omega \backslash [\![h \in f_i(f)]\!])$ is in $\mathcal{U}$. Either $\bot \in [\![h \in f_i(f)]\!]$ or not. If so, then for some $\hat{h} \in f(i)$ it holds that $\bot \in [\![h = \hat{h}]\!]$. By hypothesis, $\bot \in [\![\hat{h} \in f_i(g)$. That means that for some $\hat{g} \in g(i)$ we'd have $\bot \in [\![\hat{h} = \hat{g}]\!]$. By transitivity, $\bot \in [\![h = \hat{g}$, and we'd be done with this direction. The other possibility is $\bot \notin [\![h \in f_i(f)]\!]$. Let $q \in \omega$ be arbitrary. If $q \in [\![h \in f_i(f)]\!]$ then for some $\hat{h} \in f_i(f)(q)$ it holds that $\bot \in [\![f_q(h) = \hat{h}]\!]$. Then by the lemma $\bot \in [\![f_q(h) \in f_q(f_i(f)) = f_i(f)]\!]$, meaning that for some $\hat{h} \in f(i)$ $\bot \in [\![f_q(h) = \hat{h}$. By hypothesis, $\bot \in [\![\hat{h} \in f_i(g)]\!]$. Arguing as above with transitivity, $\bot \in [\![f_q(h) \in f_i(g) = f_q(f_i(g))]\!]$. Again via the lemma, $q \in [\![h \in f_i(g)]\!]$. This means that $[\![h \in f_i(g)]\!] \cup (\omega \backslash [\![h \in f_i(f)]\!])[\![h \in f_i(g)]\!] \cup (\omega \backslash [\![h \in f_i(f)]\!])[\![h \in f_i(g)]\!] \cup (\omega \backslash [\![h \in f_i(f)]\!])[\![h \in f_i(g)]\!] \cup (\omega \backslash [\![h \in f_i(f)]\!])$ is all of ω, and again we're done.

Now consider the right-to-left direction. Let h be in $f(i)$. By hypothesis, $\bot \in [\![h \in f_i(f) \rightarrow h \in f_i(g)]\!]$. Since $[\![h \in f_i(f)]\!]$ gets full value, so does $[\![h \in f_i(g)]\!]$.

We still need to handle the case $i = \bot$. This is similar, and left to the reader.

The final axioms IZF axioms are Collection, Separation, and Power Set. Actually, not just Collection, but even Reflection is true. Let α be such that $V_\alpha \prec_{\Sigma_n} V$. Let g_α be the construction of the model up to stage α: $g_\alpha(q) = \bigcup_{\beta<\alpha} \mathcal{CT}_\beta$. Then any formula of classical complexity Σ_n reflects to g_α. For Separation, for $\{h \in a \mid \phi(h)\}$, let $g(q)$ be $\{h \in a(q) \mid \bot \in [\![f_q(\phi)(h)]\!]\}$. We illustrated Power Set via the most critical case, $\mathcal{P}(1)$. Of course, the empty set is represented in the model by the constant function which we either suggestively or ambiguously call 0: $0(q) = \emptyset$. Similarly, $1(q) = \{0\}$. A subset of 1 is given by certain sets A, in that $1_A(q) = \{0\}$ if $q \in A$ and $1_A(q) = \emptyset$ otherwise. For such a function 1_A to be a set in the model, if $\bot \in A$ then each A_i $(i < n)$ must be in $\mathcal{U}$. Then $\mathcal{P}(1)(q) = \{1_A \mid 1_A$ is a set in the model$\}$. More generally, h represents a subset of g if $h(q) \subseteq g(q)$ for all q, and in addition h represents a set in the model by satisfying clauses 3 and 4 in the definition of $\mathcal{T}_\alpha$; $\mathcal{P}(g)(q)$ is the set of those h's representing a subset of $f_q(g)$.

For DC, it is straightforward to build a choice sequence, using DC in the meta-theory.

Finally, WLEM_{n+1} gets full value, and $\mathrm{MP}_n^\vee$ the empty set, as follows. Given pairwise incompatible $\phi_i, i \leq n$, the $[\![\phi_i]\!]$ are disjoint. So, for a fixed $k < n$, there is at most one $i \leq n$ with $[\![\phi_i]\!]_k \in \mathcal{U}$. Hence there is a fixed i with each $[\![\phi_i]\!]_k$ not in $\mathcal{U}$. For that i, $\bot \in [\![\neg\phi_i]\!]$. Of course, for each $j \notin [\![\phi_i]\!]$ (for that matter, for each $j \in \omega$), there is some i with $j \in [\![\neg\phi_i]\!]$. So $[\![\exists i\ \neg\phi_i]\!] = \{\bot\} \cup \omega$. In contrast, let α be such that $[\![\alpha(k) = 1]\!] = \omega_k$ (for $k < n$; for $k \geq n$ let $[\![\alpha(k) = 1]\!] = \emptyset$). Then α is a counter-example to $\mathrm{MP}_n^\vee$. ⊣

5.3. The other MP separations.

THEOREM 5.6. *Over IZF + DC, WMP does not imply $MP_\omega^\vee$.*

PROOF. **Model 12:** Actually, it is a misnomer to call this a model, since we argue that any of the constructions from above generalize quite easily, and a misnomer to call it a proof, since we just sketch matters quickly. For instance, take the full Kripke model with root $\bot$ and associated model V, and countably many successors with associated model some ω-non-standard M. To show that $\mathrm{MP}_\omega^\vee$ fails, let α at $\bot$ look like all 0's, and at node n have a 1 in some nonstandard position in the n^{th} block. To see that WMP holds, we need only consider $\bot$, and a sequence α there which is 0 on the standard part. If at one of the terminal nodes α is the 0 sequence (that is, $\neg\neg\exists n\ \alpha(n) = 1$ fails at $\bot$), then picking β to be α shows that the antecedent of WMP does not hold at $\bot$. Otherwise, let β be α except that on one of the terminal nodes β is the 0 sequence. Then neither disjunct holds for β at $\bot$, and again the antecedent does not hold.

This can be iterated with immediate settling. For that matter, since we're no longer trying to preserve any part of WLEM, we can take the full Kripke model over an appropriately branching tree (defined below), with base M_k for nodes of length k. Since the model beyond a node of length k has to be definable within M_k, the branching there has to be indexed by the natural numbers of M_k. $\mathrm{MP}_\omega^\vee$ fails at a node σ of length k, by considering the sequence which is 0 there and at node $\sigma^\frown n$ has a 1 in some M_k-non-standard position in the n^{th} block. WMP holds exactly as above.

Or, to get DC, we could take a topological model. Let $\{\omega_i : i \in \omega\}$ be the partition of ω into blocks of size ω used in the statement of $\mathrm{MP}_\omega^\vee$. Let f_i be a bijection between ω and ω_i for each i. Let $\mathcal{U}$ be a nonprincipal ultrafilter on ω and $\mathcal{U}_i$ the image of $\mathcal{U}$ under f_i $(i \in \omega)$. We take $X = \omega \cup \{*\}$ and form opens by taking unions of: $\emptyset$, $\{k\}$ $(k \in \omega)$, and sets of the form $\bigcup\{u_i : u_i \in \mathcal{U}_i\} \cup \{*\}$. The generic binary sequence (that is, the sequence forced by k to be 1 at index k and 0 elsewhere) is then a (weak) counterexample to $\mathrm{MP}_\omega^\vee$. For WMP, if no neighborhood of $*$ forces $\neg\neg\exists n\ \alpha(n) = 1$ then by choosing β to be α the antecedent of WMP does not hold for α; else, if $\{*\} \cup u$ is a neighborhood forcing such, let β be α except for some $k \in u$ let $k \Vdash \beta = \mathbf{0}$. DC holds by Lemma 2.2. Again, this can be iterated in any of a number of ways as above.

Model 13: As it turns out, there is another variant along the same lines that works. Let X be $\omega \cup \{*\}$. Take the topology on X for which the basic opens are of the form $\emptyset, k$, and $\{*\} \cup u$ where u is cofinite. Take the full topological model. $\mathrm{MP}_\omega^\vee$ fails, and WMP and DC hold, as above.

Model 14: Interestingly, this is a natural model one might actually wonder about for its own sake. Consider the topological model over $\mathbb{Q}$.

By Proposition 2.2, DC holds.

To see that $\mathrm{MP}^{\vee}_{\omega}$ fails, let $r \in \mathbb{Q}$. Let $\mathcal{O}_n$ be a nested sequence of open intervals with irrational endpoints and intersection $\{r\}$. Let α be such that $\mathcal{O}_n \Vdash \alpha(n) = 0$ and $\mathbb{Q}\backslash\mathcal{O}_n \Vdash \alpha(n) = 1$. Any neighborhood of r can force α to be 0 on only finitely many values. So on no particular block, each of which is infinite, can α be forced to be 0 everywhere.

In contrast, WMP holds. Let $\mathcal{O} \Vdash \forall\beta(\neg\neg\exists n\ (\beta(n) = 1) \vee \neg\neg\exists n\ (\alpha(n) = 1 \wedge \beta(n) = 0))$ and $r \in \mathcal{O}$. We must find a neighborhood of r forcing $\exists n\ \alpha(n) = 1$. For each n, let $\mathcal{O}_n$ be the maximal interval containing r which decides the value of $\alpha(i)$ for all $i < n$. If any $\mathcal{O}_n$ forces any $\alpha(i)$ to be 1, we are done. Suppose not. It cannot be the case that $\bigcap_n \mathcal{O}_n$ contains an open interval $\mathcal{I}$ with $r \in I$, because, in that case, letting β be the constant 0 function, no subinterval of $\mathcal{I}$ could force either $\neg\neg\exists n\ (\beta(n) = 1)$ or $\neg\neg\exists n\ (\alpha(n) = 1)$, contrary to hypothesis. So $\bigcap_n \mathcal{O}_n$ has r as either its left endpoint or its right endpoint (or both). Assume without loss of generality it's the first option. Define β as follows. For one, $(r, \infty) \Vdash \beta(n) = 0$ for all n. Hence, $(r, \infty) \Vdash \neg\exists n\ \beta(n) = 1$. Also, $\mathcal{O}_n \Vdash \beta(n) = 0$. Finally, $(-\infty, \inf \mathcal{O}_n) \Vdash \beta(n) = 1$. That means $(-\infty, r) \Vdash \neg\exists n\ (\alpha(n) = 1 \wedge \beta(n) = 0)$, because whenever α is forced to be 1 so is β. So no neighborhood of r can force either disjunct to the hypothesis of WMP, contrary to assumption. This contradiction finishes the proof.

It's worth pointing out that the same arguments show the same result for the model over Cantor space 2^{ω}.

It's also worth observing what happens in the model over $\mathbb{R}$. The binary sequences in that model are (the images of) binary sequences from the ground model, so MP for binary sequences holds. But by similar arguments to those here, MP for reals fails, while WMP for reals holds. (See the final section, Concluding Remarks, for further discussion on related topics.) ⊣

THEOREM 5.7. *There is a model of IZF + DC in which WLEM_{ω} holds but each $\mathrm{MP}^{\vee}_{n}$ fails.*

PROOF. We sketch the adaptations of the main constructions from earlier. For a weak separation without DC, start with a Kripke partial order with $\perp$ on bottom and successors 0, 1, 2, ... To node $\perp$ associate V and to node n associate an ω-non-standard ultrapower M of V, the same for all n. At $\perp$, take all sets that eventually become constant. That is, with k_n the transition function from $\perp$ to n, take all sets x such that, for some i and all $n > i$, $f_n(x) = f_i(x)$. It's easy to check that this satisfies IZF. WLEM_{ω} holds, as follows. Given an ω-sequence of mutually incompatible statements A_j at $\perp$, let x be the parameter used to define that family. Let i be the point beyond which x becomes constant. At each node 0, 1, ..., i, at most one A_j is true, and beyond i the only true A_j is the one true at i, if any. So there is at least one (in fact, infinitely many) A_j false at all terminal nodes, and hence also false at $\perp$. Each $\mathrm{MP}^{\vee}_{n}$ fails, because there's no bound on the i by which an x stops changing.

To get a strong separation (still without DC), iterate the previous construction. The Kripke partial order is $\omega^{<\omega}$, a set has to settle down the level after it appears (immediate settling), and considering the transition functions k_n when applied to a set, they're eventually constant.

We do not know whether DC holds in those examples. To be sure of getting a model of DC (weak separation), we have to adapt the earlier topological example. That goes as follows. Partition ω into infinitely many infinite sets. The topology on $\omega \cup \{\bot\}$ is that each $\{k\}$ is open, and $A \cup \{\bot\}$ is a neighborhood of $\{\bot\}$ iff each A_i is in $\mathcal{U}$, where $A_i = \{n \mid \text{the } n^{th} \text{ element of the } i^{th} \text{ slice of } \omega \text{ is in } A\}$. Each $j \in \omega$ induces a function f_j from the sets in the full topological model onto V, by interpreting an open set $\mathcal{O}$ as true if $j \in \mathcal{O}$ and false otherwise. Take the submodel of the full model consisting of those sets x that eventually settle down on slices: there is a j such that $f_k(x)$ is independent of the choice of k, as long as k is in ω's i^{th} slice and $i \geq j$. At this point it is routine to check that IZF + DC + WLEM_ω hold and $\exists n\ MP_n^\vee$ does not.

To get a strong separation with DC, we have to iterate the topological example. This is as in Model 11. ⊣

Regarding the nonimplication from MP to LLPO_ω, it was observed in [23] that a quite standard model (realizability using the Turing computable functions) does it. Still, we think it fitting to provide a model doing the same thing in the style of the other models of this paper.

THEOREM 5.8 ([23]). *Over IZF + DC, MP does not imply* $LLPO_\omega$.

PROOF. For a weak separation, let $\{\omega_i : i \in \omega\}$ be a partition of ω into blocks of size ω and let f_i be a bijection between ω and ω_i for each i. Let $\mathcal{U}$ be a nonprincipal ultrafilter on ω and $\mathcal{U}_i$ the image of $\mathcal{U}$ under f_i $(i \in \omega)$. We take $X = \omega \cup \{*, \infty\}$ and form opens by taking unions of: $\emptyset$, $\{k\}$ $(k \in \omega)$, $\{\infty\}$, and sets of the form $\bigcup\{u_i : u_i \in \mathcal{U}_i\} \cup \{*, \infty\}$.

Take the full model over X. As a full model, it satisfies IZF. By Proposition 2.2, DC holds. Let α be such that $\{k\} \Vdash \alpha(n) = 1$ iff $n = k$, and $\{\infty\} \Vdash \alpha(n) = 0$ for all n. It is easy to see that α is a counter-example to LLPO_ω. As for MP, suppose $* \in \mathcal{O}$ and $\mathcal{O} \Vdash \neg\neg\exists n\ \beta(n) = 1$. Then $\{\infty\} \subseteq \mathcal{O}$, and $\{\infty\} \Vdash \exists n\ \beta(n) = 1$. Let k be such that $\{\infty\} \Vdash \beta(k) = 1$. Let $\mathcal{O}' \subseteq \mathcal{O}$ be a neighborhood of $*$ forcing a value of $\beta(k)$. Since $\infty \in \mathcal{O}'$, $\mathcal{O}'$ must force $\beta(k)$ to be 1, which suffices.

To get a strong separation, this can be iterated in any of a number of ways as above. ⊣

§6. WLPO and WKL. It has been pointed out to us by Kohlenbach that WLPO and WKL are separated by functional Lifschitz realizability [1, 28] and by the monotone functional interpretation [12, 14]. They can also be separated in the style of this paper by using a Kripke model.

THEOREM 6.1. *Over IZF + DC, WKL does not imply WLPO.*

We remark that in all of the models of this theorem MP holds, as is easily verified.

PROOF. **Model 15:** (A weak separation with DC.) The Kripke poset P consists of a root node $\bot$ with two successors 0 and 1. Let $M_\bot$ and M_0 be V, and M_1 be an ultrapower of V with nonstandard integers. Take the full model over that system.

Since the universes at 0, 1 are both classical, DC can be lifted directly to the meta-level: if $\bot \Vdash x_0 \in X \wedge \forall x \in X\ \exists y \in X\ \varphi(x, y)$, then $x_0 \in S$ and $\forall x \in S\ \exists y \in S\ (\bot \Vdash \varphi(x, y))$ where $S = \{x : \bot \Vdash x \in X\}$. Applying DC at the meta-level we construct a function that can be used to construct a name witnessing the internal instance of DC.

To see that WKL holds, suppose that $g \in K_\perp$ is forced to be a decidable, infinite binary tree at $\perp$. Since g is decidable at $\perp$, $g(\perp), g(0)$, and $g(1)$ all have the same standard part. Since $K_1 = M_1$ is classical, there exists in M_1 an infinite branch α of $g(1)$. Let α_0 be the standard part of α. Then $(\{\check{n} \mid n \in \alpha_0\}, \alpha_0, \alpha)$ is an infinite branch of g at $\perp$.

WLPO fails in our Kripke model because there are names that look like the zero sequence at $\perp$ and 0, but have a one at a nonstandard input at 1. Explicitly, let α be a binary sequence in K_1 with the only nonzero terms at nonstandard places. Then $g = (\{\check{0}^n \mid n \in \omega\}, \mathbf{0}, \alpha)$ is a binary sequence at $\perp$ such that $0 \Vdash \forall n\ g(n) = 0$ and $1 \Vdash \exists n\ g(n) = 1$; hence $\perp$ does not force WLPO for g.

Model 16: We now iterate the previous construction, to get a strong separation, albeit without DC. The result looks like Model 5 or 9.

We want each node of our Kripke model to look like $\perp$ from the construction right above, so our poset will be the full binary tree $2^{<\omega}$. For a node σ containing n occurrences of 1, the associated model will be M_n, the n^{th} iterated ultrapower of V. Take the immediate settling model over that system. WKL holds, and WLPO fails, as in the immediately preceding construction.

Model 17: We now present a topological model, a small variant of Model 6, to get DC to hold (in a weak separation).

Let $\mathcal{U}$ be a nonprincipal ultrafilter on the natural numbers. Let X be $\omega \cup \{*, \infty\}$, and topologize X by letting the basic open sets be $\emptyset, \{k\} (k \in \omega), \{\infty\}$, and sets of the form $\{*, \infty\} \cup u$ for $u \in \mathcal{U}$. (So open sets are arbitrary unions of those.) Consider the full topological model M_X.

To see that WLPO fails in M_X, let α be such that $[\![\alpha(k) = 1]\!] = \{k\}$. Since $\infty \Vdash \forall n\ \alpha(n) = 0$ and $k \Vdash \exists n\ \alpha(n) = 1$ for each k, any neighbourhood of $*$ can be extended to both a neighbourhood forcing $\forall n\ \alpha(n) = 0$ and a neighbourhood forcing $\exists n\ \alpha(n) = 1$. Thus $*$ does not force $\forall n\ \alpha(n) = 0 \vee \neg\forall n\ \alpha(n) = 0$.

With regard to LLPO, since each point in $\{\infty\} \cup \omega$ behaves classically, we need only show that $* \Vdash$ LLPO. Let α be a name that, without loss of generality, X forces to be a binary sequence with at most one nonzero term. Let

$$S_E = \{k \mid k \Vdash \forall n\ \alpha(2n) = 0\}$$

and

$$S_O = \{k \mid k \Vdash \forall n\ \alpha(2n+1) = 0\}.$$

Since $S_E \cup S_O = \omega$, either S_E or S_O is in $\mathcal{U}$; without loss of generality, suppose it's S_E. If for some k we had $\infty \Vdash \alpha(2k) = 1$, then no neighbourhood of $*$ could force a value for $\alpha(2k)$. Hence $\infty \Vdash \forall n\ \alpha(2n) = 0$. Because $S_E \cup \{\infty\}$ is dense in $S_E \cup \{\infty, *\}$, for each k the latter set forces $\neg\neg\alpha(2k) = 0$. Because some neighbourhood of $*$ forces a value for each $\alpha(n)$, in particular for $n = 2k$, $S_E \cup \{\infty, *\} \Vdash \alpha(2k) = 0$. Hence $S_E \cup \{\infty, *\} \Vdash \forall n\ \alpha(2n) = 0$.

Because X is dim-zero-dimensional, by Proposition 2, DC holds.

Because both LLPO and DC hold, WKL holds.

Model 18: In order to get a strong separation (with DC), we need to iterate the previous construction. The most straightforward iteration, as in Model 7, does not work, because of the same problem as in Theorem 5.2. In some detail, consider a

simplified, one-step iteration, by which the space X can be thought of as containing the nodes of a tree of height two. On level 0 is just $*$; level one contains the successors of $*$, namely ∞ and k for each $k \in \omega$; the level-two successors of each σ on level one have the form ∞_σ and k_σ. For a set $\mathcal{O}$ to be open, if it contains σ from level one, it must contain ∞_σ, and also $\{k \mid k_\sigma \in \mathcal{O}\} \in \mathcal{U}$. Similarly if $* \in \mathcal{O}$. For better or worse, this model does not satisfy LLPO, as follows. Let α be the binary sequence such that $[\![\alpha(2n) = 1]\!] = \{n_\infty\}$, and $[\![\alpha(2n+1) = 1]\!] = \{n, \infty_n\} \cup \{k_n \mid k \in \omega\}$. In words, the set that forces the odd entries of α to be 0 is the tree beneath ∞, and the set that forces the even entries to be 0 is the tree beneath all the k's; since any neighborhood of $*$ has to contain both, no neighborhood of $*$ can force either.

Still, an iteration for a strong separation is possible, with immediate settling.

This is like Model 11. The space consists of finite sequences from $\{\infty\} \cup \omega$. ($*$ is represented by the empty sequence.) For the topology, if $\mathcal{O}$ is a basic open set and $\sigma \in \mathcal{O}$, then $\sigma^\frown\infty \in \mathcal{O}$, and $\{k \mid \sigma^\frown k \in \mathcal{O}\} \in \mathcal{U}$. The sets at a node have to stop changing after evaluation at any child of the node, as in Theorem 5.2. We leave the verification that this is as desired to the reader. ⊣

§7. Concluding remarks. In work not included here, we have shown that to get Theorem 5.1 the use of an ultrafilter-based space is actually necessary, and, similarly, to get Theorem 5.6, the use of nonultra, filter-based spaces is necessary. We would like to see other such necessity proofs, for the topologies here as well as elsewhere. We would also like to see some converses to these theorems: what general properties of topological spaces could imply that the corresponding models satisfy and falsify certain principles? For instance, Richman [22] shows that the topological model over any metric space satisfies the real-number version of WMP, and Hannes Diener has observed that Richman's proof extends to any first countable, completely regular space. Is Richman's theorem optimal? Is it acutally an equivalence? What would guarantee that binary WMP holds? There's a certain incompatibility between a topological space being metric and being ultrafilter-based; is there a nice theorem here that whether the model over a space satisfies WMP or not corresponds to a nice ultrafilter-vs.-metric style division among spaces?

Since we're talking about the real-number version of WMP, what would a space have to look like to separate binary WMP from real WMP?

It bears observation that the constructions of the last two theorems are incompatible: one produces ultrafilters, the other filters that are not ultra. This is really no surprise, for the principles in the former case—LLPO and not LPO—and in the latter—WMP and not MP—taken together are contradictory, as is easily seen in Figure 1. Could the kind of topology-based proofs used here be extended to show any of the equivalences in Figure 1?

With one exception, we provided models for all of the separations in Figure 1. Even when redundant, we thought it of interest to provide models in the style of this paper. The one exception is the nonimplication from LLPO to WKL. Is there a model of the kind exposited here providing that separation? Or are there provably none?

In the first few models for Theorems 5.1, 5.2, 5.6, 5.7, and 5.8, does DC hold?

We mentioned some work of Kohlenbach and collaborators, who achieve results similar to some of ours. They use realizability over higher-type arithmetic HA^{ω}. Can those constructions be adapted to provide independence proofs over IZF?

There are plenty of other foundational principles than those we have been discussing here: BD-N, continuity axioms, Church's Thesis, Kripke's Schema, the Fan Theorem and its weakenings, to name a few. It would be nice to see them included in an expanded version of the scheme we present here, to get an overview of all their interrelationships.

One in particular that we will mention, because it fits squarely into our scheme, is a weakening of WKL: *weak weak König's lemma* (WWKL) states that if every level of a binary tree T has at least half the nodes of that same level of the full tree, then T has an infinite path. It is easy to see that WKL implies WWKL and WWKL implies LLPO. We conjecture that neither implication reverses. Of course the Lifschitz realizability of [5] separating WKL and LLPO provides one of these separations, we just do not know which. What has been proved [33] is that WWKL is strictly weaker than WKL over RCA_0, a weak (classical) subsystem of second order arithmetic. That proof was then adapted to a constructive setting in [16], but applied to the classical contrapositives of WKL and WWKL, there called D-FAN and W-D-FAN, respectively. It is then shown that the Yu-Simpson argument that WWKL does not imply WKL translates to a proof that over IZF W-D-FAN does not imply D-FAN. In fact, it is not hard to see that both WKL and WWKL fail in this model. So apparently the classical proof separating WKL and WWKL does not translate to such a separation constructively. (Admittedly there could be other adaptations of the Yu-Simpson proof, so perhaps we should say does not so translate yet.) If one is looking to classical constructions for inspiration, one might think to look at LLPO. To be sure, LLPO itself is a classical triviality, but still one could imagine variants, such as with computability or uniformity hypotheses. We are however unaware of any classical study of any correlate of LLPO, and so have no guidance there.

§8. Acknowledgment. We thank Andreas Blass for a careful reading of an earlier draft of this paper, and for the ensuing corrections and suggestions which have served to improve it.

REFERENCES

[1] Y. Akama, S. Berardi, S. Hayashi, and U. Kohlenbach, ***An Arithmetical Hierarchy of the Law of Excluded Middle and Related Principles***, Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science (LICS '04), IEEE Press, New Jersey, 2004, pp. 192–201.

[2] E. A. Bishop, ***Foundations of Constructive Analysis***, McGraw-Hill, New York, 1967.

[3] V. Brattka and G. Gherardi, *Weihrauch degrees, omniscience principles and weak computability*, this Journal, vol. 76 (2011), no. 1, pp. 143–176.

[4] V. Brattka, M. Hendtlass, and A. P. Kreuzer, *On the uniform computational content of computability theory*, http://arxiv.org/abs/1501.00433.

[5] R-M. Chen and M. Rathjen, *Lifschitz realizability for intuitionistic Zermelo-Fraenkel set theory*. ***Archive for Mathematical Logic***, vol. 51 (2012), no. 7, 8, pp. 789–818.

[6] D. van Dalen, *An interpretation of intuitionistic analysis*. ***Annals of Mathematical Logic***, vol. 13 (1978), no. 1, pp. 1–43.

[7] H. M. Friedman, *Set theoretic foundations for constructive analysis*. ***Annals of Mathematics***, vol. 105 (1977), no. 1, pp. 1–28.

[8] H. Friedman and A. Scedrov, *The lack of definable witnesses and provably recursive functions in intuitionistic set theories*. ***Advances in Mathematics***, vol. 57 (1985), pp. 1–13.

[9] R. J. Grayson, *Heyting-valued semantics*, ***Logic Colloquium '82*** (G. Lolli, G. Longo, and A. Marcja, editors), Studies in Logic and the Foundations of Mathematics 112, North Holland, Amsterdam, 1984, pp. 181–208.

[10] H. Ishihara, Markov's principle, Church's thesis and Lindelöf's theorem. ***Indagationes Mathematicae***, vol. 4 (1993), no. 3, pp. 321–325.

[11] H. Ishihara, *Constructive reverse mathematics: compactness properties*, ***From Sets and Types to Topology and Analysis: Towards Practicable Foundations for Constructive Mathematics*** (L. Crosilla and P. Schuster, editors), Oxford University Press, Oxford, 2005, pp. 245–267.

[12] U. Kohlenbach, *Relative constructivity*, this Journal, vol. 63 (1998), pp. 1218–1238.

[13] ———, *On weak Markov's principle*. ***Mathematical Logic Quarterly***, vol. 48 (2002), Suppl. 1, pp. 59–65.

[14] ———, ***Applied Proof Theory: Proof Interpretations and their Use in Mathematics***, Springer, Berlin, 2008.

[15] M. Krol, *A topological model for intuitionistic analysis with Kripke's Scheme*. ***Zeitschrift für mathematische Logik und Grundlagen der Mathematik***, vol. 24 (1978), pp. 427–436.

[16] R. Lubarsky and M. Rathjen, *Realizability models separating various fan theorems*, ***The Nature of Computation*** (P. Bonnizoni, V. Brattka, and B. Löwe, editors), Proceedings of Computability in Europe 2013, LNCS 7921, Springer, Heidelberg, 2013, pp. 306–315.

[17] M. Mandelkern, *Constructively complete finite sets*. ***Zeitschrift für mathematische Logik und Grundlagen der Mathematik***, vol. 34 (1988), pp. 97–103.

[18] U. Mylatz, *Vergleich unstetiger Funktionen: "Principle of Omniscience" und Vollständigkeit in der C-Hierarchie*, Ph.D. dissertation, Fernuniversität Hagen, 2006.

[19] J. Moschovakis, *A topological interpretation of second-order intuitionistic arithmetic*. ***Compositio Mathematica***, vol. 26 (1973), pp. 261–276.

[20] J. van Oosten, ***Realizability***, Studies in Logic and the Foundations of Mathematics 152, Elsevier, Amsterdam, 2008.

[21] M. Rathjen, *Constructive Zermelo-Fraenkel set theory and the limited principle of omniscience*. ***Annals of Pure and Applied Logic***, vol. 165 (2014), pp. 563–572.

[22] F. Richman, *Weak Markov's principle, strong extensionality, and countable choice*, unpublished manuscript, 2000, available at http://math.fau.edu/richman/HTML/DOCS.HTM.

[23] ———, *Polynomials and linear transformations*. ***Linear Algebra and its Applications***, vol. 131 (1990), pp. 131–137.

[24] D. Scott, *Identity and existence in intuitionistic logic*, ***Applications of Sheaves*** (M. Fourman, C. Mulvey, and D. Scott, editors), Lecture Notes in Mathematics 753, Springer, Berlin, 1979, pp. 660–696.

[25] S. Simpson, ***Subsystems of Second Order Arithmetic***, Association for Symbolic Logic/Cambridge University Press, Cambridge, 2009.

[26] M. Toftdal, *Calibration of ineffective theorems of analysis in a constructive context*, Master's thesis, University of Aarhus, 2004; also as *A Calibration of ineffective theorems of analysis in a hierarchy of semi-classical logical principles*, ICALP 2004, LNCS 3142, 2004, pp. 1188–1200.

[27] A. S. Troelstra, ***Metamathematical Investigation of Intuitionistic Arithmetic and Analysis***, Lecture Notes in Mathematics 344, Springer, Berlin, 1973.

[28] ———, *Realizability*, ***Handbook of Proof Theory*** (S. Buss, editor), Studies in Logic and the Foundations of Mathematics 137, Elsevier, Amsterdam, 1998, pp. 407–474.

[29] A. S. Troelstra and D. van Dalen, ***Constructivism in Mathematics***, vol. 1, Studies in Logic and the Foundations of Mathematics 121, North Holland, Amsterdam, 1988.

[30] ———, ***Constructivism in Mathematics***, vol. 2, Studies in Logic and the Foundations of Mathematics 123, North Holland, Amsterdam, 1988.

[31] R. E. Vesley, *Realizing Brouwer's sequences*. ***Annals of Pure and Applied Logic***, vol. 81 (1996), pp. 25–74.

[32] K. Weihrauch, ***The TTE-Interpretation of Three Hierarchies of Omniscience Principles***, Informatik Berichte Nr. 130, Fernuniversität Hagen, Hagen, 1992.

[33] X. Yu and S. G. Simpson, *Measure theory and weak König's lemma*. ***Archive for Mathematical Logic***, vol. 30 (1990), pp. 171–180.

DEPARTMENT OF MATHEMATICS AND STATISTICS
UNIVERSITY OF CANTERBURY
PRIVATE BAG 4800, CHRISTCHURCH, NEW ZEALAND
E-mail: m.hendtlass@gmail.com

DEPARTMENT OF MATHEMATICAL SCIENCES
FLORIDA ATLANTIC UNIVERSITY
BOCA RATON, FL 33431, USA
E-mail: rlubarsk@fau.edu