Summer 2018 Notes

Randall Holmes

6/29/2018

I'm trying to work on providing a framework to understand my thinking... In this version I have rather unexpectedly resolved a major problem with the argument: I have removed all need for the recursive arguments in the "wrong direction" which proceed from litters with atomic parents to those atomic parents, which can be made to work but led to very unpleasant features in previous versions. A modification to the notion of local bijection removes all need for such recursions and vastly simplifies matters.

Some metalevel remarks.

1. This version contains a much simpler account of the basic clan machinery. There is some indication in "possibly useful remarks" of a more abstract approach to clans, which probably allows, for example, the extension property to be proved independently of current details of the specific system of clans used. I'm not sure a grand abstract treatment of clans is needed – in my view all the machinery here is a one-off to achieve a particular desired result. But perhaps it might be useful for constructing patterns of cardinals in the exponential function in other choice-free contexts?

2. There is an account of what the higher iterated power sets of a clan are like implicit in the argument that none of these sets are very large. It is worth noting in this connection that what happens with the *double* power set of a clan is most of what happens – and this is already quite complicated.

3. I believe that the argument here has a definite arc: there is a series of lemmas that must be handled, and it is reasonably clear why each one is needed.

4. This version may have lowered the required consistency strength: it no longer appears that μ needs cofinality at least κ, which was an artifact of recursions of the kind here avoided.

5. (agenda) Give an explicit description of how to build a system of just two clans C, D with $|\mathcal{P}^2(C)|_* = |\mathcal{P}^3(D)|_*$. This might be quite enlightening.

where we are working: We are just doing mathematics – the ambient theory is ZFA with Choice, if it must be made concrete. We do want atoms. I will attempt to provide different notation for notions native to FM interpretations, all notation really referring to things in the ambient ZFA. We are at no point actually doing set theory in NF.

our aims: We aim to construct a tangled web of cardinals. We briefly describe what this is. Fix a limit ordinal λ (λ could be taken to be ω).

We refer to nonempty subsets of λ as *extended type indices*.

For any extended type index A with $|A| \geq 2$, we define A_1 as $A \setminus \{\mathtt{min}(A)\}$. We define A_0 as A and A_{n+1} as $(A_n)_1$ when this is defined.

A tangled web on λ is a function τ from nonempty finite subsets of λ to cardinals such that

1. $2^{\tau(A)} = \tau(A_1)$ when $|A| \geq 2$.
2. the first order theory of a model of the simple typed theory of sets with n types with type 0 having cardinality $\tau(A)$ ($|A| \geq n+1$) and type $i+1$ being the power set of type i depends only on $A \setminus A_{n+1}$, the smallest $n+1$ elements of A (the smallest n elements would be more natural and may actually work in this present construction, but $n+1$ works).

We know that existence of a tangled web in a model of ZF entails consistency of NF. We know that existence of a tangled web contradicts AC, so we need to do something to kill choice (FM constructions suggest themselves). We know that existence of a tangled web is wildly unnatural, so we do not expect the construction to be "natural": this very strange situation must be carefully arranged, and we forthwith describe machinery which can arrange it.

smallness: Choose an uncountable regular cardinal κ (κ could be taken to be ω_1). Sets of cardinality $< \kappa$ are to be called *small*: other subsets are to be called *large*. [The intention is that the world of our FM interpretation will contain all of its small subsets].

the size of clans: Choose a strong limit cardinal μ of cofinality $\geq \mathtt{max}(\kappa, \lambda)$. I preserve the usual statement here but I am not sure the cofinality condition is needed in the latest version: it may be sufficient that μ is strong limit and greater than κ and λ, so it could take the desirable value $\beth_\omega$.

clans: We provide an initially indeterminate number of pairwise disjoint sets of atoms of size μ which we call *clans*. Atoms in clans are called *regular atoms*. There might be some other atoms, which we would call *irregular atoms*.

We further provide a strict well-ordering $<_c$ on clans.

litters: Each clan C is provided with a partition Λ_C into sets of size κ: the elements of sets Λ_C are called *litters*. Note that Λ_C is of size μ.

local cardinals: The *local cardinal* $[L]$ of a litter L in Λ_C is defined as the collection of all subsets of C of size κ with small symmetric difference from L. The set of local cardinals of litters in Λ_C is called K_C.

near-litters: The elements of $\bigcup(K_C)$ are called near-litters, and if $N \in K \in K_C$ we write $[N] = K$ and refer to K as the local cardinal of N.

We allow N° to denote the litter with small symmetric difference from a near-litter N.

parent sets: Each clan C is associated with a *parent set* $\Pi(C)$. The set $\Pi(C)$ is of size μ: it may contain both sets and atoms, and its precise nature is unspecified for the moment. A bijection π_C is provided from K_C to $\Pi(C)$. We may refer to $\pi_C([L])$ as the *parent* of a litter L.

We add the possibly useful remark (which might figure in a more abstract treatment) that the set $\Pi_c(C)$ consisting of all atoms in $\Pi(C)$ is either a clan or a set of irregular atoms of size μ). Some $\Pi_c^n(C)$ is a set of irregular atoms for each C. This will hold in the specific structure we are building.

general considerations about permutations: A permutation of the atoms is extended to a permutation of the universe by the rule $\pi(A) = \pi``A$ as usual in the theory of FM constructions.

allowable permutations: A permutation of the atoms is an *allowable permutation* iff it fixes each map π_C.

This requires some unfolding to see that it does everything we want. If ρ is an allowable permutation, it fixes each π_C. This means that it fixes K_C, the domain of π_C, so it maps each $[L]$ for $L \in \Lambda_C$ to some $[M]$ for $M \in \Lambda_C$ (and obviously it maps C to C). This means that it maps each litter L to a near-litter belonging to $\rho([L])$, that is, a subset of C with small symmetric difference from $\rho(L)^\circ$.

So this definition ensures that allowable permutations fix clans and map litters to near-litters as we would expect from earlier versions of our proof. There is further information involved in the correlation between $\Pi(C)$ and K_C which is imposed by fixing π_C: notably, each parent set $\Pi(C)$ is fixed by each allowable permutation.

A permutation of the atoms is a C-allowable permutation for a clan C iff it fixes each map π_D for $D <_c C$ and fixes K_C.

supports: A support is a small well-ordering of a set of regular atoms and near-litters. An object x is said to have support S iff S is a support and any allowable permutation which fixes S also fixes x. An object is said to be symmetric iff it has a support.

Note that a permutation which fixes a well-ordering fixes each element of its domain individually. We will make other uses of the fact that for us a support carries a well-ordering.

Note that it is evident that any object with a support has a support consisting only of atoms and litters. For some purposes, supports of this kind are convenient. We do want near-litters in supports because we want the image of a support under an allowable permutation to be a support.

A C-support set and C-symmetry [for a specific clan C] are defined in the same way with the restriction that all elements of C-support sets are elements or subsets of C or of some $D <_c C$, and a set X has C-support S iff S is a C-support and all C-allowable permutations fixing S also fix X.

We add the possibly useful remark that we may expect that all elements of $\Pi(C) \setminus \Pi_c(C)$ (all sets in the parent set of C) are C-symmetric, and that there are μ of them if there are any. This holds in the specific system of clans we build, but might be a useful remark in a more abstract treatment.

completely abstract FM interpretation: The hereditarily symmetric sets make up a model of ZFA by quite standard considerations. We cannot say too much about it without specific information about the maps π_C (though certain basic results probably can be proved just from the "possibly useful remarks" above which we prove for our specific clan structure directly from the details).

We introduce notation $\mathcal{P}_*(X)$ for the collection of hereditarily symmetric subsets of a hereditarily symmetric set X: this is the power set operation of the FM interpretation. Similarly, we introduce notation $|X|_*$ for the collection of all Y of minimal rank such that there is a hereditarily symmetric bijection from X to Y (the Scott cardinal of X in the FM interpretation; of course this only makes sense if X is hereditarily symmetric). We will allow $\leq$ or $\geq$ without adornment to represent the natural partial order with the appropriate sense on Scott cardinals in the FM interpretation: $|X|_* \leq |Y|_*$ will hold iff there is a hereditarily symmetric bijection from X to a subset of Y.

One thing we *can* say is that the FM interpretation has C as a set, Λ_C as a set, K_C as a set, $\Pi(C)$ as a set and π_C as a set for each clan C. Further, we have $|\Pi(C)|_* = |K_C|_* \leq |\mathcal{P}^2_*(C)|_*$ for each clan C. in fact we have a stronger condition. The sets in K_C are pairwise disjoint, so the map $(B \in \mathcal{P}^*(K_C) \mapsto \bigcup B)$ is a bijection from $\mathcal{P}_*(K_C)$ into $\mathcal{P}^2_*(C)$ which is invariant under allowable permutations, so actually $|\mathcal{P}_*(\Pi(C))|_* = |\mathcal{P}_*(K_C)|_* \leq |\mathcal{P}^2_*(C)|_*$ holds.

In the domain of possibly useful remarks, if we subscribe to the properties of the operation Π_c stated above, we have $|\mathcal{P}_*(\Pi(C))|_* = |\mathcal{P}_*(K_C)|_* \leq |\mathcal{P}^2_*(C)|_*$, so we have $|\mathcal{P}^2_*(\Pi(C))|_* = |\mathcal{P}^2_*(K_C)|_* \leq |\mathcal{P}^3_*(C)|_*$, so we have $|\mathcal{P}_*(\Pi(\Pi_c(C)))|_* \leq |\mathcal{P}^2_*(\Pi_c(C))|_* \leq |\mathcal{P}^2_*(\Pi(C))|_* = |\mathcal{P}^2_*(K_C)|_* \leq |\mathcal{P}^3_*(C)|_*$ if $\Pi_c(C)$ is a clan, and more generally $|\mathcal{P}_*(\Pi(\Pi^n_c(C)))|_* \leq |\mathcal{P}^{n+2}_*(C)|_*$ if $\Pi^n_c(C)$ is a clan. If structure is put into clans nested in parent sets to a certain number of iterations, it pops out in a form visible in the FM interpretation after a related number of iterations of the power set operation. Explicit thoughts along these lines play a role in the intellectual history of this argument, and related results appear in the development of our specific system of clans, though not in this format. A related key idea is that set structure hidden in nested parent sets does *not* appear in any way detectable in the FM interpretation in iterated power sets of the clan of lower index: this of course requires justification (and a result of this kind is proved, though not with this terminology, in the existing argument).

Notice for example that were I to manage to embed $\mathcal{P}_*(C)$ into $\Pi(D)$ and $\mathcal{P}_*(D)$ into $\Pi(C)$, I would have coerced the equation $|\mathcal{P}^2_*(C)|_* = |\mathcal{P}^2_*(D)|_*$. And, in fact, building a concrete model with just two clans C, D in which these relations obtain is not particularly difficult. Similarly, if I want $|\mathcal{P}^3_*(C)|_* = |\mathcal{P}^5_*(D)|_*$, a way to do this would be to have $\mathcal{P}^4_*(D)$ a subset of $\Pi(\Pi_c(C))$ and $\mathcal{P}^2_*(C)$ a subset of $\Pi(\Pi^3_c(D))$. This is part of the idea behind the structure of parent sets in our specific system of clans.

We note that there is no reason whatsoever to think that the fact that all the sets C, $\Pi(C)$ etc. are of the same size μ is preserved in the passage to the FM interpretation, and in fact we are counting on this not to be the case.

Similarly, the hereditarily D-symmetric sets make up a model of ZFA, in which the remarks just above hold with limitations on what clans are considered. The facts that C as a set, Λ_C as a set, K_C as a set hold for $C \leq_c D$. π_C is a set, $\Pi(C)$ is a set and $|\mathcal{P}_*(\Pi(C))|_* = |\mathcal{P}_*(K_C)|_* \leq |\mathcal{P}^2_*(C)|_*$ hold for $C <_c D$.

we show a little more of our hand: Our aim is to define for each extended type index A a clan $\texttt{clan}[A]$ with associated structure in such a way that $\tau(A) = |\mathcal{P}^2_*(\texttt{clan}[A]|$ is a tangled web in the FM interpretation.

We intend to arrange this by suitable construction of maps $\pi_{\texttt{clan}[A])}$ so that the very unlikely looking set of relations in the definition of a tangled web is enforced.

For extended type indices A, B, we define $B \ll A$ as meaning that B is a strict downward extension of A: $B \setminus A$ is nonempty and all elements of $B \setminus A$ are less than all elements of A. We provide that our abstract strict well-ordering $<_c$ on clans extends this relation.

The intention is that $\Pi(\mathtt{clan}[A])$, where $|A| \geq 2$, will be

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

and the superscript $|B| - |A| + 1$ indicates finite iteration of the power set operation of the FM interpretation.

The parent set $\Pi(\mathtt{clan}[\{\alpha\}])$ will be taken to be

$$\mathtt{clan}[\emptyset_\alpha] \cup \bigcup_{B \ll \{\alpha\}} \mathcal{P}_+^{|B|}(\mathtt{clan}(B)).$$

The set $\mathtt{clan}[\emptyset_\alpha]$ is a set of irregular atoms of size μ.

The clans $\mathtt{clan}[A]$ and the sets $\mathtt{clan}[\emptyset_\alpha]$ (which are not strictly speaking clans; they could be made clans with parent sets of irregular atoms, as in earlier versions, but this structure would not be used) are taken to exhaust the atoms, and none of these sets meets any of the others.

It should be clear that this intention will not be easy to implement: it cannot be achieved by mere stipulation. The difficulty (which we will evade with some subtlety below) is that the definition of $\mathcal{P}_*$ appears to depend on information about all maps π_C, and the definition of $\Pi(\mathtt{clan}[A])$ of course figures essentially in the definition of $\pi_{\mathtt{clan}[A]}$.

Our further intention is that there is an external isomorphism (an isomorphism present in the ground ZFA in which we work but not in the FM interpretation) between natural models of type theory restricted to $n+2$ types with base types $\mathtt{clan}[A]$ and $\mathtt{clan}[B]$ in the FM interpretation (models in which type i is $\mathcal{P}_*^i(\mathtt{clan}[A])$ or $\mathcal{P}_*^i(\mathtt{clan}[B])$, i ranging from 0 to $n+1$, and the membership and equality relations of the model are the appropriate restrictions of the membership and equality of the ambient ZFA) just in case $A \setminus A_n = B \setminus B_n$.

The second condition in the definition of tangled webs is clearly enforced by our second stipulation providing external isomorphisms.

We show that the first stipulation, concerning parent sets, enforces the first condition on tangled webs.

We want to show that $2^{\tau(A)} = \tau(A_1)$ when $|A| \geq 2$.

This translates to $|\mathcal{P}_*^3(\texttt{clan}[A])| = |\mathcal{P}_*^2(\texttt{clan}[A_1])|$. We show the inequalities in each direction.

$$|\mathcal{P}_*^3(\texttt{clan}[A])| = |\mathcal{P}_*(\mathcal{P}_*^2(\texttt{clan}[A]))| \geq |\mathcal{P}_*(\mathcal{P}_*(\texttt{clan}[A_1]))| = |\mathcal{P}_*^2(\texttt{clan}[A_1])|$$

establishes one direction, first using the basic cardinal inequality relating double power sets of clans and power sets of parent sets, then using the embedding of $\texttt{clan}[A_1)$ in the parent set of $\texttt{clan}[A]$.

$$|\mathcal{P}_*^2(\texttt{clan}[A_1])| \geq |\mathcal{P}_*(\Pi(\texttt{clan}[A_1]))| \geq |\mathcal{P}_*(\mathcal{P}_*^2(\texttt{clan}[A]))| = |\mathcal{P}_*^3(\texttt{clan}[A])|$$

establishes the other direction, first using the basic cardinality inequality relationg double power sets of clans and power sets of parent sets, then using the embedding of $\mathcal{P}_*^{|A|-|A_1|+1}(\texttt{clan}[A])$ into $\Pi(\texttt{clan}[A_1])$ which follows from our stipulation of parent sets above (as certainly $A \ll A_1$).

It should be clear from this that we expect the size of each $\mathcal{P}_*^{|B|-|A|+1}(\texttt{clan}(B))$ for $B \ll A$ to be just μ, the same size as the size of $\texttt{clan}[B]$, in the sense of our original ground interpretation of ZFA. These iterated power sets in the sense of the FM interpretation are to be extremely impoverished from the standpoint of the ground interpretation.

It can also be noted that the only value of $|B| - |A| + 1$ used in the argument above is 2: we could restrict our definition of the parent set to the iterated power sets of clans with exponent 2. But this is actually just as horribly tangled as the historically original definition, so I am sticking with it unless some powerful reason to do otherwise suggests itself.

We now abandon the stipulations concerning the definition of parent sets and existence of isomorphisms: we will make different stipulations below which turn out to enforce these.

Further restrictions imposed: We are going to stipulate that the clans are indeed just the clans $\texttt{clan}[A]$ which we have provided, and that for any extended type indices A, B we have $B \ll A \rightarrow \texttt{clan}[B] <_c \texttt{clan}[A]$. This can be enforced by letting the order on the clans be determined by the unique order on finite subsets of λ under which $\emptyset$ is maximal, $\{\alpha\} < \{\beta\}$, $A < B$ if A and B are nonempty and $\mathtt{max}(A) < \mathtt{max}(B)$ in the usual sense, and otherwise $A < B \leftrightarrow A \setminus \{\mathtt{max}(A)\} < B \setminus \{\mathtt{max}(B)\}$. Notice that this order enforces an extended type index appearing after all its downward extensions.

Fixing the circularity problem in the definition of parent sets: We define an A-support, where A is an extended type index, as a $\mathtt{clan}[A]$-support with the additional property that each element of the support belongs to a $\mathtt{clan}[B]$ or $\bigcup K_{\mathtt{clan}[B]}$ where $B \lll A$. We say that a set X is A-symmetric iff there is a $\mathtt{clan}[A]$-support for X which is an A-support. We call a $\mathtt{clan}[A]$-support for X which is an A-support, an A-support for X.

We define a *strongly symmetric set* as a subset of a set $\mathcal{P}^n(\mathtt{clan}[A])$, where $|A| \geq n \geq 1$, which is A_{n-1}-symmetric and is such that its elements are either atoms (obviously implying $n = 1$) or strongly symmetric sets. We define $\mathcal{P}^n_+(\mathtt{clan}[A])$ as the set of strongly symmetric elements of $\mathcal{P}^n(\mathtt{clan}[A])$. Notice that the elements of $\mathcal{P}^n_+(\mathtt{clan}[A])$ are elements of $\mathcal{P}^{n-1}_+(\mathtt{clan}[A])$ and so are A_{n-2}-symmetric, that is, symmetric in a stronger sense, and this continues as we consider elements of elements, etc. Clearly $\mathcal{P}^n_+(\mathtt{clan}[A]) \subseteq \mathcal{P}^n_*(\mathtt{clan}[A])$: these sets turn out to be the same, but it requires work to show this.

We stipulate that $\Pi(\mathtt{clan}[A])$, where $|A| \geq 2$, is

$$\mathtt{clan}[A_1] \cup \bigcup_{B \lll A} \mathcal{P}^{|B|-|A|+1}_+(\mathtt{clan}(B)).$$

The parent set $\Pi(\mathtt{clan}[\{\alpha\}])$ will be taken to be

$$\mathtt{clan}[\emptyset_\alpha] \cup \bigcup_{B \lll \{\alpha\}} \mathcal{P}^{|B|}_+(\mathtt{clan}(B)).$$

To define $\mathcal{P}^{|B|-|A|+1}_+(\mathtt{clan}(B))$ requires information about π_D only for $\mathtt{clan}[D] <_c \mathtt{clan}[B_{|B|-|A|+1-1}] = \mathtt{clan}[A]$, so the definition of the parent set of A no longer depends on knowledge of $\pi_{\mathtt{clan}[A]}$.

To define $\mathcal{P}^{|B|}_+(\mathtt{clan}(B))$ requires information about π_D only for $\mathtt{clan}[D] <_c \mathtt{clan}[B_{|B|-1}] = \mathtt{clan}[\{\alpha\}]$.

Thus the problem with apparent circularity of the definition of parent sets is removed. This definition will succeed if two things are established:

1. We need to show that $\mathcal{P}^n_+(\texttt{clan}[A])$ is in fact equal to $\mathcal{P}^n_*(\texttt{clan}[A])$.
2. We need to show that $|\mathcal{P}^n_+(\texttt{clan}[A])| = \mu$

Once these results are established, the system of clans can be constructed by constructing each map $\pi_{\texttt{clan}[A]}$ in turn along the order $<_c$, with a further refinement indicated in the discussion of isomorphisms at the end of the paper.

Analysis of supports: We note a special way in which an A-support can be extended. Any element of an A-support either is an atom belonging to a $\texttt{clan}[B]$ with $B \underline{\ll} A$, or a near-litter with parent either an atom belonging to a $\texttt{clan}[B]$ with $B \underline{\ll} A$ or an atom belonging to $\texttt{clan}[A_1]$ or an element of a $\mathcal{P}^{|D|-|C|+1}(\texttt{clan}[D])$ where $D \ll C \underline{\ll} B$. Only in the case $|A| = 1$ do we need to take into account the possibility of a near-litter with parent an irregular atom.

We organize the support so that near-litters in the support are disjoint (this can be done readily by making every near-litter a litter and adding further atoms to handle elements of symmetric differences of near-litters and the litters which replace them, but we do not require that near-litters in the supports we are building be litters, as we want the class of supports to be closed under application of suitably indexed allowable permutations). We provide that each atom appears after any near-litter in the support to which it belongs.

To add more information to the A-support, we enhance it by adding a C-support of each $\mathcal{P}^{|D|-|C|+1}(\texttt{clan}[D])$ with $C \ll B$ which occurs as the parent of a near-litter in the support; we add this before the near-litter in the order on the support. Note that we can do this because the minimum ordinal in extended type indices involved decreases: we apply the reorganization indicated in the previous paragraph to the new C-support added, then add further supports in the same way we indicate in this paragraph until the complete support contains a support for the parents of any suitable near-litter before it in the order. Further, if an item appears more than once in the order, place it as early as possible.

strong support: An A-strong support is an A-support in which distinct near-litters in the domain of the support are disjoint, in which any atom is preceded by any near-litter in the support which contains it, and any near-litter with parent a strongly symmetric set in $\mathcal{P}^{|D|-|C|+1}(\texttt{clan}[D])$ is preceded by a C-support for this set (which can be taken to be a C-strong support since the same closure conditions apply).

The discussion above reveals that any A-support can be extended to an A-strong support.

An [A-]extended strong support is an A-strong support in which each near-litter element is a litter and each atom is preceded by the litter containing it, and each regular atom which is the parent of an atom in the support [and not in A_1] is also in the support (with no order stipulation).

local bijection: An A-local bijection is an injective map whose domain is the same as its range, which contains all of $K_{\texttt{clan}[A])}$ in its domain and further contains small subsets (empty being a special case of small) of each $\texttt{clan}[B]$ with $B \lll A$, the local bijection acting as a permutation on each of the intersections of its domain with a clan and on the intersection of its domain with $K_{\texttt{clan[A]}}$.

A local bijection is an injective map whose domain is the same as its range, whose domain is the union of the set of local cardinals of litters with irregular parents with a set of regular atoms with small intersection with each clan, and which acts as a permutation on the intersection of its domain with each clan and on the intersection of its domain with the set of local cardinals with irregular atomic parents associated with each clan that includes litters with irregular parents.

There is a further technical requirement on both species of local bijection that if an atom in the domain belongs to a litter L whose parent is a regular atom, the parent must also belong to the domain of the local bijection. Note that an atom can only have finitely many iterated parents which are atoms.

exception: An exception of a [C-]allowable permutation ρ is a regular atom x [in $D \leq_c C$] in a litter L such that $\rho(x) \notin \rho(L)^\circ$ or $\rho^{-1}(x) \notin \rho^{-1}(L)^\circ$

Theorem (extension property): Each A-local bijection can be extended to a $\mathtt{clan}[A]$-allowable permutation with no exceptions other than elements of its domain. Each local bijection can be extended to an allowable permutation with no exceptions other than elements of its domain.

We do allow ourselves to abbreviate $\mathtt{clan}[A]$-allowable permutation to A-allowable permutation.

Proof of the (A-)extension property: Let ρ_0 be a[n A]-local bijection.

For each L, M litters included in the same clan, choose a map $\rho_{L,M}$, a bijection from $L \setminus \mathtt{dom}(\rho_0)$ to $M \setminus \mathtt{dom}(\rho_0)$. We will show that there is a unique [A-]allowable permutation extending ρ_0 and each $\rho_{L,\rho(L)^\circ}$, using the inductive hypothesis that the B-extension property with the additional strong relation to maps $\rho_{L,M}$ holds for each $B[< A]$ (this is induction on the smallest element of the extended type index, thinking of the non-indexed case as stage λ in the induction).

We show this by indicating how to compute ρ at any atom x by a recursive computation along an extended strong support of x. Suppose we have an extended strong support S of x in which x is the last element.

If L is a near-litter in S and $[L] \in K_A$, the value of $\rho([L]) = \rho_0([L])$ and so of $\rho(L)^\circ$ is known, and we can compute $\rho(L)$ as the elementwise image of L under the union of ρ_0 and $\rho_{L,\rho(L)^\circ}$.

If L is a near-litter in S and the parent of L is an atom not in A_1 (including an irregular atom), the value of ρ at the parent can be computed as a value of ρ_0, whence the value of $\rho([L])$ and so of $\rho(L)^\circ$ is known, and we can compute $\rho(L)$ as the elementwise image of L under the union of ρ_0 and $\rho_{L,\rho(L)^\circ}$.

If L is a near-litter in S and the parent of L is a set, we can by inductive hypothesis compute the value of ρ at all elements of a B-support of L, $B \ll A$. We can determine a B-local bijection ρ_0' an extension of which must send each element of the B-support to the correct value. We set values of ρ_0' at atoms to the values of ρ at those atoms already computed. At each litter M in the B-support, the value of the extension of ρ_0' at $[M]$ is already fixed, as all elements of a B-support of L have values fixed; all exceptions of the extension are already handled by including all elements of orbits in ρ_0 meeting L in the domain of ρ_0'.

Since the value of the extension ρ' at each element of the B-support is fixed by our choices of domain elements of ρ'_0 and prior computations of ρ, the value of ρ' at $[L]$ is fixed, and moreover is the only possible value for the extension ρ at $[L]$: the value of $\rho([L])$ and so of $\rho(L)^\circ$ is known, and we can compute $\rho(L)$ as the elementwise image of L under the union of ρ_0 and $\rho_{L,\rho(L)^\circ}$.

Now for any atom x in a litter L we can compute $\rho(x)$ as either $\rho_0(x)$ or $\rho_{L,\rho(L)^\circ}(x)$, as we have already computed $\rho(L)$.

We note that computations of ρ along different supports will not give different values, basically because supports can be merged.

It should be clear from the method of computation that the total map computed is $[\mathtt{clan}[A]$-$]$allowable and has no exceptions outside its domain.

elements of appopriate iterated power sets of clans are strongly symmetric: We show that for each n with $1 \leq n \leq |A|$, we have $\mathcal{P}^n_+(\mathtt{clan}[A] = \mathcal{P}^n_*(\mathtt{clan}[A]$.

We show this by induction on n.

We begin with the case $n = 1$. We aim to show that any symmetric subset of $\mathtt{clan}[A]$ (which will of course be hereditarily symmetric) has an A-support. In fact, we can show something stronger: if $X \in \mathcal{P}_*(\mathtt{clan}[A])$ has extended strong support S, then it has support $S \cap (\mathtt{clan}[A] \cup \bigcup K_{\mathtt{clan}[A]})^2$, which is of course an A-support.

We first note that each set X of the form $Y \Delta \bigcup \Lambda$ or $\mathtt{clan}[A] \setminus (Y \Delta \bigcup \Lambda)$, where Y is a small collection of atoms in $\mathtt{clan}[A]$ and Λ is a small collection of litters included in $\mathtt{clan}[A]$, has a support included in $(\mathtt{clan}[A] \cup \bigcup K_{\mathtt{clan}[A]})^2$. In English we say, sets with small symmetric difference from small or cosmall unions of litters.

Now let X be an arbitrary element of $\mathcal{P}^*(\mathtt{clan}[A])$ with extended strong support S. Choose any two distinct atoms x, y in $\mathtt{clan}[A]$ which do not belong to S and which both belong to the same element of S or both belong to no element of S. Define a local bijection fixing each atom in S and mapping x to y and y to x. Extend it to an allowable permutation ρ_{xy} using the extension property. Notice that each litter L in S will be fixed by ρ_{xy}: suppose L first to be moved; $|L|$ will be fixed because every element of a support for the parent of L is fixed, and the litter L will be fixed because if it were moved there would be an exception mapped into or out of it, and the only possible exceptions moved by ρ_{xy} are x and y, which are either both in the same element of S or both not in any element of S (other elements of the domain of the local bijection are fixed by it). These maps transposing atoms x and y illustrate that any two atoms in a litter $L \in S$ included in $\mathtt{clan}[A]$ or in the set $\mathtt{clan}[A] \setminus (S \cup \bigcup S)$ can be exchanged without moving the set X. But this means that X is of the form $Y \Delta \bigcup \Lambda$ or $\mathtt{clan}[A] \setminus (Y \Delta \bigcup \Lambda)$, where Y is a set of atoms in S and Λ is a set of litters in S, and this establishes our result for the case $n = 0$ in quite a strong form.

Suppose $\mathcal{P}^k_+(\mathtt{clan}[A] = \mathcal{P}^k_*(\mathtt{clan}[A]$. Let X be an arbitrary element of $\mathcal{P}^{k+1}_*(\mathtt{clan}[A]$ with extended strong support S. Our aim is to show that X has an A_k-support. We claim, to be exact, that the set S^- defined as the maximal A_k-support included in S is an A_k-support for X.

Let ρ be an A_k-allowable permutation which fixes each element of S^-. Our aim is to show that ρ fixes X.

We choose $x \in X$. By inductive hypothesis, x has an A_{k-1}-extended strong support T. We define a local bijection ρ'_0. This map fixes each element of S. We extend T to T^*, its closure under application of ρ and ρ^{-1} (to atoms and near-litters alike); T^* is small, and $S \cup T^*$ can be presented as a strong support (it is not an extended strong support but it can be verified that it has useful closure properties) by imposing a suitable order: ρ'_0 sends each atomic element of T^* to its image under ρ. The extension ρ' of ρ'_0 obtained from the extension property sends each litter in T^* to its image under ρ: if this failed to be true, there would be a first litter L in T^* moved by $\rho' \circ \rho^{-1}$; $|L|$ would be fixed by this composition because every element of a support for the parent of L would be so fixed. This means that an exception would be mapped into or out of the litter by this map. But all exceptions of ρ are mapped by ρ and ρ' to the same value. This means that $\rho(x) = \rho'(x)$ since ρ and ρ' have the same values on an A_{k-1}-support of x and are both A_{k-1}-allowable (certainly A_k-allowable implies A_{k-1}-allowable). But ρ' fixes X, since it fixes each element of S, so $\rho(x) \in X$. The same argument applied to ρ^{-1} shows that $\rho^{-1}(x) \in X$. But then $\rho(X) = X$ and S^- is an A_k-support of X as required.

coding functions and orbit specifications: For any support S for an object x, we define $\chi^S_x(\rho(S))$ as $\rho(x)$ for each allowable permutation ρ. This might not look like a function definition, but if $\rho(S) = \rho'(S)$ we have $\rho(x) = \rho'(x)$ because S is a support for x (consider that $\rho' \circ \rho^{-1}$ fixes S, so fixes x). We call the functions χ^S_x *coding functions*.

The domain of a coding function is an orbit in the strong supports under the allowable permutations. We analyze these orbits. We claim that the orbit to which a strong support S belongs can be specified by a function $\sigma(S)$ from the order type of S to certain data.

1. We use the notation S_α for the element of the domain of S in position $\alpha < \kappa$ in the order on S. We use the notation $\iota(S_\alpha)$ for the extended type index such that $S_\alpha \in \texttt{clan}[A]$ or $S_\alpha \subseteq \texttt{clan}[A]$ We use the notation S^α_B for the restriction of S to $\{S_\beta : \beta < \alpha \wedge \iota(S_\beta) \ll B\}$.
2. If S_α is an atom, $\sigma(S)(\alpha) = (1, A, \beta)$, where $S_\alpha \in \texttt{clan}[A]$ and either $S_\alpha \in S_\beta$, or $\beta = \kappa$ and there is no S_γ containing S_α.
3. If S_α is a near-litter with parent a regular atom, $\sigma(S)(\alpha) = (2, A, \gamma)$, where $S_\alpha \in \bigcup K_{\texttt{clan}[A]}$ and either $\pi_{\texttt{clan}[A]}([S_\alpha]) = S_\gamma \in \texttt{clan}[A_1]$ or $\gamma = \kappa$ and the relevant parent is not in the domain of S.
4. If S_α is a near-litter with parent a set, $\sigma(S)(\alpha) = (3, A, g)$, where $S_\alpha \in \bigcup K_{\texttt{clan}[A]}$, g is a coding function, and $g(S^\alpha_A) = \pi_{\texttt{clan}[A]}([S_\alpha])$.
5. If S_α is a near-litter with parent an irregular atom, $\sigma(S)(\alpha) = (4, A)$, where $S_\alpha \in \bigcup K_{\texttt{clan}[A]}$.

Values of σ are called *orbit specifications*.

We justify the name "orbit specification" by showing that two strong supports are in fact in the same orbit iff they have the same orbit specification. One direction is obvious: it should be clear that for any allowable permutation ρ and strong support S, $\sigma(\rho(S)) = \sigma(S)$.

What remains to be shown is that if we have S and T strong supports with the same orbit specification, we can find an allowable permutation ρ such that $\rho(S) = T$. We establish this by constructing an appropriate local bijection ρ_0 and applying the extension property.

1. if S_α and S_β are atoms, set $\rho_0(S_\alpha) = S_\beta$.
2. if S_α and S_β are near-litters, for any $x \in S_\alpha \Delta S_\alpha^\circ$, we designate a y such that $\rho_0(x) = y$, and for any $y \in T_\alpha \Delta T_\alpha^\circ$, we designate an x such that $\rho_0(x) = y$. Further, we need to designate values $\rho_0(x)$ and $\rho_0^{-1}(x)$ for each x in the domain of ρ_0 for which other conditions do not specify these values, under the constraints that ρ_0 and ρ_0^{-1} are injective and that any atom belonging to a near-litter S_α must be mapped to something in T_α by ρ_0, any atom belonging to a near-litter T_α must be mapped to something in S_α by ρ_0^{-1}, and anything which is in no S_α must be mapped by ρ_0 to something not in any T_α, and anything which is in no T_α must be mapped by ρ_0^{-1} to something not in any S_α. Only a small collection of new values will be needed, countable orbits being filled out for each atom S_α or T_α and each atom in a symmetric difference of near-litter and litter $S_\alpha \Delta S_\alpha^\circ$ or $T_\alpha \Delta T_\alpha^\circ$. We may need to assign values at local cardinals of litters with irregular atomic parents in a similar way, and we may need to choose images and inverse images for regular atomic parents of near-litters following the same rules as above. Further, if we assign values at any element of a litter with regular atomic parent, we are required to assign values at the parent.

The extension of the map ρ_0 thus defined to an allowable permutation ρ will send S to T. It clearly does so for atomic values, and clearly does so in all cases of near-litter values, with comment needed only in the case of near-litters with set parents. In the case of near-litters with set parents, we can use the inductive hypothesis that the procedure works for shorter orbit specifications to get ρ to have the correct value T_S^α at S_S^α, and thus to send $g(S_S^\alpha)$ to $g(T_S^\alpha)$, and so send $[S_\alpha]$ to $[T_\alpha]$. That it sends S_α exactly to T_α follows from the handling of anomalous atoms in near-litters indicated above.

Each appropriate iterated power set of a clan in the FM sense is of size μ: We show that sets $\mathcal{P}_*^n(\texttt{clan}[A])$ for $|A| \geq n$ are of size μ, by an analysis of the size and number of orbits in the allowable permutations.

What we actually show is that the elements of each $\mathcal{P}_*^n(\texttt{clan}[A])$ are precisely the union of the ranges of a family F_n^A of coding functions, and that $|F_n^A| < \mu$. The collection of A_{n-1} supports, which is a superset of each domain of a coding function in F_n^A, is exactly of size μ, so $\mathcal{P}_*^n(\texttt{clan}[A])$ is of size no more than μ. There are μ iterated singletons of atoms in $\mathcal{P}_*^n(\texttt{clan}[A])$, so its size is exactly μ.

For the case $n = 1$, we describe the family of functions F_1^A which we will use: the domain orbits in the strong supports are all the ones consisting entirely of atoms in $\texttt{clan}[A]$ and litters included in $\texttt{clan}[A]$. There are κ distinct domains of this type (with different cross-referencing of which atoms belong to which near-litters or to none of them). For each domain with specification of order type $\alpha < \kappa$, a function F in F_1^A is specified by a sequence f of bits (values 0 or 1) of order type $\alpha + 1$: $F(S)$ contains each atom S_β iff $f_\beta = 1$, and either includes or does not meet each near-litter S_β, including it iff $f_\beta = 1$; it includes each element of $\texttt{clan}[A]$ not in $S \cup \bigcup S$ iff $f_\alpha = 1$. We have seen above that all symmetric subsets of $\texttt{clan}[A]$ can be described in this way. It is then evident that there are no more than $2^\kappa < \mu$ elements in F_1^A for each A, and that the union of the ranges of the elements of F_1^A is $\mathcal{P}_*^1(\texttt{clan}[A])$.

In showing that $\mathcal{P}_*^n(\texttt{clan}[A])$ has cardinality μ, $n > 1$, we assume that $\mathcal{P}_*^m(\texttt{clan}[B])$ has been shown in all cases where $\texttt{min}(B_{m-1}) < \texttt{min}(A_{n-1})$, by construction of families of coding functions F_m^B.

Let $X \in \mathcal{P}_*^n(\texttt{clan}[A])$. Let S be a strong support for X. For each $x \in X$, choose an A_{n-2} support T which is an end-extension of the maximum A_{n-2} support included in S (if $n = 2$ we further cut down to just atoms and litters in $\texttt{clan}[A]$ included in S). We write $T \leq_{n-2}^A S$ to express that T is an end extension of the maximal A_{n-2}-strong support included in S (with the additional reduction in case $n = 2$). The object x will be an image $F_x(T)$ for some $F_x \in F_{n-1}^A$. The collection of such F_x's will determine our coding function, as we spell out in the next paragraph.

We explicitly define the coding function which will have X in its range and belong to F_n^A: for any U with $\sigma(U) = \sigma(S)$, $F_X(U) = \{F_x(V) : F_x \in F_{n-1}^A \wedge (\exists x \in X : (\exists T : F_x(T) = x \wedge T \leq_{n-2}^A S \wedge \sigma(V) = \sigma(T) \wedge V \leq_{n-2}^A U))\}$. We define F_n^A as the set of all such functions for $X \in \mathcal{P}_*^n(\texttt{clan}[A])$.

Note that F_X is exactly determined by $\sigma(S)$, the orbit specification of S, and the set $\{F_x \in F_{n-1}^A : (\exists x \in X : (\exists T : F_x(T) = x \wedge T \leq_{n-2}^A$

$S))\}$. The latter set is of size $< \mu$ because $|F^A_{n-1}| < \mu$ by inductive hypothesis and μ is a strong limit cardinal. The collection of all orbit specifications of A_{n-1}-strong supports is also of size $< \mu$, because each such specification is a small structure built from extended type indices, ordinals less than κ, and coding functions belonging to F^B_m's already known to be of size $< \mu$ by our inductive hypotheses. Thus the set F^A_n of all F_X's constructed as indicated is of size $< \mu$ as desired. It is important to notice that in fact *all* coding functions (of suitable extended type index) belong to the sets of coding functions we are constructing, except in the case $n = 1$.

It should be clear by examination that F_X actually is a coding function. Its domain is an orbit in the strong supports and it satisfies $F_X(\rho(U)) = \rho(F_X(U))$ by inspection of the details of the definition (basically by properties of orbit specifications).

It should be evident from the construction that $X \subseteq F_X(S)$: we explicitly described how to build F_x and T so that $F_x(T) = x$ would fall in this set. It remains to show that $F_X(S) \subseteq S$. A general element of $F_X(S)$ is of the form $F_x(T')$ where there are $x \in X, T \leq^A_{n-2} S$, such that $F_x(T) = x$, and $T' \leq^A_{n-2} S$ (of course T and T' have the same specification). $S \cup T$ and $S \cup T'$ can be presented with the same order specification (end extending the original order on S and agreeing on the extension with the order on T, T'), so there is an allowable permutation sending $S \cup T$ (suitably ordered) to $S \cup T'$, which fixes X and sends $x = F_x(T)$ to $F_x(T')$, so in fact $X = F_X(S)$, which completes the argument for size of iterated power sets of clans.

Existence of isomorphisms: Our intention is to present maps I_α preserving all structure we are interested in which send each $\texttt{clan}[A]$ for which α strictly dominates A to $\texttt{clan}[A \cup \{\alpha\}]$.

Arranging for these isomorphisms to exist is simple. For each B with α dominating B, define $I_\alpha \lceil \texttt{clan}[B]$ as an arbitrarily chosen bijection from $\texttt{clan}[B]$ to $\texttt{clan}[B \cup \{\alpha\}]$ whose action sends $K_{\texttt{clan}[B]}$ to $K_{\texttt{clan}[B\cup\{\alpha\}]}$ (it has to preserve the litter structure). Each map $\pi_{\texttt{clan}[B\cup\{\alpha\}]}$ is then taken to be the result of the action of the union of all these maps on $\pi_{\texttt{clan}[B]}$ (which we can do, because $\texttt{clan}[B \cup \{\alpha\}]$ appears later in the order on clans than $\texttt{clan}[B]$): the action of I_α on any set all atoms in the transitive closure of which are in suitable clans is defined in the usual way ($I_\alpha(X) = I_\alpha``X$). The map $I_\alpha\lceil\mathtt{clan}[\emptyset_\beta])$ from $\mathtt{clan}[\emptyset_\beta]$ [which has no structure of interest] to $\mathtt{clan}[\{\alpha\}]$, for $\beta < \alpha$, is also needed and may be taken to be quite arbitrary. Everything commutes exactly with the relationships built into the model at the outset. The result that all elements of iterated power sets of clans (with suitably low index) are strongly symmetric has a lot to do with this working: the strongly symmetric power sets are clearly copied isomorphically in the way indicated: this would be unlikely to be true if there was less symmetry in the iterated power sets.

Chasing down isomorphisms of power sets: The model of TST with n types with base type of cardinality $|\mathcal{P}^2_*(\mathtt{clan}[A])|$ has top type $|\mathcal{P}^{n+1}_*(\mathtt{clan}[A])|$. The top type is A_n-symmetric, that is, its definition depends only on maps $\pi_{\mathtt{clan}[B]}$ for $B \ll A_n$ and on $K_{\mathtt{clan}[A_n]}$ (and iterated singletons in the top type contain all information about lower types conveniently). All such models with n types with base type of cardinality $|\mathcal{P}^2_*(\mathtt{clan}[B])|_*$ are easily seen to be isomorphic if $A \setminus A_{n+1} = B \setminus B_{n+1}$: start with the model based on $|\mathcal{P}^2_*(\mathtt{clan}[A \setminus A_{n+1}])|_*$ and apply the isomorphisms of the previous paragraph repeatedly to get to the models based on $|\mathcal{P}^2_*(\mathtt{clan}[A])|_*$ and $|\mathcal{P}^2_*(\mathtt{clan}[B])|_*$ respectively. Having the theory in a tangled web depend on the smallest $n+1$ elements of the extended type index *does* work to establish Con(NF), though it seems inelegant. I think that the theory actually depends only on the smallest n elements of the extended type index, but tracing out out this works would involve additional effort and is not required for our main result.