A new pass at the NF consistency proof

M. Randall Holmes

5/28/2020 8:30 am

0.1 Version notes

5/28/2020: commented out the digression on coding, which breaks the narrative. A complete development along this lines might eventually have use.

5/7/2020: Further notes about how the proof might be carried out using the scheme of infinitary notation.

5/6/2020: proofreading pass. Did discover some misstatements and editing errors. 10:30 pm added a little more to the digression on infinitary notation.

5/5/2020: I do not think the "material difficulty" was real. I believe the paper is correct as it stands (at least in the respect I was worried about). I was making a fencepost error. I think it might be interesting to add as an appendix the material on coding of elements of parent sets which I was working on to fix the supposed problem.

I added a section describing a scheme for coding the parent sets, which I thought I needed to fix the supposed error, but which might have some intrinsic interest. It is clearly marked as not part of the development; I may eventually extend it to a complete approach.

5/1/2020: I am working on a substantial correction to the argument. There is a difficulty with the construction of parent sets described in this version (and possibly in other recent versions) though I was aware of this problem in earlier versions. I believe I have it in hand, but things are provisional just now.

5/1/2020: Editing pass. Do I want to add an appendix describing the FM properties of a single clan?

I added some explicit language about how parent maps are chosen, spelling out in detail something already said in the text, but in mind-numbing detail.

4/30/2020 11:30 am: 8 am Corrected $\pi_{L,M}$ to $\pi_{L,\pi(L)^\circ}$ in a number of places in the discussion of the Freedom of Action theorem.

11 am parenthetical remark added in proof of Freedom of Action. 11:30 minor tweaks.

Adding some language about FM models.

4/29/2020 2 pm: I added technical appendices expanding certain definitions and certain paragraphs in the argument in various sections. This required actual debugging of the argument for cardinalities: the high-level description in one paragraph was wrong, and it (it is to be hoped) corrected in the text and expanded on in section 10.3. I also improved the statement of results about elementarity.

This is now eligible to be a flagship proof document: it has everything in it. I am morally certain that more debugging is needed, though.

4/28/2020 1:30 pm: Fixed misstatement of the definition of local cardinal. Proofreading is needed.

More proofreading. Added references (the same bibliography as previous versions, citations inserted as needed).

At 2:30, slight edit to the conclusion.

4/28/2020 8:40: Fixed up misstatements in the section on elementary equivalence between natural models of TST_{n+2}. Mod silliness which certainly must be present (it is quite hard to port this proof to a new version without dropping the ball on some detail) the argument is complete here. Some of it should be documented in more detail with bullet points, and I will be doing that.

There are things that readers of previous versions should notice. I have omitted any $\mathtt{clan}[\emptyset]$, allowing parent sets to be undefined for clans with singleton index. This causes the definition of parent sets to be uniform without weird exceptions. It also causes the tangled web to be localized to clans with index with a common maximum element, which means that λ has to be a limit ordinal greater than ω: $\omega \cdot 2$ would work. I avoided problems with the set of near-litters being possibly larger than μ by making the cofinality of μ strictly greater than κ, so the minimum possible value of μ is $\beth_{\omega_2}$ rather than $\beth_{\omega_1}$ as in previous versions.

Contents

1 Introductory remarks

This argument has a vexed history. I had the initial idea in 2010, though early versions were certainly formally incorrect. Later versions I believe are correct, but the evidence suggests that they are unreadable. The underlying ideas are not really that monstrous; after a time away from this I am going to try to lay it out again here.

1.1 The issue of preliminaries

There are two prerequisites, some understanding of the model theory of *NF* and *NFU* (we suggest [1] for general introduction to NF, and Jensen's NFU paper [7] and Specker's ambiguity paper [13] for crucial background results) and a basic understanding of Frankel-Mostowski permutation models (a tool originally developed for proofs of independence of the Axiom of Choice from theories with atoms: see [6] for an account we used for reference).

NF disproves the Axiom of Choice: this is an old result of Specker ([12]), and the reason that I was thinking of FM models. The details of this proof, and indeed of any mathematics conducted internally to NF, are irrelevant here: all of our business is conducted in ordinary set theory.

2 Preliminaries on type theory with urelements and NFU

The theory TSTU is a many-sorted first order theory with sorts indexed by the natural numbers and primitive predicates of equality and membership. The empty set $\emptyset$ can be provided as a primitive constant (actually, a suite of constants $\emptyset^{\mathbf{i+1}}$, one of each positive type).

We write $\texttt{type}(\text{"}x\text{"})$ for the sort of the variable x. $x = y$ is well-formed iff $\texttt{type}(\text{"}x\text{"}) = \texttt{type}(\text{"}y\text{"})$. $x \in y$ is well-formed iff $\texttt{type}(\text{"}x\text{"}) + 1 = \texttt{type}(\text{"}y\text{"})$.

TSTU has three axiom schemes (the one involving $\emptyset$ is merely a convenience). All sentences of the given shapes are axioms, for each possible assignment of types to the variables.

the empty set is empty: $(\forall x : x \notin \emptyset^{\texttt{type}(\text{"}x\text{"})+1})$.

weak extensionality: $(\forall xyz : z \in x \rightarrow (x = y \leftrightarrow (\forall u : u \in x \leftrightarrow u \in y)))$: nonempty objects with the same elements are equal.

comprehension: For each formula ϕ in which the variable A is not free, the universal closure of $(\exists A : (\forall x : x \in A \leftrightarrow \phi))$ is an axiom: any condition on objects of sort $\texttt{type}(\text{"}x\text{"})$ defines a set of sort $\texttt{type}(\text{"}A\text{"}) = \texttt{type}(\text{"}x\text{"}) + 1$.

The theory NFU is a single-sorted theory with equality, membership and $\emptyset$ as primitive notions (the theory was originally described by Jensen in [7] without distinguishing $\emptyset$ from other nonempty objects). For convenience, we provide a partial function $\texttt{type}$ from variables to natural numbers, with the preimage of each natural number and the complement of the domain of $\texttt{type}$ all countably infinite sets of variables. We say that an atomic formula $x = y$ is well-typed iff $\texttt{type}(\text{"}x\text{"}) = \texttt{type}(\text{"}y\text{"})$. $x \in y$ is well-typed iff $\texttt{type}(\text{"}x\text{"}) + 1 = \texttt{type}(\text{"}y\text{"})$.

We can then describe the axioms of *NFU* neatly:

the empty set is empty: $(\forall x : x \notin \emptyset)$

weak extensionality: $(\forall xyz : z \in x \rightarrow (x = y \leftrightarrow (\forall u : u \in x \leftrightarrow u \in y)))$: nonempty objects with the same elements are equal.

comprehension: For each formula ϕ in which the variable A is not free, and in which each atomic formula is well-typed (this can be weakened to apply only to atomic formulas in which both variables are bound), the universal closure of $(\exists A : (\forall x : x \in A \leftrightarrow \phi))$ is an axiom. Do recall that *NFU* is in fact unsorted, so versions of the comprehension axioms with atomic formulas not well-typed are readily proved (as for example by renaming bound variables).

Supplement the language of TSTU with a bijection $x \mapsto x^+$ from variables to variables of positive type, with the property that $\texttt{type}("x^+") = \texttt{type}("x") + 1$. For any formula ϕ, define ϕ^+ as the result of replacing each variable in ϕ with its image under this map. It should be evident that if ϕ is a theorem, so is ϕ^+.

Define the Ambiguity Scheme as the scheme asserting each $\phi \leftrightarrow \phi^+$ for ϕ without free variables.

Specker proved in [13], 1962, that NFU is consistent iff TSTU + Ambiguity is consistent. Well, in fact it is anachronistic to say this. Specker was talking about the specific version NF of NFU which Quine proposed in [8], 1937. Jensen proposed NFU in [7], 1969, and observed that Specker's proof applies to NFU as well. We will have no occasion to look under the hood of Specker's proof here: it allows us to confine our attention to TSTU + Ambiguity, and not to reason in or about NFU directly at all. It is further worth noting that the model of NFU obtained will satsify the same sentences as the model of TSTU with distinctions of type dropped.

TST is obtained by strengthening TSTU to assert strong extensionality $(\forall xy : x = y \leftrightarrow (\forall u : u \in x \leftrightarrow u \in y))$. NF is obtained from NFU in the same way. The original form of Specker's 1962 theorem was the assertion that NF is consistent iff TST + Ambiguity is consistent.

We present a version of Jensen's 1969 proof from [7] that TSTU + Ambiguity (and so NFU) is consistent.

Let λ be a limit ordinal (ω would do for our immediate purposes, but larger values of λ can be technically useful). Suppose we have a sequence of sets X_α for each $\alpha < \lambda$ and injective maps $f_{\alpha,\beta}$ for each $\alpha < \beta < \lambda$ from $\mathcal{P}(X_\alpha)$ into X_β. This is clearly a possible situation: for example $X_\alpha = V_\alpha$ and $f_{\alpha,\beta}$ the identity on $V_{\alpha+1}$ would work.

Such a sequence gives us many models of TSTU. We work in ordinary set theory (our working theory is actually ZFA with a set of atoms, but there is no need for that much detail here). For each strictly increasing sequence s

taken from λ (by which we mean a strictly increasing function from ω into λ), we point out the model of TSTU in which type i is implemented as X_{s_i}, equality on type i is implemented by equality on X_{s_i}, and membership of type i objects in type $i+1$ objects is represented thus: $x \in_{s,i} y$ is defined as

$$x \in X_{s_i} \wedge y \in X_{s_{i+1}} \wedge y \in \mathtt{rng}(f_{s_i,s_{i+1}}) \wedge x \in f^{-1}_{s_i,s_{i+1}}(y).$$

The symbol $\emptyset^{i+1}$ can be interpreted as $f_{s_i,s_{i+1}}(\emptyset)$. It is easy to see that this gives a model of TSTU for each s; notice that in most cases there are many urelements, as $X_{s_{i+1}}$ is almost always larger than $\mathcal{P}(X_{s_i})$.

For any formula ϕ of the language of TSTU, let ϕ^s be its translation into the model of TSTU determined by s.

Let Σ be a finite set of formulas of the language of TSTU. Let n be greater than the sort of any variable appearing in any formula of Σ. The set of formulas Σ determines a partition of $[\lambda]^n$ (the set of n element subsets of λ) into no more than $2^{|\Sigma|}$ compartments, by considering for each $A \in [\lambda]^n$ the truth values of formulas ϕ^s for $\phi \in \Sigma$ and the range of the restriction of s to n being A.

Now by Ramsey's theorem there is an infinite homogeneous set H for this partition, which includes the range of a strictly increasing sequence h. The theory of the model determined by the sequence h will satisfy $\phi \leftrightarrow \phi^+$ for each $\phi \in \Sigma$. It then follows that the Ambiguity Scheme is consistent with TSTU by compactness, so NFU is consistent. It follows readily that NFU is consistent with Infinity and with Choice (and does not prove Infinity): if Choice holds in the metatheory, it will hold in the model of TSTU with Ambiguity obtained, and Infinity holds iff X_0 is infinite.

It is quite unclear how to adapt the argument just given to avoid producing many atoms, and it is ominous that NF disproves the Axiom of Choice (Specker, [12], 1953). In any event, while NFU was shown to be consistent, and there have been further investigations of what is consistent with and independent of NFU backed by consideration of actual models, the question of the consistency of NF (which was by unfortunate historical accident described first) has remained open for a long time.

We note that TST is easily modelled. We define a *natural model* of TST as a model in which each type i is implemented by a set X_i, equality and membership are implemented by subsets of the equality and membership relations of the metatheory, and $X_{i+1} = \mathcal{P}(X_i)$ in the metatheory. It is important to observe that the first order theory of a natural model of TST is determined by the cardinality of the set X_0 representing type 0, as two such models with type 0 represented by sets of the same size are clearly isomorphic. This result does not depend on choice. The natural model of TST_n with a given base set X_0 consists of the X_i's for $i < n$. TST_n itself is the restriction of TST to the language with all variables having type less than n.

3 Oh what a tangled web we weave...

Without an undue burden of motivation, we present an approach to constructing a model of TST + Ambiguity and so proving the consistency of NF which we developed in our paper [3], 1995, by a perhaps strained analogy with Jensen's 1969 argument. There is no need to consult that paper: in fact, the results needed are presented much more clearly here.

We work in ZFA with a set of atoms. This is our metatheory for the entire paper. The atoms play no particular role here, but they are essential to the actual construction of the situation described.

Let λ be a limit cardinal. An *extended type index* is defined as a nonempty finite subset of λ. If A is an extended type index, we define A_1 as $A \setminus \{\mathtt{min}(A)\}$. We define A_0 as A and A_{n+1} as $(A_n)_1$ where this is defined.

We define a *tangled web* as a function from extended type indices to cardinals with the following properties:

power types: for each extended type index A with $|A| > 1$, $\tau(A_1) = 2^{\tau(A)}$

elementarity: The first order theory of the natural model of TST_n with base type A is determined by $A \setminus A_n$ (the n smallest elements of A), if $|A| \geq n$.

We demonstrate that the existence of a tangled web of cardinals implies the consistency of TST + Ambiguity and so of NF. The proof should be reminiscent of (our version of) Jensen's proof of the consistency of NFU above.

Let τ be a tangled web of cardinals.

Let Σ be a finite set of formulas of the language of TST. Let n be greater than the sort of any variable appearing in any formula of Σ. The set of formulas Σ determines a partition of $[\lambda]^n$ (the set of n element subsets of λ) into no more than $2^{|\Sigma|}$ compartments, by considering for each $A \in [\lambda]^n$ the truth values of formulas ϕ for $\phi \in \Sigma$ in natural models of TST wth type 0 implemented as a set of size $\tau(B)$ with $B \setminus B_n = A$.

Now by Ramsey's theorem there is an homogeneous set H of size $n + 1$ for this partition. Notice that type 1 of a natural model with base type of size $\tau(H)$ is of size $\tau(H_1)$. Thus the two natural models of TST_n with base types the sets implementing type 0 and type 1 of a natural model of TST with base type $\tau(H)$ have the same truth values for formulas in Σ, by homogenity of H. The natural models of TST with base type of sizes $\tau(H)$

satisfy $\phi \leftrightarrow \phi^+$ for $\phi \in \Sigma$, so the entire Ambiguity Scheme is consistent with TST by compactness, so NF is consistent by Specker's 1962 result.

I obtained this result in 1995 and sat on it for fifteen years, as the possibility of existence of such a pattern of cardinals (clearly not consistent with choice) is not obvious at all.

Fix a limit cardinal λ for the rest of the paper. It could be ω for our exact purposes here, but the greater generality is important to understand the scope of our results. For technical reasons we actually want $\lambda > \omega$, but not by much: $\omega \cdot 2$ will suffice. We will work with extended type indices relative to the fixed λ in our construction, but we will actually define tangled webs relative to limit cardinals $\delta < \lambda$, as will be seen, thus λ cannot be the smallest limit cardinal.

4 We set out to construct a tangled web

We now proceed with a quite complicated construction of a Fraenkel-Mostowski model of ZFA with a set of atoms in which it will be seen that there is a tangled web of cardinals.

We are working initially in ZFCA.

This argument has some parameters. We have fixed a limit ordinal $\lambda > \omega$ already, which could be $\omega \cdot 2$.

We fix an uncountable regular cardinal κ. Sets of size $< \kappa$ we term "small" and other sets we term "large". It can be noted that all small subsets of our FM model will be elements of the FM model (as will be readily seen from the definitions). For our immediate purposes, $\kappa = \omega_1$ would work, but the additional abstraction is useful for further purposes.

The cardinality μ of the set of atoms is strong limit with cofinality greater than the maximum of κ and λ. With the minimum choices of λ and κ, $\mu = \beth_{\omega_2}$ will work.[1]

We describe some truly mysterious structure on the atoms whose purpose really only becomes evident by seeing how the argument works.

The atoms are partitioned into *clans*, each of size μ. There is a clan $\texttt{clan}[A]$ for each extended type index A.[2]

A linear order on the extended type indices will be important to us. The order $\ll$ is the unique order on finite subsets of λ (including the empty set) under which $\emptyset$ is last, if $\texttt{max}(A) < \texttt{max}(B)$ we have $A \ll B$, and if $\texttt{max}(A) = \texttt{max}(B)$ we have $A \ll B$ iff $A \setminus \{\texttt{max}(A)\} \ll B \setminus \{\texttt{max}(B)\}$. This is a well-ordering. Its crucial feature is that all downward extensions of an index appear before it does.

Each of the clans $\texttt{clan}[A]$ is partitioned into *litters* of size κ. Let the partition be denoted by $\Lambda[A]$.

We define a *near-litter* as a subset of a clan with a small symmetric difference from a litter. If N is a near-litter, define N° as the litter with small symmetric difference from N, and define the set of *anomalies* of N as $N \Delta N^\circ$.

[1] In earlier versions I said that the cofinality of μ could be κ, and so the minimal value could be $\beth_{\omega_1}$, but I did not properly appreciate problems with counting "near-litters".

[2] In earlier versions of the argument there was a $\texttt{clan}[\emptyset]$ or many copies of such an additional clan. We think it is not needed; it resulted from extending the "formula for parent sets" introduced below one step farther than necessary.

That the cofinality of μ is strictly greater than κ ensures that there are no more than μ near-litters.

We define the *local cardinal* $[N]$ of any near-litter as the set of all near-litters M with $M^\circ = N^\circ$. Notice that $N^\circ \in [N] = [N^\circ]$. For each extended type index A, define $K[A]$ as the set of all local cardinals of near-litters included in $\mathtt{clan}[A]$.

5 A glimpse of impossible things

I am trying to attack the problem of the opacity of this construction. It seems to me that the best approach is to flatly describe the unbelievable thing which happens in the end of the construction. Nothing in this section may be used in subsequent sections, except the final subsection on the verification of the power type property of the purported tangled web: this is entirely a premonition, and the implementation of everything described here will be given independently in following sections. The reasons that it is difficult should be made clear, and it should also be clearer in the course of the later construction why I want to do certain unlikely things.

We define $\mathcal{P}_*(X)$ as the power set of X in the FM construction which we are describing in this section (and we will continue to use this notation for power sets in FM interpretations).

For each extended type index A with $|A| > 1$, we provide a bijection Π_A whose domain is $K[A]$, the set of local cardinals over $\mathtt{clan}[A]$, and whose range is

$$\mathtt{clan}[A_1] \cup \bigcup_{A\cup\{\alpha\}\ll A} \mathcal{P}_*^2(\mathtt{clan}[A \cup \{\alpha\}]).$$

We call this map the parent map of $\mathtt{clan}[A]$.[3][4] $\Pi_A([N])$ might be called the parent of $[N]$, of N, or of any element of N° in various connections. Its bizarre range we call the parent set of $\mathtt{clan}[A]$. Note that $A \cup \{\alpha\} \ll A$ holds iff all elements of A dominate α.

Now we describe the permutation group which defines the FM interpretation (and yes, we know that it is a prerequisite to describing the ranges of the parent maps!)

A permutation of the set of atoms determines a permutation of the entire universe by the rule $\pi(A) = \pi``A$. We systematically confuse permutations of sets of atoms and the associated class permutations of the universe.

[3]The form of this description of parent sets was originally the truly appalling

$$\mathtt{clan}[A_1] \cup \bigcup_{B \text{ s.d.e. } A} \mathcal{P}_*^{|B|-|A|+1}(\mathtt{clan}[B]),$$

where s.d.e. abbreviates "strictly downward extending". The simplification is due to an observation of Nathan Bowler.

[4]By not providing parent maps for clans with singleton index, we avoid the need for $\mathtt{clan}[\emptyset]$ or variations.

An allowable permutation is a permutation of the atoms whose action on sets fixes each Π_A and each $K[\{\alpha\}]$. Let G be the group of allowable permutations. Notice that if L is a litter with parent $\Pi_A([L]) = p$ and π is an allowable permutation then $\pi(L)$ is a near-litter with parent $\pi(p)$: there may be a small collection of elements of L mapped outside $\pi(L)^\circ$, or non-elements of L mapped into $\pi(L)^\circ$. Also notice that the condition given ensures that an allowable permutation fixes clans and parent sets.

We define a *support order* as a small well-ordering on atoms and near-litters with the property that distinct near-litters in the range of the support order are disjoint. [5]

For each support order S, we define G_S as the set of allowable permutations which fix S (and so fix each element of the domain of S).

We say that a set or atom has support S iff it is fixed by all elements of G_S. We say that a set is symmetric if it has a support.

It should not be hard to see that any symmetric set has a support whose domain consists entirely of atoms and litters and we will often want to consider such supports. The reason that we allow near-litters in domains of support orders is that in general it is good for support orders to be mapped to support orders by allowable permutations, and more particularly this makes it easier to see that this actually is an FM construction.

Our FM model then consists of all hereditarily symmetric sets and atoms. Entirely standard considerations ([6] is a reference for these methods, and the basic results are reviewed in a subsection below) show that this is a model of ZFA (in which choice is clearly false). Atoms, near-litters, local cardinals, and parent maps are elements of the FM model for straightforward reasons.

We give only highlights of results about this situation: the details will be found in the explicit construction below illustrating that this is possible.

A crucial result is that the allowable permutations act quite freely. This seems intuitively plausible, but as will be seen below it requires a lot of care in the construction to get it to work correctly. A *local bijection* is an injective partial function on atoms with the same domain and range, and whose domain has small intersection with each litter. The Freedom of Action Theorem asserts that each local bijection can be extended to an allowable permutation (actually, the theorem is stronger than that, allowing freer action on some parent sets, but we leave that for the detailed development).

[5] If I say "support set", I mean support order; "in a support set" means "in the domain of a support order".

This implies satisfying results about litters and clans. Litters are κ-amorphous in the FM model (they have exactly their small and co-small subsets). The subsets of a clan in the FM interpretation are exactly the sets with small symmetric difference from small or co-small unions of litters included in the clan.

5.1 Verification of the power type property of a purported tangled web, on the "impossible" hypotheses

Define $|X|_*$ as the cardinality of X in the FM interpretation. Since we work in the absence of choice, we define cardinals following Scott in [10].

The result of crucial interest (the aim of our quest) is that if δ, a limit ordinal, belongs to λ, $\tau(A) = |\mathcal{P}_*^2(\mathtt{clan}[A \cup \{\delta\}])|_*$ with δ dominating A, defines a tangled web (on subsets of δ rather than λ).[6]

There are two things to verify to see this. We will be able to give the justification for the power type property of a tanged web here: this may give some hint as to how we came up with the bizarre configuration of parent sets. The elementarity condition holds because the structure consisting of the first $n+2$ iterated power sets of $\mathtt{clan}[A]$ is externally isomorphic to the structure consisting of the first $n+2$ power sets of $\mathtt{clan}[B]$ if $A \setminus A_n = B \setminus B_n$ and A_n and B_n are nonempty ("externally" because the isomorphism is a function in the ambient ZFCA which does not belong to the universe of the FM interpretation); this should not seem any more preposterous than what has already been stated but verifying it requires careful attention to how the system of clans and parent maps is actually constructed.

We want to show that $|\mathcal{P}_*^3(\mathtt{clan}[A])|_* = |\mathcal{P}_*^2(\mathtt{clan}[A_1])|_*$, when $|A| > 2$ and $\mathtt{max}(A) = \delta$: this is the translation of the power type condition for the purported tangled web into our context.

Observe that Π_A is a set of the FM interpretation, a bijection from $K[A]$ to

$$\Pi_A``K[A] = \mathtt{clan}[A_1] \cup \bigcup_{A\cup\{\alpha\} \ll A} \mathcal{P}_*^2(\mathtt{clan}[A \cup \{\alpha\}]).$$

$K[A]$ is a subset of $\mathcal{P}_*^2(\mathtt{clan}[A])$.

So we have $|\Pi_A``K[A]|_* \leq |\mathcal{P}_*^2(\mathtt{clan}[A])|_*$, but in fact we have the stronger condition $|\mathcal{P}(\Pi_A``K[A])|_* \leq |\mathcal{P}_*^2(\mathtt{clan}[A])|_*$, because $K[A]$ is a family of disjoint sets, so unions of subsets of $K[A]$, belonging to the double power set of the clan, correspond one-to-one to subsets of $K[A]$.

Now by examination of subsets of $K[A]$. we have $\mathcal{P}_*(\mathtt{clan}[A_1])$ smaller than $\mathcal{P}_*^2(\mathtt{clan}[A])$ [there is a subset of the parent set $K[A]$ the same size as the

[6]The introduction of δ here is a complication of removing $\mathtt{clan}[\emptyset]$ or many indexed versions of such a clan; it still seems to be a gain.

clan $\mathtt{clan}[A_1]$, and the power set of any subset of the parent set is smaller than the double power set $\mathcal{P}_*^2(\mathtt{clan}[A])]$, so $|\mathcal{P}_*^3(\mathtt{clan}[A])|_* \geq |\mathcal{P}_*^2(\mathtt{clan}[A_1])|_*$, one direction of our desired conclusion.

Observe further that $\mathcal{P}_*^2(\mathtt{clan}[A]) = \mathcal{P}_*^2(\mathtt{clan}[A_1 \cup \{\mathtt{min}(A)\}])$ is smaller than Π_{A_1}“$K[A_1]$ (note that the role of δ here is to ensure that A_1 has at least two elements and so Π_{A_1} exists), whence $\mathcal{P}(\mathcal{P}_*^2(\mathtt{clan}[A]))$ is smaller than $\mathcal{P}_*^2(\mathtt{clan}[A_1])$, which establishes $|\mathcal{P}_*^3(\mathtt{clan}[A])|_* \leq |\mathcal{P}_*^2(\mathtt{clan}[A_1])|_*$, the other direction of the desired conclusion.

6 The actual construction

Now we need to disentangle impossible things.

We need to be able to describe the parent maps, which are essential to defining the allowable permutations, without a prior definition of the FM interpretation and thus of allowable permutations.

We need to be able to show that the double power sets of clans are (externally) of cardinality no greater than μ, to avert cardinality obstructions to the description of parent sets.

We need to arrange the external isomorphisms needed for the elementarity condition for the purported tangled web.

Here we exploit the order we described above. For each extended type index A, we define an A-allowable permutation as a permutation of the atoms which fixes each Π_B for $B \ll A$, while also fixing every $K[C]$ (for *any* extended type index C). We can then define support sets and an indexed FM interpretation just as above, (we recapitulate this below to stick with our resolution that nothing in the previous section is to be used).

For each extended type index A we provide a bijection Π_A whose domain is $K[A]$, the set of local cardinals over $\texttt{clan}[A]$, and whose range is

$$\texttt{clan}[A_1] \cup \bigcup_{A\cup\{\alpha\}\ll A} \mathcal{P}^2_{*A}(\texttt{clan}[A \cup \{\alpha\}]).$$

The notation $\mathcal{P}^2_{*A}(\texttt{clan}[A\cup\{\alpha\}])$ refers to the double power set of $\texttt{clan}[A\cup\{\alpha\}]$ in the sense of the FM interpretation determined by the A-allowable permutations. Its definition depends on Π_B's only for $B \ll A$.

We call this map the parent map of $\texttt{clan}[A]$. $\Pi_A([N])$ might be called the parent of $[N]$, of N, or of any element of N° in various connections. Its bizarre range we call the parent set of $\texttt{clan}[A]$.

Now we describe the permutation group which defines the FM interpretation indexed by A.

A permutation of the set of atoms determines a permutation of the entire universe by the rule $\pi(A) = \pi``A$. We systematically confuse permutations of sets of atoms and the associated class permutations of the universe.

An A-allowable permutation is a permutation of the atoms whose action on sets fixes each Π_B with $B \ll A$ and each $K[C]$. Let G^A be the group of A-allowable permutations. Notice that if L is a litter with parent $\Pi_C([L]) = p$ and π is an A-allowable permutation then $\pi(L)$ is a near-litter with parent

$\pi(p)$: there may be a small collection of elements of L mapped outside $\pi(L)^\circ$, or non-elements of L mapped into $\pi(L)^\circ$. Also notice that the condition given ensures that an A-allowable permutation fixes clans and parent sets.

We define an *A-support order* as a small well-ordering on atoms belonging to $\texttt{clan}[A]$ or clans with index downward extending A and near-litters included in $\texttt{clan}[A]$ or clans with index downward extending A with the property that distinct near-litters in the range of the support order are disjoint. Note that the condition on clan indices here is much stronger than the condition of appearing before A in the order.

For each A-support set S, we define G^A_S as the set of A-allowable permutations which fix S (and so fix each element of S).

We say that a set or atom has A-support S iff it is fixed by all elements of G_S. We say that a set is A-symmetric if it has an A-support.

It should not be hard to see that any A-symmetric set has a A-support consisting entirely of atoms and litters and we will often want to consider such supports. The reason that we allow near-litters in support sets is that in general it is good for support sets to be mapped to support sets by allowable permutations, and more particularly this makes it easier to see that this actually is an FM construction.

Our FM model indexed by A then consists of all hereditarily A-symmetric sets with transitive closure containing no atoms in clans distinct from $\texttt{clan}[A]$ with index not downward extending A. Entirely standard considerations (see [6]) show that this is a model of ZFA (in which choice is clearly false). An appendix to this section details why the construction satisfies the conditions to determine a model of ZFA in this way.

It is useful to observe immediately that atoms and near-litters in suitable clans are elements of any of these FM models, and so are local cardinals and parent maps. The order on the singleton of an atom or near-litter is a support for it; the singleton of an element of a local cardinal is a support for it. It is also useful to observe that any small subset of the domain of one of these FM models is an element of the model (by concatenating support orders, first arranging for all near-litters in the domains of the support orders to be litters, by replacing each near-litter N with N° and the anomalies for N, the atoms in $N \Delta N^\circ$). A parent map is invariant under allowable permutations and hereditarily symmetric by previous considerations.

Further, we can define an allowable permutation as one which is A-allowable for every A, and define the final model as before. There are more details to be added of exactly how Π_A is constructed, but these have no bearing on the applicability of FM methods. We can conveniently view an allowable permutation in the strict sense as an $\emptyset$-allowable permutation, since $\emptyset$ follows all the extended type indices in the order we use on them.

6.1 Applicability of FM methods

We follow [6] in our account of FM methods. We briefly recapitulate what is said there.

Start in ZFA (this could be ZFCA).

Any permutation π of the atoms extends to a class permutation of the entire universe via the rule $\pi(A) = \pi``A$.

Choose a group G of permutations of the atoms and a set Γ of subgroups of G with the following properties:

1. For each atom a, Γ contains the subgroup G_a consisting of the permutations in G which fix a.
2. For any $H \in \Gamma$ and any subgroup K of G, $H \subseteq K \rightarrow K \in \Gamma$.
3. For any $H, K \in \Gamma$, $H \cap K \in \Gamma$.
4. For any $\pi \in G$ and $H \in \Gamma$, $\pi H \pi^{-1} \in \Gamma$.

Γ is a normal filter on the subgroups of G including the stabilizers of the atoms.

An object (set or atom) is Γ-symmetric if for some $H \in \Gamma$, every function in H fixes the object.

The master theorem, which we are citing not proving here, is that the collection of atoms and hereditarily Γ-symmetric sets is a model of ZFA.

In our case, the group G is the collection of (A-)allowable permutations and the set Γ consists of all groups G_S, the set of (A-)allowable permutations fixing the (A-)support S, and all subgroups of G which include a G_S as a subset.

The first condition obviously holds (the order on the singleton of an atom is a support).

The second condition holds because of the way Γ is defined.

The third condition requires slight finesse: if S and T are support orders, first produce orders S' and T' by replacing each near-litter N in the domain with N° and the atoms in $N \Delta N^\circ$. Then define an order $S' + T'$ produced by appending T' to S' and removing all but the earliest occurrence of duplicated items from the order. $G_{S'+T'}$ is included in both G_S and G_T; this is enough to verify that intersections of elements of Γ are in Γ.

The fourth condition holds because $\pi G_S \pi^{-1}$ is $G_{\pi(S)}$: this is why we allow near-litters in supports.

7 Fine control of the choice of parent maps

In order to get the crucial Freedom of Action Theorem, it seems to be necessary to impose some fine control on the choice of parent maps.[7]

A *strong support order* is defined as an order on atoms and near-litters in which each atom is preceded by any near-litter in the strong support which contains it (there might not be a near-litter containing the atom) and the segment preceding any near-litter includes a support for the parent of the near-litter (the parent of a near-litter N included in $\texttt{clan}[A]$ being $\Pi_A([N])$) (if the near-litter has a parent, which it does not in case the index of the clan is a singleton). An A-strong support order is restricted to suitable clans as in the case of an A-support order, and does not need to include a support of the parent of a near-litter whose parent is in $\texttt{clan}[A]$. An expanded definition appears at the end of the section.

An *extended strong support order* is a strong support order with the further property that every near-litter in it is a litter and every atom in it belongs to a near-litter in it. An expanded definition appears at the end of the section.

Now we describe the method of construction of Π_A which gives us the desired fine control.

We assume as an inductive hypothesis that each element of a $\mathcal{P}^2_*(\texttt{clan}[B])$ for $B \ll A$ has a B-extended strong support. We well-order $K[A]$ in some arbitary way with order type μ and similarly well-order the intended range of Π_A. When assigning values of Π_A at $[N] \in K[A]$ which belong to a set $\mathcal{P}^2_{*A}(\texttt{clan}[A \cup \{\alpha\}])$, we ensure that $\Pi_A([N])$ has in each case a A-support with the property that any $[M] \in K[A]$ which contains an element of the strong support precedes $[N]$ in the well-ordering given for $K[A]$. We do note that this also requires the inductive hypothesis that the size of the predetermined set Π_A"$K[A]$ is μ (it is clearly at least μ, but we need to exclude the possibility that it is greater). We can arrange that A-supports chosen are in fact A-extended strong supports, by ensuring that atoms in the support are preceded by litters containing them and inserting missing B-extended strong supports for parents of litters where $B \ll A$ as required.

This will work. Notice that we have μ elements of $\texttt{clan}[A_1]$ on which we do not have to impose special conditions. The order on $K[A]$ should be of

[7]Even finer control will be exerted in a subsequent section, but without compromising the conditions added here.

order type exactly μ. At each even stage, choose the first set in Π_A“$K[A]$ (on which we have supposed an order chosen too) which satisfies the condition or choose an atom if there is no such set; at odd stages always choose an atom. Every set in Π_A“$K[A]$ is eventually chosen because it has a small strong support. Every atom is clearly chosen. The only possible obstruction is the existence of too many elements of some double power set of a clan in the FM interpretation.

We recapitulate the statement of the previous paragraph in more detail. Provide a strict well-ordering $<_1^A$ of $K[A]$ of order type μ. Then provide a strict well-ordering $<_2^A$ of the set already defined which is intended to be Π_A“$K[A]$ which is of order type μ and has the property that if X is the item at position δ in the order $<_2^A$ and $[N]$ is the item at position δ in the order $<_1^A$, then X has an A-support whose domain contains only atoms and litters, in which each litter included in `clan`$[A]$ is M° for an $[M] <_1^A [N]$ and each atom belonging to `clan`$[A]$ belongs to M° for an $[M] <_1^A [N]$. We may for some purpose want to record the choice of such a *designated support* for each parent. This can be done: $<_1^A$ can be chosen arbitrarily along with an arbitrary order $<_3^A$ of type μ on Π_A“$K[A]$, then $<_2^A$ chosen stage by stage, at each odd stage choosing an atom and at each even stage choosing the $<_3^A$-first set satisfying the required condition, or an atom if there is no such set. Every set will eventually be chosen because it has a small support and the cofinality of μ is greater than κ. Then define Π_A so that it maps the item at position δ in $<_1^A$ to the item at position δ in $<_2^A$ for each $\delta < \mu$. This paragraph makes no mention of strong support orders, but notice that any litter in $K[A]$ (and so any atom in `clan`$[A]$) is assigned an extended strong support order in this way: this is constructed by prepending to $[N] \in K[A]$ the A-support respecting the orders $<_1^A$ and $<_2^A$ as described above, then repeating this process for each element of `clan`$[A]$ or $\Lambda[A]$ appearing in this support, inserting the new A-support immediately before the item of which it is the support, then resolving duplications by preserving the earliest item. We could assume without loss of generality that all designated supports are A-extended strong supports.

We have now forced the condition that every atom in `clan`$[A]$ has an A-extended strong support.

We then need to verify that any A-support in which all near-litters are litters can be extended to an A-extended strong support, in order to get extended strong supports for double power sets. We have to correct order so that atoms appear before litters they lie in. Each atom or near-litter has an extended support of appropriate index. Inserting extended supports for each item in the given support which does not have such a support included in the given support gives a support in which all failures are of the same kind: they are all down to litters in some fixed clan with index B downward extending A wth parent in $\texttt{clan}[B_1]$. Each time we make an insertion, we insert an extended support of appropriate index immediately before the item requiring it, then repeated items are eliminated, with the earliest one remaining. We can then proceed through ω steps of making these additional insertions, with the process terminating after finitely many steps at each particular failure. The danger of ending up with a non-well-ordering is avoided because the number of future insertions to be made at any given point is finite and predicted.

7.1 Definition of strong support order

A support order is a small well-ordering $<_S$ of a set S of atoms and near-litters in which distinct near-litters are disjoint. It is said to be an A-support order if each $\texttt{clan}[B]$ which contains or includes an element of the order has B equal to or downward extending A. In the body of the text we will typically use the letter S for a support order rather than its domain, but in appendices of this kind we will use this convention.

For any support order $<_S$ we define S_α as the element of S such that the restriction of $<_S$ to $\{x \in S : x <_S S_\alpha\}$ is of order type α.

An A-support order is said to be strong if the following conditions hold:

1. If S_α is an atom, either there is S_β with $\beta < \alpha$ such that $S_\alpha \in S_\beta$ (this β has to be unique) or S_α belongs to no element of S.

2. If S_β is a near-litter belonging to $\texttt{clan}[B]$ with B strictly downward extending A, then either the parent of S_β is an atom and either equal to S_γ for some $\gamma < \beta$ or not an element of S at all, or there is a subset T of β such that the restriction of $<_S$ to T is a B-strong support for the parent of S_β.

7.2 Definition of extended strong support order

An A-extended strong support order is a strong support order with $<_S$ with three additional properties:

1. Every atom in S is an element of a near-litter in S.

2. Every near-litter in S is a litter.

3. For every near-litter in S with atomic parent, the parent belongs to S as well, unless $|A| > 1$ and the parent is in $\texttt{clan}[A_1]$.

8 The Freedom of Action Theorem

If π is an A-allowable permutation, an *exception* of π is an atom x in a $\texttt{clan}[B]$ with $B = A$ or $B \ll A$ such that $x \in L \in \Lambda[B]$ but either $\pi(x) \notin \pi(L)^\circ$ or $\pi^{-1}(x) \notin \pi^{-1}(L)^\circ$. We know that a permutation has only a small collection of exceptions in each litter.

We define an A-local bijection as an injective map with the same domain and range whose domain contains all of $K[A]$ and has small intersection (empty is a species of small) with each litter in each clan at or before $\texttt{clan}[A]$ in the order, and whose domain contains no other sort of item.

Let π_0 be an A-local bijection. For each pair of litters L, M in a relevant clan, let $\pi_{L,M}$ be a bijection from $L \setminus \texttt{dom}(\pi_0)$ to $M \setminus \texttt{dom}(\pi_0)$. We prove that there is an A-allowable permutation π which extends π_0 and each $\pi_{L,\pi(L)^\circ}$.

Note that the map constructed by the Freedom of Action theorem has no exceptional action on atoms other than that imposed by π_0: if an element of a litter L is mapped out of $\pi(L)^\circ$ by π or out of $\pi^{-1}(L)^\circ$ by π^{-1} then it is in the domain of π_0.

We show how to compute π at any relevant atom by a recursion on an extended strong support of the atom.

If we know how to compute π at a local cardinal in $K[B]$ for $B = A$ or $B \ll A$ then we know how to compute π at each atom in the litter L belonging to the local cardinal, by applying either π_0 or $\pi_{L,\pi(L)^\circ}$: we compute $\pi(L)^\circ$ not by computing $\pi(L)$ but by identifying it as the litter belonging to $\pi([L])$.

This means that we can compute π at any atom in a strong support and at any litter whose parent is an atom, and at any litter belonging to an element of $K[A]$.

There remains the case of computation of π at a litter whose parent is a set X. If we can compute π at this set parent, we can then compute π at each element of the litter and so at the litter.

We observe that the set X belongs to a $K[B]$ for $B \ll A$ and we have a B-extended strong support for X as part of our given support, on which we have already computed values for π. We use the inductive hypothesis that our theorem works for $B \ll A$ to choose a B-allowable permutation π_1 which extends the restriction of π as computed so far to atoms and elements of $K[B]$ appearing in the given support (an element of $K[B]$ "appears in the support" via the litter belonging to it) and extends each relevant $\pi_{L,\pi'(L)^\circ}$ and the restriction of π_0 to relevant litters. We extend π to agree with π' at this litter, which it must, as it agrees with the values already computed for

π at a B-support, which is also an A-support. There is something to check: how do we know that π' agrees with π on litters in the support? Consider the first litter at which it fails to agree: since π and π' agree at all elements of a support of this first litter, they agree at the parent of the litter, and can only disagree by mapping some small number of atoms in the litter differently. But in fact once the value at the parent of the litter is determined, both π and π' compute values in the litter in the same way, by consulting the maps π_0 and $\pi_{L,\pi(L)^\circ}$, so they agree at the litter as well.

This gives a method of computing π along an extended strong support for any atom; one has to argue that this is well-defined in the sense that computation along any extended strong support will give the same value.

Suppose that two different extended strong support orders S and T for an atom x gave different computed values of π by the procedure above. Merge them by putting T after S then eliminating items in T which appear in S. This will be an extended support order, and any computation along it agrees with the computation along S and the computation along T of necessity.

All of this works for allowable permutations in general, with the remark that a local bijection in this case will include all sets $K[\{\alpha\}]$ in its domain (since the clans with singleton indices have no parent sets).

It is important to notice that the proof of the Freedom of Action Theorem for A-allowable permutations does not depend on the construction of Π_A succeeding: this proof goes through if each map Π_B for B downward extending A is constructed successfully, without any reference to the existence of Π_A itself.

9 Agreement of iterated power sets of clans across indexed FM models

In this section we argue that $\mathcal{P}^2_{*^A}(\texttt{clan}[A\cup\{\alpha\}])$ is the same as $\mathcal{P}^2_*(\texttt{clan}[A\cup\{\alpha\}])$. This confirms the validity of the formula for parent sets in the impossible things section.

We argue first that $\mathcal{P}_{*^A}(\texttt{clan}[A\cup\{\alpha\}])$ is the same as $\mathcal{P}_*(\texttt{clan}[A\cup\{\alpha\}])$. We show this by arguing from the Freedom of Action Theorem that every subset of any $\texttt{clan}[A\cup\{\alpha\}]$ in the FM interpretation has small symmetric difference from a small or co-small union of litters, and any such set has an $A\cup\{\alpha\}$-support (which is certainly an A-support) and so is in $\mathcal{P}_{*^A}(\texttt{clan}[A\cup\{\alpha\}])$.

Suppose that X is a subset of a litter L included in $\texttt{clan}[B]$ with both X and $L\setminus X$ large, and that parent maps are defined for all indices before B in the order. Let C be the first item in the order on extended type indices $B=C$ or B downward extends C such that Π_C is not defined (or $\emptyset$ if there is no such C). Choose a C-extended strong support of X and a C-extended strong support of L. Let $a\in X$ and $b\in L\setminus X$ be chosen to be in the range of neither support. Define a local bijection sending a to b, b to a and fixing each atom in either support, and extend this to an allowable permutation with no exceptional actions except at these values, as the Freedom of Action Theorem allows. The permutation will fix each litter in either of the supports as well: the first litter moved will have its parent or local cardinal fixed by the allowable permutaton, so if it fails to be fixed it must have an exception which is moved by the permutation from inside the litter to its outside or vice versa, and in fact each exception other than a or b is fixed, and a,b lie in the same litter. Now the permutation fixes X because it fixes all elements of its support, but moves X because it exchanges a and b. This contradiction shows that each litter has no subsets in the FM interpretation other than its small and co-small subsets, which are clearly sets in the FM interpretation.

We say that a set X cuts a litter L if $L\setminus X$ and $L\cap X$ are both inhabited. Suppose that a set X wth a C-extended strong support S cuts each of a large collection of litters included in $\texttt{clan}[B]$. Choose a litter cut by X which is not in the support of X, and choose an element a of this litter not in X and an element b of this litter which is in X. Neither of these atoms can be in the extended support. Now extend the local bijection swapping a and b and fixing each atom in the strong support of X to a C-allowable permutation with no other exceptional actions (litters in the support are also fixed by essentially the same argument given above, with the different point about a and b that they do not belong to any litter in the support). This permutation will fix X because it fixes all elements of a support of X, and not fix X because it exchanges a and b. Thus we see that a set of the FM interpretation can only cut a small number of litters, and so has small symmetric difference from a union of litters.

Suppose that a set X is a union of a large collection of litters in $\mathtt{clan}[B]$ and there is also a large collection of litters in $\mathtt{clan}[B]$ which are disjoint from X. Choose a C-extended strong support of X. Choose a litter L included in $\mathtt{clan}[B]$ and in X and a litter M included in $\mathtt{clan}[B]$ and disjoint from X, neither being in the extended strong support. Choose $a \in L$ and $b \in M$ and as in each earlier step define a C-allowable permutation extending the local bijection swapping a and b and fixing all atoms in the support (and all litters in the support by a kind of argument already given), and so all elements of the support, with no other exceptional actions. As before, this permutation both fixes and does not fix the set X. This shows that any union of litters in $\mathtt{clan}[B]$ is either the union of a small collection of litters or the relative complement with respect to the clan of such a union.

Thus every subset of $\mathtt{clan}[B]$ in the FM interpretation has small symmetric difference from the union of a small or co-small collection of litters, and so has a support consisting of atoms and litters taken from $\mathtt{clan}[B]$: since this is an B-support this set is also in the FM support indexed by B: the identity of this power set is stable in all the FM interpretations we consider.

Note that this argument does not depend on existence of Π_B, and extends to all models we consider up to the first one where the construction of parent maps fails.

Suppose that all clans with index appearing before A in our order have parent maps defined and that Π_A is defined. Let C be the first index which has A as a downward extension such that Π_C does not exist, or $\emptyset$ if there is no such C.

Let X be a subset of $\mathcal{P}_*(\mathtt{clan}[A \cup \{\alpha\}])$ with C-extended strong support S. Let S_0 be the restriction of S to items eligible to be in an A-support (which will be an A-extended strong support). We argue that S_0 is also an A-support for X, so X already belonged to $\mathcal{P}^2_{*A_1}(\mathtt{clan}[A \cup \{\alpha\}])$. We need to show that any A-allowable permutation π which fixes each element of S_0 fixes X.

Choose any element Y of X, and choose for it a strong A-support by starting with a support consisting of atoms and litters in $\texttt{clan}[A \cup \{\alpha\}]$ (we know this exists because Y is a subset of a clan): these can be supplied with strong $A \cup \{\alpha\}$ supports by inductive hypothesis except in the case where we need to supply the atomic parent in $\texttt{clan}[A]$ of a litter. We define a local bijection agreeing with π on each atom in the support for Y, at each exception of π, and at each atomic element of S_0 and fixing each atom in $S \setminus S_0$.

Let π' extend this local bijection without additional exceptions. The map $\pi' \circ \pi$ fixes each litter in S_0: at a first near-litter moved, it would fix the parent, and so it would have to have an exception mapped into or out of the litter. But π and π' agree at all their exceptions in relevant clans so this cannot happen. Thus π' agrees with π at Y. The map π' fixes all litters in S: at a first litter in S that it moved, there would have to be an exception of π' mapped into or out of the litter, and no element of the domain of the local bijection defining π' is mapped into a litter in S from outside, or vice versa. So π' fixes X.

The C-allowable permutation π' agrees with π at Y and must fix X, from which it follows that π maps Y into X; the same argument shows that π^{-1} must map Y into X as well, so π fixes X as desired.

This establishes that the unlikely description of the parent sets of clans is correct – subject of course to the cardinality issue, which is handled in the next section.

We can prove a more general stability result. For any A with $|A| > n$, the definition of $\mathcal{P}_*^{n+1}(\texttt{clan}[A_n])$ in the interpretation indexed by A_n is the same as in the final interpretation. The strategy is very similar to the previous argument. What we need to show is that if we have $X \in \mathcal{P}_*^{n+1}(\texttt{clan}[A_n])$ with extended strong support S, then the restriction S_0 of S to clans with index at or downward extending A_n is an A_n-support for X, so X exists in the FM interpretation indexed by A_n. We need to show that any A-allowable permutation π which fixes each element of S_0 fixes X.

Choose any element Y of X, and choose for it an extended strong A_{n-1}-support (we know it has such a support by inductive hypothesis or at the basis by the argument just above). We define a local bijection agreeing with π on each atom in the support for Y and at each element of S_0 and fixing each atom in $S \setminus S_0$. This local bijection extends to an allowable permutation which agrees with π at Y and must fix X, from which it follows that π maps Y into X; the same argument shows that π^{-1} maps Y into X, so π fixes X as desired. The allowable permutation extending the local bijection behaves correctly at litters in the supports by the same kind of argument given above, because it fails to have any exceptions which might break this.

This argument assumes the existence of Π_{A_n}.

10 Cardinality of iterated power sets of clans in the FM models

We need to show that each set $\mathcal{P}^2_*(\texttt{clan}[A])$ is of size no more than μ in the sense of the ambient ZFCA (these sets are obviously of size at least μ, as they contain double singletons of all atoms in a clan).

We prove this by an analysis of orbits under allowable permutations.

Let x be an object with strong support (not extended strong support) S. Note that strong supports are sent to strong supports by allowable permutations (this is not true of extended strong supports). For any allowable permutation π, x has strong support $\pi(S)$. We define a map on the orbit of S by $\xi_{x,S}(\pi(S)) = \pi(x)$. This is a definition of a function because S is a support for x. We call the maps $\xi_{x,S}$ *coding maps*.

Every object is in $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$ is an image of a A-support under an A-coding map. There are μ supports in the entire structure (here it is important that the cofinality of μ is greater than κ), so certainly μ supports relevant to the set $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$ we are trying to count. Our aim is to show that there is a collection of coding maps for elements of $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$ of size $< \mu$ whose range covers $\mathcal{P}^2_*(\texttt{clan}[A])$, from which it follows that $\mathcal{P}^2_*(\texttt{clan}[A])$ is of size μ, which is the final piece in the argument that our construction of parent maps succeeds for each clan index.

We indicate a formal description of orbits in A-support sets. An orbit specification for a strong support order S is a well-ordering of the same length as S which contains codes for formal information about each element of S: whether each item is an atom or near-litter; the index of the clan containing or including the item; if the item is an atom, the ordinal position of the earlier near-litter containing it, or an indication that there is no such item; if the item is a near-litter not included in $\texttt{clan}[A]$, an indication of the ordinal position of its atomic parent, if it has one, or an indication of the ordinal positions of the B_1-support for the parent of the item, an element of $\mathcal{P}^2_*(\texttt{clan}[B])$ with B strictly downward extending A and a coding function sending the B_1-support to the parent, taken from a covering set of coding functions of size $< \mu$, assumed to exist as an inductive hypothesis. If we have such covering families, we can ensure that all supports we use in our strong support order are in the domain of the covering family. A more detailed formal description of this appears as an appendix to this section.

We argue that A-supports with the same orbit specification belong to the same orbit under the A-allowable permutations: given two supports with the same orbit specification, recursively construct a A-local bijection sending the one support to the other, which can then be extended by the Freedom of Action Theorem to map the one support to the other. Thus the orbit specifications determine the orbits. The only difficult case in this recursion is the case in which a near-litter in one support is to be mapped to the near-litter in the corresponding position in the other support. It might be necessary to introduce a countable set of new atoms into the domain of the local bijection as images or preimages of anomalous elements of these near-litters, and iterated images and preimages of these new atoms. The only restriction on choosing such atoms is if they belong to a near-litter whose image is known (under the local bijection or its inverse) that their image should belong to the correct image near-litter, or that if they belong to no litter in either support, their assigned images and preimages should not: no new exceptional actions are to be introduced. A more detailed description of this process is given in an appendix to this section.

There are $< \mu$ orbit specifications for A-support sets if there are covering sets of coding functions of cardinality $< \mu$ for each $\mathcal{P}^2_*(\texttt{clan}[B])$ strictly downward extending A. This is because an orbit specification (however we formulate it precisely) is a small wellordering of items built from ordinals less than λ, extended type indices, ordinals $< \kappa$, subsets of ordinals $< \kappa$, and the aforesaid covering functions.

Let X be an element of $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$, with strong A-support S. The idea of the next part of the argument is that we can formally describe a coding function with X as a value using the orbit specification for the support order S and a set of $A \cup \{\alpha\}$-coding functions for elements of X. We do this by choosing a support for each element Y of X which is in the domain of a coding function generating Y and belonging to the covering family we have already constructed, then appending the whole to S, resolving duplications, and eliminating elements not appropriate in $A \cup \{\alpha\}$-supports, then providing the coding function for each Y relative to this support. Details are in an appendix to this section.

Now we can suppose the $A \cup \{\alpha\}$-coding functions taken from a set of size less than μ by inductive hypothesis, and there are $< \mu$ sets of such coding functions because μ is strong limit, and of course there are $< \mu$ possible orbit specifications for supports S. So fewer than μ coding functions are needed to cover $\mathcal{P}^2_*(\texttt{clan}[A])$, which, applied to μ supports, give μ elements of $\mathcal{P}^2_*(\texttt{clan}[A])$.

This completes the argument that there are μ elements of $\mathcal{P}^2_*(\texttt{clan}[A])$, obtained as values of $< \mu$ covering functions at μ supports, completng the argument that our construction of the system of clans and parent sets succeeds.

10.1 Formal definition of orbit specification:

Let $<_S$ be an A-strong support order. Conventions from the definition of strong support above will be used in this subsection.

The orbit specification of $<_S$ is a well-ordering $<_s$, and we define s_α in the domain of $<_s$ so that the order type of the restriction of $<_s$ to $\{x \in \texttt{dom}(<_s) : x <_s s_\alpha\}$ is α.

The orbit specification is uniquely specified by the following conditions:

1. If S_α is an atom in $\texttt{clan}[B]$, s_α is $(1, B, \beta)$, where either S_β contains S_α or $\beta = \kappa$ and S_β does not belong to any element of S.

2. If S_α is a near-litter in $\texttt{clan}[A]$ then $s_\alpha = (2, A)$.

3. If S_α is a near-litter in $\texttt{clan}[B]$ with B downward extending A, and parent a set then $s_\alpha = (3, B, T, \chi)$, where T is a subset of α, the restriction of $<_S$ to T is a B-strong support for the parent of S_α, and $\chi = \chi_{\Pi_B([S_\alpha]), <_S \lceil T}$, the coding function for the parent of S_α using the given strong support. If we are using a covering family of coding functions, we require the support $<_S \lceil T$ to be in the covering family (and we assume that we are using covering families for all B downward extending A).

4. If S_α is a near-litter in $\texttt{clan}[B]$ with B strictly downward extending A, and the parent of S_α is an atom, s_α is $(4, B, \delta)$ where either S_δ is the parent of S_α or $\delta = \kappa$ and the parent does not belong to S.

10.2 Construction of a local bijection from strong supports with the same orbit specification

Let $<_S$ and $<_T$ be A-strong support orders with the same orbit specification. We indicate how to construct an A-local bijection the extension of which under the Freedom of Action theorem will take $<_S$ to $<_T$. The construction is recursive: when we carry out the action determined by S_α and T_α, we presuppose knowledge of the actions indexed by $\beta < \alpha$.

1. If S_α and T_α are atoms, the local bijection maps S_α to T_α.

2. If S_α and T_α are near-litters in $\texttt{clan}[A]$, the local bijection maps $[S_\alpha]$ to $[T_\alpha]$.

3. If S_α and T_α are near-litters in $\texttt{clan}[B]$ with B strictly downward extending A with parents which are atoms not in S or T respectively (the fact that the order specifications are the same ensures that these conditions coordinate) then the local bijection maps the parent of S to the parent of T.

4. If S_α and T_α are near-litters in $\texttt{clan}[B]$ with parents sets (which will be in the range of the same coding function because the orbit specifications are the same) we need to ensure that anomalous elements of S_α and T_α are treated correctly. To do this, we map each element of $S_\alpha \Delta S_\alpha^\circ$ to an element of T_α° not in S or T or otherwise already put in the domain or range of the local bijection under the local bijection and each element of $T_\alpha \Delta T_\alpha^\circ$ in the domain of the local bijection to an element of element of S_α° not in S or T or otherwise already put in the domain or range of the local bijection under the *inverse* of the local bijection.

5. Further, we need to extend the local bijection so that its domain is the same as its range. We do this by assigning a value under the local bijection to each x so far put into its domain without an assigned value under the local bijection with the constraint that the new value be strictly new (not in the domain already) and that if x belongs to a litter in S that the new value be taken from the appropriate litter in T, and that if x belongs to no litter in S the new value belongs to no litter in T, and then for each x put into the range of the local bijection without a preimage assigned, choose a new preimage under the same conditions

with the roles of litters in S and T reversed. The point here is that we complete all orbits under the local bijection without introducing any exceptional actions involving litters in S or T. We should also avoid choosing new items from litters which have atomic parents which belong to S or T except where we have to because such litters belong to the supports themselves. This is where it is important that κ be uncountable. A local bijection will be obtained, and the application of the Freedom of Action theorem will give an A-allowable permutation taking $<_S$ to $<_T$ for standard reasons. We also need to extend it to be a bijection from $K[A]$ to $K[A]$, which presents no difficulties.

10.3 Construction of a coding function for a set

Let X be an element of $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$, with strong A-support S.

For each element Y of X choose an $A \cup \{\alpha\}$-strong support S^0_Y for which we can express Y as $\chi_{Y,S^0_Y}(S^0_Y)$, χ_{Y,S^0_Y} being an element of the covering family of coding functions which we posit as an inductive hypothesis.

For each Y, we augment S^0_Y to S_Y by prepending S, dropping duplicated items from S_Y, then dropping items not appropriate for an $A \cup \{\alpha\}$-support.

We then claim that we can determine a coding function from the orbit specification of the support S and the set Σ of coding functions $\{\chi_{Y,S_Y} : Y \in X\}$. The function χ we define will send $T = \pi(S)$ in the orbit of S under A-allowable permutations to the set of all $\chi_{Y,S_Y}(T_Y)$, where T_Y is in the domain of χ_{Y,S_Y} and agrees with T where S agrees with S_Y. That is, $\chi(T)$ is the set of all $\xi(T')$ where $\xi \in \Sigma$ and T' is in the domain of ξ and has as an initial segment the result of dropping all items from T which are not appropriate in an $A \cup \{\alpha\}$-support (we need to be clear that the definition does not depend on reference to S or to elements of X).

To verify that χ is a coding function (in fact, $\chi = \chi_{X,S}$) we need to verify that $\pi(X) = \chi(T)$.

Each element Y of X has $\pi(Y) \in \chi(T)$, because $\pi(Y) = \chi_{Y,S_Y}(\pi(S_Y))$, and $\pi(S_Y)$ has the correct properties as a support.

Suppose $Z \in \chi(T)$. We have $Z = \chi Y, S_Y(T')$ for some Y and T' extending the truncated version of T. $\pi^{-1}(Z) = \chi_{Y,S_Y}(\pi^{-1}(T'))$. The support $\pi^{-1}(T')$ has an initial segment agreeing with the appropriate truncation of S and the same orbit specification as S_Y, so a permutation fixing S and so fixing X will send $\pi^{-1}(T')$ to S_Y, and so send $\pi^{-1}(Z)$ to $Y \in X$, so in fact $\chi(T) = \pi(X)$ as desired.

Each coding function for an element of $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$ with given strong A-support S is thus determined by the orbit specification of S and a set of coding functions which are not literally taken from a covering family of size $< \mu$ provided by inductive hypothesis, but from a family of coding functions obtained by a uniform transformation of the coding functions in the covering family induced by changing the supports which are their domains by prepending S, and eliminating duplicates and inappropriate items. This family of coding functions is of size $< \mu$ just as the covering family is.

This establishes that there is a covering family of $< \mu$ coding functions for elements of $\mathcal{P}^2_*(\texttt{clan}[A \cup \{\alpha\}])$, determined by the orbit specification of the support S and a subset of a set of $< \mu$ covering functions determined by a covering family previously discussed and the orbit specification of S.

11 External isomorphisms between iterated power sets in the FM models (via even finer control of parent maps)

The existence of external isomorphisms between the natural models of TST_{n+2} with base types $\texttt{clan}[A]$ and $\texttt{clan}[B]$ with $A \setminus A_n = B \setminus B_n$ and A_n, B_n nonempty is arranged by further fine tuning of the way the parent maps Π_A are constructed, causng such isomorphisms to exist between the natural models of TST_{n+2} with base types $\texttt{clan}[A]$ and $\texttt{clan}[A \cup \{\alpha\}]$ where $|A| \geq n$ and α dominates all elements of A.

We do this by defining maps χ_α acting on $\texttt{clan}[A]$'s with α dominating A, and on all sets whose transitive closures contain only appropriate atoms by the rule $\chi_\alpha(X) = \chi_\alpha``X$. These maps belong to the ambient ZFCA: they are certainly not present in any of the FM interpretations.

Define χ_α on each $\texttt{clan}[A]$ with α greater than each element of A as a bijection onto $\texttt{clan}[A \cup \{\alpha\}]$ which sends litters to litters. Define Π_A using the procedure above only when $|A| = 2$; define Π_A when $|A| > 2$ as $\chi_A(\Pi_A)$, with the same results but with finer control. In terms of the very fine description of this process using specific well-orderings, choose $<_1^{A\cup\{\alpha\}} = \chi_\alpha(<_1^A)$ and $<_3^{A\cup\{\alpha\}} = \chi_\alpha(<_3^A)$: then everything proceeds exactly in parallel in corresponding clans.

The resulting map χ_α commutes with equality, membership and parenthood.

It follows that the natural model of TST_{n+2} with base type $\texttt{clan}[A]$ and top type $\mathcal{P}_*^{n+1}(\texttt{clan}[A])$, all elements of all types in which have support in clans downward extending A_n, is isomorphic to the natural models of TST_{n+2} with base type $\texttt{clan}[A \cup \{\alpha\}]$ and top type $\mathcal{P}_*^{n+1}(\texttt{clan}[A \cup \{\alpha\}])$ with α dominating A via χ_α. Repeated application of this result gives isomorphism between the models of TST_{n+2} with base types $\texttt{clan}[A]$ and $\texttt{clan}[B]$ with $A \setminus A_n = B \setminus B_n$. This is sufficient to give the elementary equivalence of these models, which gives the elementary equivalences needed to verify the elementarity conditions for the purported tangled web.

This completes the proof that New Foundations is consistent.

12 Conclusions, extended results, and questions

This is a rather boring resolution of the NF consistency problem.

NF has no locally interesting combinatorial consequences. Any fact about sets of a bounded standard size which holds in ZFCA will continue to hold in models constructed using this strategy with the parameter κ chosen large enough. That the continuum can be well-ordered or that the axiom of dependent choices can hold, for example, can readily be arranged. Any theorem about familiar objects such as real numbers which holds in ZFCA can be relied upon to hold in our models (even if it requires Choice to prove), and any situation which is possible for familiar objects is possible in models of *NF*: for example, the Continuum Hypothesis can be true or false. It cannot be expected that *NF* proves any strictly local result about familiar mathematical objects which is not also a theorem of ZFCA (or even of ZFC).

Questions of consistency with NF of global choice-like statements such as "the universe is linearly ordered" cannot be resolved by the method used here (at least, not without major changes).

NF with strong axioms such as the Axiom of Counting (introduced by Rosser in [9], an admirable textbook based on *NF*), the Axiom of Cantorian Sets (introduced in [2]) or my axioms of Small Ordinals and Large Ordinals (introduced in my [4] which pretends to be a set theory textbook based on *NFU*) can be obtained by choosing λ large enough to have strong partition properties, more or less exactly as I report in my paper [5] on strong axioms of infinity in NFU: the results in that paper are not all mine, and I owe a good deal to Solovay (unpublished conversations and [11]).

That NF has α-models for each standard ordinal α should follow by the same methods Jensen used for NFU in his original paper [7]. No model of NF can contain all countable subsets of its domain; all well-typed combinatorial consequences of closure of a model of TST under taking subsets of size $< \kappa$ will hold in our models, but the application of compactness which gets us from TST + Ambiguity to NF forces the existence of externally countable proper classes, a result which has long been known and which also holds in NFU.

We mention some esoteric problems which our approach solves. The Theory of Negative Types of Hao Wang (TST with all integers as types, proposed in [14]) has ω-models; an ω-model of NF gives an ω-model of TST immediately. This question was open.

In ordinary set theory, the Specker tree of a cardinal is the tree in which the top is the given cardinal, the children of the top node are the preimages of the top under the map ($\kappa \mapsto 2^\kappa$), and the part of the tree below each child is the Specker tree of the child. Forster proved using a result of Sierpinski that the Specker tree of a cardinal must be well-founded (a result which applies in ordinary set theory or in NF(U), with some finesse in the definition of the exponential map in NF(U)). Given Choice, there is a finite bound on the lengths of the branches in any given Specker tree. Of course by the Sierpinski result a Specker tree can be assigned an ordinal rank. The question which was open was whether existence of a Specker tree of infinite rank is consistent. It is known that in NF with the Axiom of Counting the Specker tree of the cardinality of the universe is of infinite rank. Our results show that Specker trees of infinite rank are consistent in ZFA. We are confident that our permutation methods can be adapted to ZFC using forcing in standard ways to show that Specker trees of infinite rank can exist in ZF.

We believe that NF is no stronger than TST + Infinity, which is of the same strength as Zermelo set theory with separation restricted to bounded formulas. Our work here does not show this, as we need enough Replacement for existence of $\beth_{\omega_2}$ at least. We leave it to others to tighten things up and show the minimal strength that we expect holds.

Another question of a very general and amorphous nature which remains is: what do models of NF look like in general? Are all models of NF in some way like the ones we describe, or are there models of quite a different character?

References

[1] Forster, T.E. [1995] Set Theory with a Universal Set, exploring an untyped Universe Second edition. Oxford Logic Guides, Oxford University Press, Clarendon Press, Oxford.

[2] Henson, C.W. [1973a] Type-raising operations in NF. Journal of Symbolic Logic 38 , pp. 59-68.

[3] Holmes, M.R. "The equivalence of NF-style set theories with "tangled" type theories; the construction of omega-models of predicative NF (and more)". *Journal of Symbolic Logic* 60 (1995), pp. 178-189.

[4] Holmes, M. R. [1998] Elementary set theory with a universal set. volume 10 of the Cahiers du Centre de logique, Academia, Louvain-la-Neuve (Belgium), 241 pages, ISBN 2-87209-488-1. See here for an on-line errata slip. By permission of the publishers, a corrected text is published online; an official second edition will appear online eventually.

[5] Holmes, M. R. [2001] Strong Axioms of infinity in NFU. Journal of Symbolic Logic, 66, no. 1, pp. 87-116.
("Errata in 'Strong Axioms of Infinity in NFU' ", JSL, vol. 66, no. 4 (December 2001), p. 1974, reports some errata and provides corrections).

[6] Jech, Thomas, *Set theory*, Academid Press 1978, pp. 199-201.

[7] Jensen, R.B. "On the consistency of a slight(?) modification of Quine's NF". *Synthese* 19 (1969), pp. 250-263.

[8] Quine, W.V., "New Foundations for Mathematical Logic". *American Mathematical Monthly* 44 (1937), pp. 70-80.

[9] Rosser, J. B. [1978] Logic for mathematicians, second edition. Chelsea Publishing.

[10] Scott, Dana, "Definitions by abstraction in axiomatic set theory", *Bull. Amer. Math. Soc.*, vol. 61, p. 442, 1955.

[11] Solovay, R, "The consistency strength of NFUB", preprint on `arXiv.org`, `arXiv:math/9707207 [math.LO]`

[12] Specker, E.P. "The axiom of choice in Quine's new foundations for mathematical logic". *Proceedings of the National Academy of Sciences of the USA* 39 (1953), pp. 972-975.

[13] Specker, E.P. [1962] "Typical ambiguity". *Logic, methodology and philosophy of science*, ed. E. Nagel, Stanford University Press, pp. 116-123.

[14] Wang, H. [1952] Negative types.