2018 notes on NF proof

M. Randall Holmes

1/15/2018 3:15 pm home laptop. Major cleanup of proof of the extension property: improved the order of presentation (and also fixed an indexing error). Fixed an error in the argument re sizes of double power sets of clans (by making it more like the previous version) but the language is now rather messy and needs further cleanup.

The purpose of this document was initially to present an abstract description of a system of clans which has exactly the features I need to establish the extension property. I am hoping that this presentation will make it clearer what is happening. It has expanded into a full account of the construction of a tangled web. This is basically the same construction as in previous documents, though the order of presentation is different and there is one slight tweak (which seems not to have been used (or not very much), though I think it is potentially useful). The notation is quite different, though some translations to earlier notation are provided.

The initial description of the clans here is *not* complete: notice that the clans are indexed by ordinals, not the finite sets of ordinals used in the main argument, and that the character of the parent sets (in particular of the collections of set parents) is left quite vague at first. Further material will make things more concrete, showing how the concrete construction in the argument for NF consistency is an example of this abstract construction.

Fixing some parameters: To begin with, we fix a regular uncountable cardinal κ in the base theory (ZFA with choice). We also fix a limit ordinal λ about which we have little to say yet. Sets of cardinality $< \kappa$ will be called small and all other sets will be called large. Our intention for κ is that all small subsets of any of our FM models also belong to the FM model.

The sequence of clans (and parameters χ and μ): We have a sequence of sets C_α for α below some ordinal χ which we call *clans*. Each C_α is a set of atoms. If $\alpha \neq \beta$, C_α and C_β are disjoint. There are atoms which do not belong to any clan. We assume that all sets C_α are of the same cardinality $\mu > \texttt{max}(\kappa, \lambda)$, where μ is a strong limit cardinal and the cofinality of μ is at least κ.

The parameter χ will eventually be seen to be determined by λ.

Each clan is partitioned into litters: With clan C_α we associate a partition Λ_α of C_α into sets of size κ. The elements of sets Λ_α are called *litters*.

local cardinals introduced: Each element L of Λ_α determines a set which we will call the local cardinal of L and sometimes write $|L|$ where confusion with other notions with similar notation will not arise: $|L|$ is the collection of all subsets of C_α with small symmetric difference from L. We use K_α to denote the set of local cardinals of elements of Λ_α. We refer to elements of $\bigcup K_\alpha$ as *near-litters* included in C_α.

atomic and set parent maps introduced: We provide a (not necessarily strictly) increasing function $\alpha \mapsto \alpha^+$ whose domain and range are included in χ. $\alpha^+ > \alpha$ holds for any α in the domain of this function.

We provide for each $\alpha < \chi$ a set of atoms Π_α of size μ. If α^+ is defined, $C_{\alpha^+} \subseteq \Pi_\alpha$. $\Pi_\alpha \setminus C_{\alpha^+}$ is either empty or of size μ, and in either case does not meet any clan. $\Pi_\alpha \cap \Pi_\beta$ for distinct α, β is C_{α^+} if $\alpha^+ = \beta^+$, and otherwise empty. If α^+ is not defined, Π_α does not meet any clan.

For each $\alpha < \chi$, we will provide a bijection f_α from K_α to Π_α (we call Π_α the parent set of the clan C_α). Further, we will provide a bijection g_α from $\Pi_\alpha \setminus C_{\alpha^+}$ (or the empty set if this is not defined) to a collection of sets Σ_α. We call Σ_α the collection of set parents of (selected) elements of C_α. Note that the set Σ_α is of size μ if it is nonempty. Further stipulations about the natures of the sets Π_α and Σ_α are stated below.

We describe later how the maps f_α and g_α are constructed by a recursion along the order on clans (in the initial pass we leave the exact way in which the range Σ_α of g_α is chosen rather vague, but this will be made precise later).

If $N \in K \in K_\alpha$, we use the notation $|N|$ for K and refer to K as the local cardinal of N (generalizing this usage to near-litters as well as litters). We refer to $f_\alpha(|N|)$ as the (atomic) parent of N and to $g_\alpha(f_\alpha(|N|))$ (if it exists) as the set parent of N. We use the notation N° for the element of Λ_α with small symmetric difference from N. We may refer to elements of $N \Delta N^\circ$ as *anomalies* of N.

extending permutations of the atoms: We follow the convention that any permutation of the atoms is extended to the entire universe of sets by the rule $\pi(A) = \pi``A$.

groups of permutations introduced: The group of permutations we consider is the group G of permutations of all atoms (those in clans and those not in clans) whose action fixes each set K_α (thus mapping each litter included in any clan to a near-litter included in the same clan) and further fixes each f_α and g_α. We define G_β for each $\beta < \chi$ as the larger group of permutations which fix each K_α for $\alpha \leq \beta$ and each f_α and g_α for $\alpha < \beta$.

The permutations in G may be called *allowable permutations*, and those in G_α may be called α-*allowable permutations*.

support sets and symmetry: We call a small set of atoms and near-litters a support set [α-support set] iff each atom in the set belongs to a clan and each pair of distinct near-litters in the set is disjoint [and no element of the set belongs to any C_β or $\bigcup K_\beta$ for $\beta > \alpha$]. We say that an object x has support S [α-support S] iff S is an [α-]support set and any permutation in G [G_α] which moves x also moves some element of S. A set is [α-] hereditarily symmetric iff each element of the transitive closure of its singleton has a[n] [α]-support. The [α-] hereditarily symmetric objects make up a model of ZFA for standard reasons (this is an example of the usual Frankel-Mostowski construction].

It is useful to note that for any near-litter, the [α-]supports of its atomic parent, its set parent [if this exists], and its local cardinal are the same.

We will see below that all objects with [α-]supports have [α-]supports which include only atoms and litters. The reason that we allow near-litters in supports is that we want to preserve the nice condition that if x has [α-]support S and ρ is in G [G_α] then $\rho(x)$ has [α-]support $\rho(S)$:

$\rho(S)$ may of course contain near-litters which are not litters even if S does not.

strong support sets: We call a [γ-]support set S with an associated well-ordering $<_S$ of S a [γ-]strong support set iff it satisfies the condition that every atom in S which belongs to a C_α [$\alpha \leq \gamma$] is preceded in the order by at most one element of $\bigcup K_\alpha$ which contains it as an element (and not followed by any such near-litter), and every element N of a $\bigcup K_\alpha$ [$\alpha < \gamma$] which belongs to S is preceded in the order $<_S$ by all elements of some α-support of $f_\alpha(|N|)$. If an object has a [γ-]strong support set S as a [γ-]support, we say that it has [γ-]strong support S. Note that the defining conditions for a strong support imply that any element x of a [γ-]strong support set (belonging to a C_β or $\bigcup K_\beta$ with $\beta < \gamma$) itself has a β-strong support included in the weak segment (including x) determined by x in $<_S$. We refer to orders $<_S$ as *strong support orders*.

symmetry condition on set parents stated: We already know that Σ_α is α-symmetric with empty support: we further stipulate that each element of Σ_α has a strong α-support. Note that the property of having a strong α-support can be verified without any knowledge of f_γ or g_γ for $\gamma \geq \alpha$ (using information about f_β's and g_β's only for $\beta < \alpha$). Thus we may suppose that Σ_α is constructed in some unspecified way on the basis of knowledge of the f_β's and g_β's for $\beta < \alpha$.

technical order condition on atomic parents: We further provide a well-ordering $<_\alpha$ on Π_α (when g_α is nonempty) and stipulate that every element $g_\alpha(x)$ of Σ_α has an α-strong support S such that for each $N \in S \cap \bigcup K_\alpha$ we have $f_\alpha(|N|) <_\alpha x$.

This condition is enforced by the procedure we use to construct f_α and g_α once the maps f_β and g_β for $\beta < \alpha$ and the collection Σ_α have been constructed: we choose (in a way as yet unspecified) a well-ordering $<_{\alpha,1}$ of K_α, a well-ordering $<_{\alpha,2}$ of C_{α^+} (we choose an order on Π_α if α^+ is not defined), a well-ordering $<_{\alpha,3}$ of Σ_α, and a well-ordering $<_{\alpha,4}$ of the domain of g_α, each of order type μ (or 0 if the set in question is empty).

We describe the procedure we follow if Σ_α is of cardinality μ. At each ordinal step we will place an element of Π_α in that ordinal position in $<_\alpha$ and associate it with its preimage under f_α: if the element of Π_α is an element of the domain of g_α, we will associate it with an element of Σ_α which will be the associated value of g_α. At each even ordinal step (the first step being step 0) we place the first still available element of C_{α^+} at that position and associate it with the first still available element of K_α. At each odd ordinal step, we place the first still available element of the domain of g_α at that position and associate it with the first still available element of K_α, further associated with the first still available element of Σ_α which has an α-strong support S such that every element N of $S \cap \bigcup K_\alpha$ has $f_\alpha(|N|)$ already defined. This process will work as long as the cofinality of μ is at least κ, since supports are small. The interleaving of atomic and set parents ensures that new sets in Σ_α satsifying the restriction are always available at odd stages.

Of course the construction of f_α and g_α is trivially easy if Σ_α is empty: g_α is empty and f_α is determined in the obvious way by the given orders on K_α and C_{α^+} (or an order on Π_α if C_{α^+} is undefined).

existence of strong supports and nice supports: It is a consequence of the conditions stated so far that every object with a support actually has a strong support. In fact, we can show that every object with an [α-]support has a *nice support*, which is a strong support S satisfying the stronger conditions that every atom in S belongs to a near-litter in S, and every near litter in S is a litter, and for any near-litter $N \in S \cap \bigcup K_\beta$, [$\beta < \alpha$] there is a β-strong support for $f_\beta(|N|)$ included in the segment determined by N in $<_S$ which has the property that any M in this support which belongs to $\bigcup K_\beta$ satisfies $f_\beta(M) <_\beta f_\beta(N)$. This is important for merging supports sensibly.

We argue that any [α-]support set S_0 can be extended (in a qualified sense: near-litters N which are not litters are replaced by the litter N° near them and the atoms in their symmetric difference from this litter) to a[n] [α-]nice support set S with associated well-ordering $<_S$ with the appropriate properties. Start by well-ordering S_0 quite arbitrarily; this order will be modified in the course of the argument. Replace each near-litter N with the litter N° immediately followed by the atoms in $N \Delta N^\circ$. For each atomic element x of S_0 as modified, add the litter containing x to S (or use the one already present in S_0 as modified) and place it just before x in the order (or move it if it is already present). For each near-litter element N of S_0 as modified so far (including those added by the previous step) which belongs to a $\bigcup K_\beta$ [with $\beta < \alpha$], add a β-nice support for its local cardinal (which exists by the inductive hypothesis that we have shown the result for all $\beta[< \alpha]$) just before N in the order. If an atom or litter (any near-litters added will be litters) in the support added is thus present in two places in the order, delete the later occurrence. This process will produce an [α-]nice support. Notice that the addition of β-nice supports will not create further obligations related to their own elements.

locally small bijections and the extension property: An [α-]locally small bijection is a bijection from a set of atoms (its field) to the same set of atoms the intersection of whose field with any C_β [for $\beta \leq \alpha$] is small, and the intersection of whose field with the domain of any g_β [for $\beta < \alpha$] is empty, and whose field includes all other atoms [including elements of the domains of g_β's for $\beta \geq \alpha$]. We claim that any [α-]locally small bijection can be extended to an element of G [G_α]. This is called the extension property.

The extension may be chosen to have a further technical property: we say that an element ρ of G [G_α] has x as an *exception* if for some $\beta[\leq \alpha]$, $x \in L \in \Lambda_\beta$, and either $\rho(x) \notin \rho(L)^\circ$ or $\rho^{-1}(x) \notin \rho^{-1}(L)^\circ$. The technical condition is that the extension of an [α-]locally small bijection ρ_0 may be chosen so that it has no exceptions not belonging to the field of ρ_0.

precise statement of the version of the extension property to be proved: We define the [α-]extension property. Let ρ_0 be an [α-]locally small bijection. For each $\beta[\leq \alpha]$ and pair of litters L, M both included in C_β choose a bijection $\rho_{L,M}$ from $L \setminus \mathtt{fld}(\rho_0)$ to $M \setminus \mathtt{fld}(\rho_0)$. Our claim is that there is a uniquely determined [α-]locally small bijection ρ which extends ρ_0 and extends each $\rho_{L,\rho(L)^\circ}$, and which has no exceptions not in the field of ρ_0.

proof of the extension property: We observe first that any atom x in a C_β [$\beta \leq \alpha$] has a [α-]strong support, obtained by appending the litter to which the atom belongs, followed by the atom, to the [α-]strong support of the parent of the litter to which it belongs which we know must exist by conditions stated above. In fact, it has a[n] [α-]nice support (defined above) and we will assume that all supports are nice in this argument. (if $\beta > \alpha$, the atom is already in the domain of ρ_0).

We will compute the value at x by a recursion along order $<_S$ on the [α-]nice support S of x, of which x itself is the last element.

We first indicate how to compute the value of ρ at a litter N belonging to $S \cap \Lambda_\gamma$, on the inductive hypothesis that we have already computed ρ for each object $<_S N$.

[If $\gamma = \alpha$, we simply stipulate that the parent of N is mapped to its image under ρ_0 (in whose domain it lies). We then know the identity of the local cardinal $\rho(|N|)$ and so of the litter $\rho(N)^\circ$, and we can compute the value $\rho(x)$ for each $x \in N$ as either $\rho_0(x)$ or $\rho_{N,\rho(N)^\circ}(x)$, and so we have computed the value of $\rho(N)$.]

If N has parent in C_{γ^+} [, with $\gamma^+ < \alpha$,] then this parent p appears in S before N and we have already computed the value of $\rho(p)$ by inductive hypothesis. We then know the identity of the local cardinal $\rho(|N|)$ and so of the litter $\rho(N)^\circ$, and we can compute the value $\rho(x)$ for each $x \in N$ as either $\rho_0(x)$ or $\rho_{N,\rho(N)^\circ}(x)$, and so we have computed the value of $\rho(N)$.

We use the inductive hypothesis that we already have the δ-extension property [for $\delta < \alpha$,] and so in particular the γ-extension property , in the case where [$\gamma < \alpha$ and] N has set parent.

By the inductive hypothesis, we have computed ρ already at each element of S before N and so at each element of $T = S \cap (C_\gamma \bigcup \Lambda_\gamma)$: this set with the order $<_T$ obtained by restriction of $<_S$ is a γ-strong support for the set parent X of N. We use this information to compute a nonce value for $\rho(X)$. Use the γ-extension property to select a γ-allowable permutation ρ'' extending the restriction of ρ as computed so far to atoms in T and to atomic parents not in T of litters in T and extending each $\rho_{L,M}$ where L, M are included in a Λ_δ with $\delta \leq \gamma$. We claim that all such permutations ρ'' have the same value at X, and we set our nonce value of $\rho(X)$ to this common value of all $\rho''(X)$'s. If there were distinct permutations ρ''_1 and ρ''_2 meeting the conditions above which had different values at X, there would be a first litter M in T at which such permutations could have distinct values. But observe that ρ''_1 and ρ''_2 would be forced to have the same value at the parent of M, because they would have the same values at all elements of a support of the parent (or at the parent itself if it were not in T), and then the value of each of ρ''_1 and ρ''_2 could be computed in the same way at each element of M, using either ρ'_0 or $\rho_{M,\rho''_1(M)^\circ}$. Note further that this unique value for all ρ'''s is also the only possible value for a permutation ρ meeting our final specifications. We now compute the value of ρ at the parent of N as the image under g_γ^{-1} of our nonce value for $\rho(X)$. We thus have computed a value for $\rho(|N|)$ and can then compute a value for $\rho(N)$ as above: we then know the identity of the local cardinal $\rho(|N|)$ and so of the litter $\rho(N)^\circ$, and we can compute the value $\rho(x)$ for each $x \in N$ as either $\rho_0(x)$ or $\rho_{N,\rho(N)^\circ}(x)$, and so we have computed the value of $\rho(N)$.

Now for any atom y in S, including x, we compute the value $\rho(N)$ for the litter N in S containing y and in the course of this computation we have already computed $\rho(y)$ as either $\rho_0(y)$ or $\rho_{N,\rho(N)^\circ}(y)$.

We do need to verify that computation of the value of the extension ρ along two different supports cannot give different values at the same item. Suppose that there is an item x which has two different values of ρ computed as above along different supports. We may assume without loss of generality that x is the first item in a given [α-]nice support S for which distinct computed values are possible. Choose another [α-]nice support T along which the computed value for S is different. Construct an strong support U containing both S and T, with the further proviso that the embedded nice α-supports of x from both S and T are inserted before x (this last appears to use the fact that the supports are nice). It is then clear that the value at x computed along U must be the same as the value at x computed along S and the same as the value at x computed along T, because the data used for both original computations is used in the merged computation, and the computations for earlier items in S must have the same results they had originally.

We have thus shown that for every atom we obtain a unique value for ρ at that atom, which at some point we saw *had* to be the value of ρ at that atom if ρ were to satisfy the desired conditions. Further, it is clear from the construction that the extended ρ is in G or G_α as appropriate, and has no exceptions other than elements of the field of ρ_0.

power sets of litters in the FM interpretation: We claim that the power set of C_α in the FM model determined by G or by any G_β with $\beta \geq \alpha$ is the collection of subsets of C_α with small symmetric difference from small or co-small unions of litters.

Suppose that X is a subset of C_α and L is a litter, and that X has β-nice support S and L has β-nice support T, so that both $L \cap X$ and $L \setminus X$ have a common β-nice support U extending $S \cup T$. Let x be an atom in $(L \cap X) \setminus S$ and let y be an atom in $(L \setminus X) \setminus S$. Let ρ be a β-allowable permutation which fixes each atom in S and exchanges x and y, and has no exceptions other than possibly elements of S, x, and y. We first observe that ρ fixes every litter in U: if this were not the case, there would be a $<_U$-first litter M not fixed, whose parent would be fixed by ρ because ρ would fix all elements of a support thereof. It follows that M would have to contain an exception of ρ, an atom mapped either into M from outside or out of M from inside, and this is impossible, since each possible exception of ρ is either fixed (elements of S) or mapped from an element of L (which would in this case have to be M) to another element of L (x and y). It follows that ρ fixes $L \cap X$ and $L \setminus X$, since it fixes every element of their common support U, and clearly this is not the case.

So in fact every subset of C_α in the FM interpretation has either small or co-small intersection with each litter. This implies that litters (which do have support, their own singletons) are κ-amorphous in the FM interpretation: they have only small and co-small subsets.

Now suppose that X is a subset of C_α for which there is a large collection of litters L such that $L \cap X$ and $L \setminus X$ are both nonempty, and that X has a β-nice support S. Choose a litter L which belongs to this large collection, does not belong to S, and contains no element of S. Choose an atom x in $L \cap X$ and an atom y in $L \setminus X$ and construct ρ which fixes each atom in S and sends x to y and y to x, with no exceptions other than elements of S, x, or y. By the same argument given above, the map ρ fixes every litter in S and so fixes X. But clearly ρ does not fix X.

So in fact every subset of C_α in the FM interpretation either includes or completely excludes each litter in a co-small collection of litters.

Suppose that X is a subset of C_α and there is a large collection of litters which are included in X and a large collection of litters which are disjoint from X (these collections are not necessarily sets in the FM interpretation). Suppose that X has a β-nice support S. Choose a litter L which is a subset of X and not an element of S and contains no element of S. Choose a litter M which is disjoint from X, not an element of S and contains no element of S. Choose an atom x in L and an atom y in M and construct ρ which fixes each atom in S and sends x to y and y to x, with no exceptions other than elements of S, x, or y. By the same argument given above, the map ρ fixes every litter in S and so fixes X. But clearly ρ does not fix X.

We now know that every subset of C_α which belongs to the FM model is either the union of a small set of litters, a small set of near-litters (the small collection of sets $L \cap X$ which are large and not litters), and a small set of atoms (the union of the small collection of nonempty sets $L \setminus X$ which are small) or the union of a large set of litters, a small set of near-litters, and a small set of atoms, and in either case such a set has small symmetric difference from a small or co-small union of litters.

It is further straightforward to see that a subset of C_α which has small symmetric difference from a small or co-small union of litters actually does have a β-support, namely the union of the collection of atoms in the small symmetric difference and either the small collection of litters with large intersection with the subset or the small collection of litters which do not have large intersection with the subset (one of these collections of litters is small, of course).

It is worth noting that no large collection of litters can be a set in the FM interpretation. We leave this as an exercise, since it plays no essential role in our argument.

"convergence" of iterated power sets of litters: We prove that $\mathcal{P}^{n+1}(C_\alpha)$ in the model determined by any G or G_β with $\beta \geq \alpha^{+^n}$ is the same set. (by α^{+^0} we mean α and by $\alpha^{+^{n+1}}$ we mean $(\alpha^{+^n})^+$).

Let X be an element of $\mathcal{P}^{n+1}(C_\alpha)$ be an element of the FM interpretation determined by G or by G_β for some $\beta > \alpha^{+^n}$. Our aim is to show that X has an α^{+^n}-support.

For $n = 0$ we verify this easily. A set of atoms in $\mathcal{P}^{0+1}(C_\alpha)$ has a support made up of atoms in C_α and litters in Λ_α for reasons discussed above.

Each of these litters either has atomic parent in C_{α^+} or set parent in Σ_α with an α-support. This support can obviously be extended to an $\alpha^+ = \alpha^{+^1}$-strong support.

We may then assume that $n > 1$ and that each $Y \in X$ has an $\alpha^{+^{n-1}}$-support. Let S be a β-nice support for X. Let S^- be the restriction of S to C_γ and $\bigcup K_\gamma$ with $\gamma \leq \alpha^{+^n}$ We claim that S^- is the desired α^{+^n}-support of X.

Let ρ be an element of $G_{\alpha^{+^n}}$ which fixes every element of S^-. Let Y be an element of X. Y has a $\alpha^{+^{n-1}}$-nice support T by inductive hypothesis. We construct a β-locally small bijection ρ_0' which sends each atom in T, each exception of ρ belonging to or mapping into an element of T, and each atom in S^- to the same value to which ρ sends it and fixes each atom in $S \setminus S^-$, and which maps any other element y of a litter L in S^- or T to an element of $\rho(L)^\circ$ (extending the map to a locally small bijection may require adding additional atoms to its domain and range, but only in clans with index $\leq \alpha^{+^{n-1}}$). Arguments of sorts already presented show that an extension of ρ_0' to an element ρ' of G_β with no exceptions outside the field of ρ_0' sends each litter in T to its image under ρ and fixes each litter in S [$\rho' \circ \rho^{-1}$ fixes all litters in T or S^- by a familiar sort of argument, and ρ' is seen directly to fix everything in S by the same sort of argument]. As a result, $\rho'(Y) = \rho(Y)$ and $\rho'(X) = X$, from which it follows that $\rho(Y) \in X$, verifying that S^- is indeed an α^{+^n}-support for X as required.

analysis of orbits: We analyze orbits in permutions in G or G_α with an eye to determining the sizes of iterated power sets of clans in the FM interpretation.

An object x with an [α-]strong support S with order $<_S$ can be expressed as a function of $<_S$ in a stereotyped way: $\chi_{x,S}(\rho(<_S)) = \rho(x)$ is the definition we intend, where $\rho \in G$ [G_α]. This does not look like a definition of a function, but if $\rho(<_S) = \rho'(<_S)$ it follows that $\rho' \circ \rho^{-1}$ fixes each element of S, so fixes x, so $\rho(x) = \rho'(x)$. The function $\chi_{x,S}$ is called an [α-]*coding function*: note that the domain of an [α-]coding function is the orbit of an [α-]strong support order under G [G_α], and the range of $\chi_{x,S}$ is the orbit of x under G [G_α].

We claim that orbits in the [α-]strong support orders $<_S$ are precisely determined by stereotyped information presented in the same order:

at each position in $<_S$ occupied by an atom in a clan C_β, provide the triple $(1, \beta, \gamma)$, where the item in the (earlier) position β in the order $<_S$ is a near-litter containing the atom, or else $\gamma = \chi$ and no near-litter in S contains the atom; at each position in $<_S$ occupied by a near-litter in $\bigcup K_\beta$, provide either a triple $(2, \beta, \gamma)$, where the item at the (earlier) position γ is the atom in $C_{\beta+}$ which is the atomic parent of the near-litter, or provide a triple $(3, \beta, f)$, where f is a coding function which, if applied to the largest β-strong support order for the parent of the near-litter embedded in $<_S$, yields the parent, or 0 if $\beta = \alpha$. We refer to the lists of data constructed in this way as *[α-]orbit specifications*. If two support orders $<_S$ and $<_T$ have the same data in this sense, it is straightforward to construct a local bijection between atoms whose extension to an element of G or G_α must send $<_S$ to $<_T$. The local bijection maps atoms in each position in $<_S$ to the atom in the corresponding position in $<_T$. The parent of a near-litter in any position in $<_S$ is mapped to the parent of the near-litter in the same position in $<_T$ by the extension of the local bijection constructed up to that point because the two near-litters are images under the same [α-]coding function of a suborder of $<_S$ and a suborder of $<_T$, where any extension of the local bijection constructed so far to a map in G [G_α] sends the suborder of $<_S$ to the suborder of $<_T$. To get the extension to send the near-litter L in the position in $<_S$ to the near-litter M in the corresponding position in $<_T$ requires that we add some additional values to the [α-]local bijection: each element of $L \setminus L^\circ$ needs to be mapped to an element of M and each element of $L^\circ \setminus L$ needs to be mapped to a non-element of M, and elements of $M^\circ \setminus M$ need to be assigned preimages not in L, while elements of $M \setminus M^\circ$ need to be assigned preimages in L, and for each of these items up to countably many iterated images and preimages need to be assigned to make up a full orbit, with the further proviso that additional atoms introduced in this way which belong to near-litters in S or T will be assigned images or preimages in appropriate corresponding litters in T or S as appropriate. This can be done, and requires the assignment of values at only a small set of additional atoms.

Suppose S is [α-]nice (which we may do wlog; or just suppose that all near-litter elements of S are litters). Suppose further that each element of S has β-support for a $\beta[\leq \alpha]$. Each element y of x has an β-strong support T with an order $<_T$ which is a suborder of the order on an [α-]strong support with $<_S$ as an initial segment, all near-litters in T being litters. In merging $<_S$ and an initial [α-]strong support $<_{T_0}$ for y extending a β-strong support for y (all near-litters in which are litters), start by placing $<_S$ before $<_{T_0}$, and observe that the only corrections needed will be deletion of elements of $<_{T_0}$, since both [α-]supports do not contain non-litter near-litters, followed by deletion of elements of $<_S$ which cannot be in a β-support. We argue that a[n] [α-]coding function for x is determined by the orbit of $<_S$ in G [G_α] and a set C of [α-]coding functions $\chi_{y,<_T}$ for each y in x with the [α-]support $<_T$ chosen as just indicated: the [α-]coding function associated with $<_S$ and C sends a[n] [α-]support $<_U$ in the orbit of $<_S$ to the set of all $\chi_{y,<_T}(<^\beta_V)$ where $\chi_{y,<_T} \in C$, $<_U$ is an initial segment of $<_V$, and $<^\beta_V$, the restriction of $<_V$ to C_γ's and $\bigcup K_\gamma$'s with $\gamma \leq \beta$, is an element of the domain of $\chi_{y,<_T}$, which is the orbit of $<_T$. First of all, this is a coding function: it sends $<_S$ to something x' and it sends $\rho(<_S)$ to $\rho(x')$ for any appropriate ρ. We claim that $x' = x$, so this function is actually $\chi_{x,<_S}$. Clearly any element y of x belongs to x'. An arbitrary element z of x' is of the form $\chi_{y,<_T}(<^\beta_V)$ where $<_V$ has $<_S$ as an initial segment. We can use the extension property (with fiddles at near-litters as above) to construct an [α-]allowable permutation fixing $<_S$ and so x and sending $<_T$ to $<^\beta_V$, so sending y to z, whence $z \in x$.

double power set lemma: It is straightforward to establish as usual that the power set of the parent set of a clan is smaller than or the same size as the double power set of the clan: the point is that there is an invariant bijection between elements of the parent set and local cardinals of litters in the clan, which are elements of the double power set of the clan, and further the local cardinals are pairwise disjoint, so there is an invariant bijection from the power set of the parent set to the collection of unions of sets of local cardinals, which is still a subset of the double power set of the clan. All power sets and cardinalities in this paragraph are to be understood in terms of the FM model. This is the double power set lemma of the usual construction, and it goes in the same way (it depends only on abstract features of clans).

We write this out as an inequality (all concepts being in terms of the FM interpretation using G or any G_β with $\beta > \alpha$): $|\mathcal{P}(\Pi_\alpha)| \leq |\mathcal{P}^2(C_\alpha)|$.

clan indices introduced: Now we will start talking about the limit ordinal λ briefly mentioned above. Finite subsets of λ are called "clan indices" for a reason shortly to be introduced. For any clan index A (which is nonempty) we define A_1 as $A \setminus \{\texttt{min}(A)\}$ and define A_0 as A and A_{n+1} as $(A_n)_1$ where this is defined.

the master order on clan indices introduced: We introduce an order $\leq_w$ on clan indices, with a curious definition. $A \leq_w B$ is defined as holding iff B is empty, if $\texttt{max}(A) < \texttt{max}(B)$ or if $\texttt{max}(A) = \texttt{max}(B)$ and $A \setminus \texttt{max}(A) \leq_w B \setminus \texttt{max}(B)$. To prove that this is a well-ordering is straightforward. It has the interesting property that $A <_w A_{n+1}$ for all A with at least $n+1$ elements: downward extension moves a set earlier in the order.

ordinal indexing of clan indices; χ specified in terms of λ: We define $\iota(A)$ as the order type of the restriction of $\leq_w$ to indices $B <_w A$ and we specify that the order type χ of the order on our clans is $\iota(\emptyset) + \lambda$. The clan indexed by a clan index A is then representable as $C_{\iota(A)}$. We can further define $\iota(A)^+$ as $\iota(A_1)$ when $|A| \geq 2$, and define $\iota(\{\alpha\})$ as $\iota(\emptyset) + \alpha$. On the ordinals $\geq \iota(\emptyset)$, the map $\alpha \mapsto \alpha^+$ is not defined.

we commence our usual construction: We can now present a description of the system of clans underlying our previously described tangled web in terms of this model (with a small technical modification: we could easily present it exactly as in the previous development, but we think the technical modification is conceptually useful [though so far we seem not to have used it]).

We note that the notation $\texttt{clan}[A]$ of previous documents is (roughly) synonymous with $C_{\iota(A)}$, the notation $\texttt{parents}[A]$ is synonymous with $\Pi_{\iota(A)}$ and the notation $\texttt{setparents}[A]$ is synonymous with $\Sigma_{\iota(A)}$. The qualification signalled by the parenthesized "roughly" is described in the next paragraph.

For each $\alpha < \lambda$, $\Pi_{\iota(\emptyset)+\alpha}$ is a set of μ atoms not in any clan. In the previous treatment, the same clan $\texttt{clan}[\emptyset]$ was embedded in every set $\texttt{parents}[\{\alpha\}]$: in this treatment, we use disjoint sets of atoms for each α for this purpose. A suggestive notation for these in the style of previous documents might be $\texttt{parents}[\emptyset_\alpha]$ for $\Pi_{\iota(\emptyset)+\alpha}$ and $\texttt{clan}[\emptyset_\alpha]$ for $C_{\iota(\emptyset)+\alpha}$. The technically convenient effect of this is that the actions of allowable permutations on iterated power sets of clans whose associated clan indices have distinct maximum elements are completely decoupled.

Notice that as in earlier presentations, for each nonempty clan index A with $|A| \geq 2$, $\texttt{parents}[A] = \Pi_{\iota(A)}$ has embedded in it the set $C_{\iota(A)^+} = C_{\iota(A_1)} = \texttt{clan}[A_1]$. The modification is that the clan embedded in $\texttt{parents}[\{\alpha\}]$ is now $\texttt{clan}[\emptyset_\alpha]$, a different clan for each α whose parent set is a collection of atoms not in any clan.

description of set parents: For each nonempty clan index A, $\Sigma_{\iota(A)}$ is the union of all sets $P^2(C_{\iota(A\cup\{\beta\})})$, where $\beta < \texttt{min}(A)$. We explain this notation: $P(C_\gamma)$ is the set of all subsets of C_γ with small symmetric difference from small or co-small unions of litters (which we have already seen is the power set of C_γ in the FM interpretation). $P^2(C_\gamma)$ is the collection of all subsets of $P(C_\gamma)$ with γ^+-strong support, and by results above $P^2(C_\gamma)$ is $\mathcal{P}^2(C_\gamma)$ in the sense of the FM interpretation.

From the definitions it follows that $P^2(C_{\iota(A\cup\{\beta\})})$ is the collection of subsets of $P(C_{\iota(A\cup\{\beta\})})$ with $\iota(A)$-strong support, and is the same set as $\mathcal{P}^2(C_{\iota(A\cup\{\beta\})})$ in the sense of the FM interpretation.

To show that we have successfully defined an instance of our abstract model at this point, what remains is to show that $P^2(C_{\iota(A\cup\{\beta\})})$ is in every case of size μ (in the sense of the ground ZFA). Everything else then falls out from the conditions given for how f and g maps are to be constructed.

double power sets are small: We argue that $P^2(C_\gamma)$ is of size μ for every γ.

Obviously $P^2(C_\alpha)$ has at least μ elements: consider double singletons of atoms.

We use coding functions to demonstrate this. Each element of $P^2(C_\gamma)$ is obtained by applying a coding function to a support order. There are μ support orders. It is sufficient to show that all elements of $P^2(C_\gamma)$ are generated by a set of $< \mu$ coding functions.

A coding function for elements of $P^2(C_\gamma)$ is determined by a specification for an orbit in strong support orders for elements of $P^2(C_\gamma)$ and a set of coding functions for elements of $P(C_\gamma)$.

A coding function for elements of $P(C_\gamma)$ is determined by a specification for an orbit in strong support orders for elements of $P^2(C_\gamma)$ *in which we can ignore all information except that about atoms in* C_γ *and near-litters in* $\bigcup K_\gamma$ and a set of coding functions for atoms, all of which can be taken to be *projections* of support orders. If we disregard all elements of the strong supports for subsets of clans which are not elements of or subsets of that clan, specifications consist of small lists of items which are triples $(1, \alpha, \gamma)$ or $(3, \alpha, 0)$; there are only a small number of such specifications. There are no more than κ projections. Thus there are no more than 2^κ coding functions to consider (it might be just κ, but certainly $2^\kappa < \mu$). Each element y of an element x of $P^2(C_\gamma)$ with x having a given γ^+-strong support S itself has an $\gamma < \gamma^+$ strong support which is a list of elements of C_γ and Λ_γ which may have the elements of these sets belonging to S as an initial segment, from which we can see that the set of coding functions component determining a coding function for elements of $P^2(C_\gamma)$ can be supposed to have all elements taken from a set of $< \mu$ γ-coding functions (not γ^+-coding functions), and so we can restrict ourselves to $< \mu$ such sets for each choice of the orbit specification component (using the fact that μ is strong limit).

The orbit specification component is a small list of items which are triples built from elements of 3, elements of $\chi + 1$, small ordinals, and coding functions for elements of $P^2(C_\delta)$'s for $\delta < \gamma$ (here we use our specific prescription for parent sets). We can assume as an inductive hypothesis that there are $< \mu$ coding functions generating each $P^2(C_\delta)$. It is then evident that there are $< \mu$ such orbit specifications.

This completes the argument that there are exactly μ elements of $P^2(C_\gamma)$.

arranging external isomorphisms: In general terms, we want to ensure that all structure over assorted clans $C_{\iota(A)}$ which we consider, with α an ordinal dominating each A involved, is exactly analogous to structure over the corresponding clans $C_{\iota(A\cup\{\alpha\})}$.

To this end, we provide at each step bijections h_α sending each $C_{\iota(A)}$ to $C_{\iota(A\cup\{\alpha\})}$ [and further extended to other related atoms as we will discuss, and to sets whose transitive closures contain no atoms not in the domain of h_α by the rule $h_\alpha(X) = h_\alpha$“X.]. We require $h_\alpha(\Lambda_{\iota(A)}) =$

$\Lambda_{\iota(A\cup\{\alpha\})}$. When we construct $f_{\iota(A\cup\{\alpha\})}$, we prescribe that $f_{\iota(A\cup\{\alpha\})}(h_\alpha(x)) = h_\alpha(f_{\iota(A)}(x))$ and $g_{\iota(A\cup\{\alpha\})}(h_\alpha(x)) = h_\alpha(g_{\iota(A)}(x))$ for all appropriate x. This works at each level on the assumption that suitable isomorphism of structure already held at all earlier stages of the construction. In particular, all structure to do with allowable permutations commutes correctly with these external isomorphisms.

The definition of the h_α's is actually made entirely precise by the assertion that $<_{\iota(A\cup\{\alpha\}),k} = h_\alpha(<_{\iota(A),k})$ for each A and α dominating A and $k = 1, 2, 3, 4$: this ensures that the construction of the f and g maps proceeds correctly in order for the equations above to hold, and an inductive hypothesis tells us that the sets $\Sigma_{\iota(A)}$ and $\Sigma_{\iota(A\cup\{\alpha\})}$ are isomorphic in structure in the required sense at each step, the appropriate restriction of h_α already constructed being the isomorphism (they are constructed as unions of double power sets of clans which will be isomorphic if each previous step has worked correctly). The restriction of h_α sending $C_{\iota(\emptyset)+\mathtt{max}(A)}$ to $C_{\iota(\{\alpha\})}$ can be chosen quite arbitrarily (apart from mapping litters to litters) since the former set lacks interesting structure.

The equation $g_{\iota(A\cup\{\alpha\})}(h_\alpha(x)) = h_\alpha(g_{\iota(A)}(x))$ requires that the transitive closure $g_{\iota(A)}(x)$ contain only atoms in the domain of h_α, that is, atoms in clans with indices dominated by α, and our stipulations about parent sets do ensure this. An element of $C_{\iota(A)}$ will have a strong support containing atoms and clans only in $C_{\iota(B)}$'s or $\bigcup K_{\iota(B)}$'s with $\mathtt{max}(B) = \mathtt{max}(A)$ or in $C_{\iota(\emptyset)+\mathtt{max}(B)}$ or $\bigcup K_{\iota(\emptyset)+\mathtt{max}(B)}$. This can be verified by induction on the ordinal indices of clans using our specification of the parent sets. This means that the actions of [α-]allowable permutations on clans and iterated power sets of clans with associated clan indices with different maxima are entirely independent of each other.

the tangled web: Finally the demonstration that the map τ sending each nonempty clan index A to the cardinal $|\mathcal{P}^2(C_{\iota(A)})|$ in the sense of the FM interpretation is a tangled web goes just as in the previous document.

What needs to be shown is that

1. $2^{\tau(A)} = \tau(A_1)$ for $|A| > 2$.

2. The first-order theory of the natural model of TST_n with base type of size $\tau(A)$ $[|A| \geq n]$ depends only on the smallest n elements of A.

$2^{\tau(A)} = |\mathcal{P}^3(C_{\iota(A)})| \geq |\mathcal{P}^2(\Pi_{\iota(A)})| \geq |\mathcal{P}^2(C_{\iota(A_1)})| = \tau(A_1)$

$\tau(A_1) = |\mathcal{P}^2(C_{\iota(A_1)})| \geq |\mathcal{P}(\Pi_{\iota(A_1)})| \geq |\mathcal{P}(\mathcal{P}^2(C_{\iota(A)}))| = |\mathcal{P}^3(C_{\iota(A)})| = 2^{\tau(A)}$

and the first point is established. The crucial points are uses of the double power set lemma $|\mathcal{P}(\Pi_\alpha)| \leq |\mathcal{P}^2(C_\alpha)|$ and the fact that $\mathcal{P}^2(C_{\iota(A)}) \subseteq \Pi_{\iota(A_1)}$.

The model of TST_n with base type of size $\tau(A) = |\mathcal{P}^2(A)|$ has top type of size $|\mathcal{P}^{n+1}(A)|$. This set is completely determined in the FM model using $G_{\iota(A)^{+n}} = G_{\iota(A_n)}$. By application of the external isomorphisms h_β successively, this model of TST_n is externally (in a way not visible to the FM interpretation) isomorphic to the model of TST_n with base type of size $\tau(A \setminus A_n) = |\mathcal{P}^2(A \setminus A_n)|$, and by reversing this process isomorphic to the model of TST_n with base type of size $\tau(B) = |\mathcal{P}^2(B)|$ for any B for which $B \setminus B_n = A \setminus A_n$ (that is, any B with the same first n elements). Thus these theories have the same first-order theory, noting that the FM interpretation has the same natural numbers as the original interpretation of ZFA.