A Fraenkel-Mostowski construction for patterns in the exponential map $\kappa \mapsto 2^\kappa$, with an application to Quine's New Foundations

M. Randall Holmes

7/18/2018: various cleanup. 7/12/2018: cleanup of definition of framework maps I^*_C, 6:48 pm. 7/11/2018: solution to the preceding noted problem in place. Corrected a glitch in the definition of strong support. Fixed silly "typo" in parenthetical remark about ambiguity. 8:45 am. 7/11/2018 noted technical issue which doesn't affect the NF proof and probably does work in the general case: the tree used in the NF proof has all downward closed subtrees rigid, and some work is required to make it clear that assertions about isomorphisms work in the general case (NF proof not affected: the condition that subtrees are rigid could be imposed, but I don't want to do this unless I have to). 7/10/2018 corrected stupidity in description of isomorphisms; this could readily have howlers in it, watch for version changes. Correction re equality of clans included in parent sets, 4:19 pm Boise time. Spelled out starred framework maps, 5:28 pm. Further remarks about isomorphisms, 6:23 pm.

Definition of TST and TST$_n$: The simple typed theory of sets TST is the sorted first order theory with logic and membership as primitive relations and with sorts indexed by the natural numbers, whose axioms are Extensionality,

$$(\forall xy.(\forall z : z \in x \leftrightarrow z \in y) \rightarrow x = y),$$

for each choice of variables x, y, z which yields a well-sorted formula, and Comprehension

$$(\exists A : (\forall x : x \in A \leftrightarrow \phi)),$$

for each formula ϕ and variables A, x chosen so that A does not occur free in ϕ and the complete formula is well-sorted. The resemblance of these axioms to the inconsistent axioms of "naive set theory" is not an accident.

The theory TST_n, for each natural number n, is the subtheory of TST in which the sorts i are restricted so that $0 \leq i < n$. A *natural model of TST_n* is a model of this theory in which each type i is implemented as $\mathcal{P}^i(X)$ for a fixed set X (implementing type 0) and equality and membership relations between appropriate sorts are implemented by intersections of the true equality relation with sets $\mathcal{P}^i(X) \times \mathcal{P}^i(X)$ and intersections of the true membership relation with sets $\mathcal{P}^i(X) \times \mathcal{P}^{i+1}(X)$. It should be clear that the first-order theory of a natural model of TST_n is completely determined by the cardinality of the set implementing type 0 and the parameter n determining the number of types.

Our background theory and aims: We work in ZFA with Choice.

Let $\mathbb{T}$ be a set and P a possibly partial function on $\mathbb{T}$ with the property that each nonempty subset X of $\mathbb{T}$ contains an element which is not in the range of $P\lceil X$. This puts a well-founded tree structure on $\mathbb{T}$: P is for "parent". We define $x \leq_T y$ as $(\exists n \in \mathbb{N} : P^n(x) = y)$; this well-founded relation defines the tree structure in a more conventional way.

We further provide a well-ordering $\leq^w_T$ of $\mathbb{T}$ which extends $\leq_T$. We define $<_T$ and $<^w_T$ in the obvious way.

We will present an FM construction and an injective function τ such that for each $t \in \mathbb{T}$, $\tau(t)$ is a Scott cardinal in the sense of the FM interpretation and $2^{\tau(t)} = \tau(P(t))$ when the latter is defined (in the sense of the FM interpretation), and, further, the natural model of TST_n [in the sense of the FM interpretation] with base type of size $\tau(t)$ (where $P^n(t)$ is defined) will be isomorphic in the sense of the ground interpretation of ZFA to the natural model of TST_n [in the sense of the FM interpretation] with base type of size $\tau(t')$ (where $P^n(t')$ is defined), if there is an isomorphism from the restriction of $\leq_T$ to $\{s \in \mathbb{T} : s \leq_T P^n(t)\}$ to the restriction of $\leq_T$ to $\{s \in \mathbb{T} : s \leq_T P^n(t')\}$ which sends t to t'. This isomorphism between natural models is not an object of the FM interpretation; it does have the effect that the natural models are seen in the FM interpretation to have the same first-order theory extending TST_n.

We will further point out that this allows us to establish the consistency of Quine's set theory New Foundations, which has been a long-standing unsolved problem.

cardinal parameters: Let κ be an uncountable regular cardinal. Let μ be a strong limit cardinal which is greater than κ and greater than $|\mathbb{T}|$.

clans: Let $\mathbb{C}$ be a bijection from $\mathbb{T}$ to a pairwise disjoint family of sets, with $|\mathbb{C}(t)| = \mu$ for each $t \in \mathbb{T}$. We call the sets $\mathbb{C}(t)$ *clans*. If C is a clan $\mathbb{C}(t)$, we define $P_c(C)$ as $\mathbb{C}(P(t))$. We define $C \leq_c D$, where $C = \mathbb{C}(s)$ and $D = \mathbb{C}(t)$, as holding iff $s \leq_T t$. We define $C \leq_c^w D$, where $C = \mathbb{C}(s)$ and $D = \mathbb{C}(t)$, as holding iff $s \leq_T^w t$. We define $<_c$ and $<_c^w$ in the obvious way. We refer to the elements of clans as *regular atoms*; there may be some other atoms, termed *irregular atoms*.

litters and local cardinals: For each clan C, we provide a partition $\Lambda(C)$ of C into sets of cardinality κ. Elements of $\Lambda(C)$ are called *litters*. For each $L \in \Lambda(C)$, we define $[L]$, the *local cardinal of* L, as the collection of all subsets of C with small symmetric difference from L. We define $K(C)$ as $\{[L] : L \in \Lambda(C)\}$. Elements of any of the sets $\bigcup K(C)$ are called *near-litters*. For each near-litter N, we define N° as the litter with small symmetric difference from N. We may write $[N]$ for $[N^\circ]$ and refer to this as the local cardinal of N. We may refer to elements of $N \Delta N^\circ$ as *anomalies* of N.

parent sets and maps: With each clan C, we will associate a set $\Pi(C)$ whose nature is for the moment unspecified (it may contain sets or atoms or both) and a bijection $\pi_C : K(C) \to \Pi(C)$ (from which we can see $|\Pi(C)| = \mu$). We call $\Pi(C)$ the *parent set* of the clan C, and refer to $\pi_C([N])$ as the parent of the near-litter $N \subseteq C$. We may employ the notation $\pi(N)$ for the parent of a near-litter N, since C can be deduced from N.

permutations: We extend permutations ρ of the atoms to sets by the usual rule $\rho(A) = \rho``A$. We define an *allowable permutation* as a permutation which fixes all the maps π_C. We define a C-allowable permutation as a permutation which fixes each π_D for $D <_c C$ and also fixes $K(C)$.

Note that a [C-]allowable permutation ρ will map litters included in a clan D to near-litters included in the same clan [for $D \leq_c C$], since the permutation must fix $K(D)$. Such a litter L will be mapped to a near-litter $N \in \rho([L])$ with small symmetric difference from the litter $N^\circ = \rho(L)^\circ$. It is useful to note that $\rho(L)^\circ$ is completely determined by $[L]$, the local cardinal of L.

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

supports and symmetry: A *support* is a well-ordering S (which we may write $\leq_S$ if we have occasion to use a relation symbol) on a small set of atoms and near-litters. A C-support is a support all of whose domain elements are elements or subsets of clans $D \leq_c C$. An object x is [C-]*symmetric* if there is a [C-]support S such that any [C-]allowable permutation which fixes S fixes x; we also say that x has support S in this situation. Note that a permutation which fixes a support S fixes each element of its domain individually.

We may say "x in S" where S is a support, meaning $x \in \mathtt{dom}(S)$; in general, supports may be confused with their domains where we are not concerned with the order on them. If S is a support and $T \subseteq S$, $S \setminus T$ refers to $S \cap (\mathtt{dom}(S) \setminus \mathtt{dom}(T))^2$. We cannot neatly define the union of two supports S and T, as this involves imposing a suitable order on the union of their domains, which contains extra information: we might say *a* support $S \cup T$ to indicate a support whose domain is the unions of the domains of S and T and which extends each of S and T as orders.

FM interpretations: We note that the atoms and hereditarily [C-]symmetric sets make up a model of ZFA (presumably without choice) for standard reasons. We introduce the notation $\mathcal{P}_*(X)$ for the collection of hereditarily symmetric subsets of a hereditarily symmetric set X (the power set operation of the FM interpretation). We introduce the notation $|X|_*$ for the set of all Y of minimal rank such that there is a hereditarily symmetric bijection from X to Y (the Scott cardinal in the FM interpretation). We reserve the right to use unadorned $\leq$ and related symbols for the partial order on Scott cardinals in the FM interpretation.

Cardinal inequalities in the FM interpretation: It is evident that $|\Pi(C)|_* = |K(C)|_* \leq |\mathcal{P}^2_*(C)|_*$, because the bijection π_C is invariant, as are $K(C)$ and $\Pi(C)$. Further, because of the existence of the invariant $(B \in \mathcal{P}_*(K(C)) \mapsto \bigcup B)$, which is injective because $K(C)$ is a pairwise disjoint collection of sets, $|\mathcal{P}_*(\Pi(C))|_* = |\mathcal{P}_*(K(C))|_* \leq |\mathcal{P}^2_*(C)|_*$.

Atoms in parent sets: For each clan C, we provide that the intersection of $\Pi(C)$ with the set of atoms is a set of atoms $\Pi_*(C)$ of size μ, which coincides with $P_c(C)$ if this is defined, and otherwise meets no clan. For distinct C, D, $\Pi_*(C)$ and $\Pi_*(D)$ are disjoint unless $P_c(C)$ and $P_c(D)$ are defined and equal. We provide that every atom belongs either to a clan or to a parent set or both: it should be clear that the cardinality of the set of atoms is μ once this provision is made.

Sets in parent sets: We define $\mathcal{P}^0_+(C)$ as C, and further define $\mathcal{P}^{n+1}_+(C)$ as the set of all $P^n_c(C)$-symmetric elements of $\mathcal{P}(\mathcal{P}^n_+(C))$, when $P^n_c(C)$ is defined ($P^0_c(C)$ of course is C). Note that $\mathcal{P}^n_+(C) \subseteq \mathcal{P}^n_*(C)$ is evident.

We refer to elements of sets $\mathcal{P}^{n+1}_+(C)$ as *strongly symmetric* elements of iterated power sets of clans.

We will provide that $\Pi(C) \setminus \Pi_*(C) = \bigcup_{D \in P_c^{-1}``\{C\}} \mathcal{P}^2_+(D)$ (we will explain below in detail why we can do this).

framework maps: We start with a

Definition: For any clan C, we define C^* as the $\leq^w_c$ first clan C' such that there is an isomorphism between the restrictions of $\leq_c$ to $\{D : D \leq_c C\}$ and $\{D : D \leq_c C'\}$.

We provide a bijection I_C from C to C^* (I_{C^*} is the identity map) such that the action of I_C on sets defined in the natural way sends $K(C)$ to $K(C^*)$ (it preserves litter structure). We call these "framework maps": informally, they provide a scheme of "identifications" of atoms in clans in analogous positions in the tree structure on clans. The functions I_C are not elements of the domain of the FM interpretation: this scheme of identifications is "external".

We also provide a framework map I^*_C from $\Pi_*(C)$ to $\Pi_*(C^*)$:

1. If $C = C^*$, or if $P_c(C) = P_c(C^*)$, I^*_C is the identity map on $\Pi_*(C) = \Pi_*(C^*)$.

2. If $C \neq C^*$ and $\Pi_*(C)$ is not a clan, I^*_C from $\Pi_*(C)$ to $\Pi_*(C^*)$ is chosen arbitrarily.

3. If $C \neq C^*$ and $P_c(C)$ is defined and distinct from $P_c(C^*)$ and $\Pi_*(C)$ is the clan $P_c(C) = P_c(C)^*$, choose I^*_C from $\Pi_*(C)$ to $\Pi_*(C^*)$ arbitrarily [choosing the same map as I^*_D for each D with $P_c(C) = P_c(D)$], except in the case where $P_c(C) = P_c(C^*)^*$, in which I^*_C must be $I^{-1}_{P_c(C^*)}$.

4. If $C \neq C^*$ and $P_c(C)$ is defined and distinct from $P_c(C^*)$ and $\Pi_*(C)$ is the clan $P_c(C) \neq P_c(C)^*$, we define $I^*_C = I^*_D \circ I_{P_c(C)}$, where $D^* = C^*$ and $P_c(D) = P_c(C)^*$ (noting from the previous clause that choice of D under these conditions does not affect what I^*_D is).

These clauses enforce a coherence condition: if $C^* = D^*$ and $P_c(C)$, $P_c(D)$ are defined and $P_c(C)^* = P_c(D)^*$, we must have $(I^*_D)^{-1} \circ I^*_C = I^{-1}_{P_c(D)} \circ I_{P_c(C)}$: "identifications" of atoms in parent sets coordinate with "identifications" of the same atoms in clans.

We also refer to compositions (and inverses) of framework maps as framework maps: if $C^* = D^*$, $I^{-1}_D \circ I_C$ is a framework map from C to D.

Description of the construction: We proceed along the well-ordering $<^w_c$. Let C be the first clan at which we have not yet defined π_C.

If $C^* = C$ we define π_C as an arbitrarily chosen bijection from $K(C)$ to $\Pi_*(C) \cup \bigcup_{D \in P_c^{-1}``\{C\}} \mathcal{P}^2_+(D)$, if the latter set is of size $\mu$. If the proposed parent set $\Pi_*(C) \cup \bigcup_{D \in P_c^{-1}``\{C\}} \mathcal{P}^2_+(D)$ is of cardinality $> \mu$ (it will certainly be of cardinality $\geq \mu$), the construction fails at this point.

Note that the definition of the parent set $\Pi(C)$ depends on information about maps π_E only for $E <_c C$, since the sets $\mathcal{P}^2_+(D)$ involved are all $P_c(D) = C$-symmetric. Since the maps π_E for $E <_c C$ have by inductive hypothesis already been constructed, we can do this (subject to the cardinality condition).

If $C^* \neq C$, we choose a particular isomorphism ι from $\leq_c$ on $\{D : D \leq_c C^*\}$ to $\leq_c$ on $\{D : D \leq_c C\}$, and define π_C as the image of π_{C^*} under the action on sets of the union of all the framework maps $I^{-1}_{\iota(D)} \circ I_D$ between corresponding clans (on clans $D \leq_c C^*$ suitable structure has already been provided and is simply copied to clans $\iota(D) \leq_c C$), and the framework map $(I^*_C)^{-1}$ from $\Pi_*(C^*)$ to $\Pi_*(C)$. It is worth noting that the only D for which $I^{-1}_{\iota(D)} \circ I_D$ sees use are the immediate children of C^* and C^* itself: this has the effect that we do not need to concern ourselves with coherence of choices of maps ι at different clans C.

Note that if we used a different ι' from $\leq_c$ on $\{D : D \leq_c C^*\}$ to $\leq_c$ on $\{D : D \leq_c C\}$ we would get the same π_C. This reflects an automorphism in the structure being defined parallel to the automorphism $\iota^{-1} \circ \iota$ of $\leq_c$ on $\{D : D \leq_c C^*\}$.

The intended equation $\Pi(C) = \Pi_*(C) \cup \bigcup_{D \in P_c^{-1} \text{``} \{C\}} \mathcal{P}^2_+(D)$ will hold for each C for which the definition succeeds: the purpose of the framework maps is to make the structure produced parallel in every respect where the underlying structure of $\mathbb{T}$ is parallel.

existence of strong supports: A [C-]support S is a [C-]*strong support* iff it is a [C]- support and satisfies certain additional conditions.

1. Distinct near-litters in the domain of the support are disjoint. It is straightforward to take an arbitrary S and modify it to satisfy this condition, by replacing each near-litter N in S with N° and the elements of $N \Delta N^\circ$.
2. Each atom in the support is preceded in the support by the near-litter in the support which contains it as an element, if there is one. There is no requirement that the near-litter containing an atom in the support be in the support; this condition just concerns its placement if it is present.
3. For each near-litter in the domain of the support with parent a regular atom [not in $\Pi_*(C)$], the parent atom also belongs to the domain of the support (with no order constraint).
4. Each near-litter in the domain of the support with parent a set in $\mathcal{P}^2_+(D)$, [$P_c(D) <_c C$] is preceded in the order by all elements of a $P_c(D)$-strong support for this set. Note that this only applies to near-litters included in clans $E <_c C$.

Every object has a strong support, because one can iterate closure under the conditions through ω stages, and one will not have an infinite regress, because any chain of objects each of which occasions the introduction of the next one has the clans in which elements of the chain are elements or subsets "decreasing" in the sense of $<_c$, so all such chains are finite. When an atom is repeated, keep its earliest occurrence after a litter which contains it, and when a near-litter appears with one that it meets, replace each occurrence with the corresponding litter and its anomalies (placing the anomalies after any litter which includes them and otherwise after the given litter) and keep the earliest occurrence only of any litter which is repeated. Note that one is not obligated to precede an atom with a near-litter containing it, and we do not do this in the closure process.

local bijection: A [C-]local bijection is a bijection ρ whose domain (also its range) meets each clan [$\leq_c C$] in a small subset (empty is a case of small) and whose range includes any local cardinals with parents which are irregular atoms [or in $\Pi_*(C)$]. ρ must act as a permutation on the intersection of its domain with any clan or any set $K(D)$ (and it has no other domain elements). In addition, if a [C-]local bijection contains an atom in its domain which belongs to a litter with atomic parent in a clan $D[\leq_c C]$, then it also has the atomic parent in its domain.

Theorem (extension property): Any [C]-local bijection can be extended to a [C]-allowable permutation with no exceptions other than elements of the domain of the local bijection.

Proof of the extension property: Let ρ_0 be a [C]-local bijection.

For each pair of litters L, M in the same clan [$\leq_c C$], construct a bijection $\rho_{L,M}$ from $L \setminus \texttt{dom}(\rho_0)$ to $M \setminus \texttt{dom}(\rho_0)$. We claim that there is a unique [C]-allowable permutation extending ρ_0 and all maps $\rho_{L,\rho(L)^\circ}$, and moreover we assume that we already have this result for D-local bijections for all $D[<_c C]$ (inductive hypothesis).

We indicate how to compute $\rho(x)$ for any atom x by recursion along a strong support S of x in which x appears last, and which is such that every near-litter in S is a litter.

If L is a near-litter in the domain of S and we can compute $\rho([L])$, then we can compute $\rho(L)$ as the elementwise image of L under the union of

ρ_0 and $\rho_{L,\rho(L)^\circ}$, noting that we know $\rho(L)^\circ$ as soon as we know $\rho([L])$.

If L is a near-litter in the domain of S and its parent is an atom, then we can compute $\rho([L])$, either because the local cardinal $[L]$ is in the domain of ρ_0 or because the parent of L is a regular atom in the domain of ρ_0 (by the condition on parents of litters in the definition of local bijection).

If L is a near-litter in the domain of S and its parent is a set with D-strong support T appearing earlier in S, we have by inductive hypothesis already computed ρ at each element of T, and we can by inductive hypothesis (because $D <_c C$) produce a unique D-allowable permutation ρ' extending a D-local bijection ρ'_0 extending the restriction of ρ_0 to appropriate clans and their power sets and the computed values of ρ at atomic elements of T and at anomalies of near-litter elements of T, as well as relevant maps $\rho_{M,\rho'(M)^\circ}$ (M required to be included in a clan $\leq_c D$). There is no near-litter in T at which ρ and ρ' disagree, because if a litter $M \subseteq E$ (a clan) were the first bad one, ρ and ρ' would have the same values at an E-support of $[M]$, and so would send $[M]$ to the same value, and if ρ and ρ' disagreed at M, $\rho^{-1} \circ \rho'$ would have an exception mapped into or out of M (it is important here that ρ and ρ' agree at all anomalies of M), but ρ and ρ' must agree at all their exceptions. The value of $\rho'([L])$ is the only possible value for $\rho([L])$ if there is a suitable ρ, as both maps would be D-allowable and they agree on a D-support for $[L]$: we report $\rho'([L])$ as the computed value of $\rho([L])$.

Computations along different strong supports cannot disagree, basically because supports can be merged.

The map produced in this way is clearly $[C]$-allowable and has no exceptions other than domain elements of ρ_0.

elements of appopriate iterated power sets of clans are strongly symmetric: We show that we have $\mathcal{P}^n_+(C) = \mathcal{P}^n_*(C)$, when $\mathcal{P}^n_+(C)$ is defined.

We show this by induction on n.

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

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

Now let X be an arbitrary element of $\mathcal{P}^*(C)$ with strong support S. Choose any two distinct atoms x, y in C which do not belong to S and which both belong to the same element in S or both belong to no element in S. Define a local bijection fixing each atom in S and each anomalous element of a near-litter element 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 in 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 in S or both not in any element in 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 near-litter L in S included in C or in the set $C \setminus (S \cup \bigcup S)$ can be exchanged without moving the set X. But this means that X is of the form $Y \Delta \bigcup \Lambda$ or $C \setminus (Y \Delta \bigcup \Lambda)$, where Y is a set of atoms in S and Λ is a set of litters in S (near-litters in S can be replaced in this expression by the litters near them, adjoining their anomalies to Y), and this establishes our result for the case $n = 0$ in quite a strong form.

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

Let ρ be an $P^k_c(C)$-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 P^{k-1}_c-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 and T^* can be merged into a single strong support by imposing a suitable order on the union of their domains: ρ'_0 sends each atomic element in T^* and each anomalous element for a near-litter in T^* to its image under ρ. The extension ρ' of ρ'_0 obtained from the extension property sends each near-litter in T^* to its image under ρ: if this failed to be true, there would be a first near-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 nearby litter L° by this map (note that anomalies of the near-litter are fixed, so we can really think about the nearby litter). 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 P_c^{k-1}-support of x and are both $P_c^{k-1}(C)$-allowable (certainly $P_c^k(C)$-allowable implies $P_c^{k-1}(C)$-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 $P_c^k(C)$-support of X as required.

A crucial factor in the background here is that there will be no interaction between elements in T^* (elements appropriate for a $P_c^{k-1}(C)$ support) and elements in $S \setminus S^-$ (elements not appropriate for a $P_c^k(C)$-support) which conflicts with the conditions we impose on them.

The map τ defined and relationships to the exponential map established: We now know that $\Pi(C) = \Pi_*(C) \cup \bigcup_{D \in P_c^{-1}``\{C\}} \mathcal{P}_*^2(D)$ for any C for which the construction of π_C succeeds.

We now define the desired map τ from $\mathbb{T}$ to cardinals of the FM interpretation: $\tau(t) = |\mathcal{P}_*^2(\mathbb{C}(t))|_*$. We can verify the desired relation $\tau(P(t)) = 2^{\tau(t)}$ in the sense of the FM interpretation (subject to the demonstration below that the construction of π_C will not fail for any C).

We want to show that $2^{\tau(t)} = \tau(P(t))$.

This translates to $|\mathcal{P}_*^3(\mathbb{C}(t))|_* = |\mathcal{P}_*^2(\mathbb{C}(P(t)))|_*$. We show the inequalities in each direction. We let $\mathbb{C}(t)$ be represented as C, and $\mathbb{C}(P(t))$ be represented as $P_c(C)$.

$$|\mathcal{P}_*^3(C)|_* = |\mathcal{P}_*(\mathcal{P}_*^2(C))|_* \geq |\mathcal{P}_*(\mathcal{P}_*(\Pi(C)))|_* \geq |\mathcal{P}_*(\mathcal{P}_*(P_c(C)))|_* = |\mathcal{P}_*^2(P_c(C))|_*$$

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 $P_c(C)$ in the parent set of C.

$$|\mathcal{P}_*^2(P_c(C))|_* \geq |\mathcal{P}_*(\Pi(P_c(C))|_* \geq |\mathcal{P}_*(\mathcal{P}_*^2(C))|_* = |\mathcal{P}_*^3(C)|_*$$

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 $\mathcal{P}_*^2(C)$ into $\Pi(P_c(C))$.

Isomorphisms between natural models of type theory: Any $t \in \mathbb{T}$ for which $P^n(t)$ is defined corresponds to a natural model of the first n types of the simple typed theory of sets in which type i is implemented as $\mathcal{P}_*^{i+2}(\mathbb{C}(t))$ and the membership and equality relations are inherited from our FM interpretation. Note that the cardinality of type i in this natural model will be $\tau(P^i(t))$.

If $\leq_T$ restricted to $\{u : u \leq_T P^n(t)\}$ is isomorphic to $\leq_T$ restricted to $\{u : u \leq_T P^n(t')\}$, via an isomorphism sending t to t', then all the structure we have presented here to do with clans $\leq_c P_c^n(\mathbb{C}(t))$ is exactly isomorphic to the structure we have presented here to do with clans $\leq_c P_c^n(\mathbb{C}(t'))$, via framework maps. Now we observe that each type i is represented by $\mathcal{P}_*^{i+2}(\mathbb{C}(t))$ in the first structure and $\mathcal{P}_*^{i+2}(\mathbb{C}(t'))$ in the other, but these are respectively identical to $\mathcal{P}_+^{i+2}(\mathbb{C}(t))$ in the first structure and $\mathcal{P}_+^{i+2}(\mathbb{C}(t'))$, and $i + 2 \leq (n-1) + 2$ in the second, and information about $\mathcal{P}_+^{n+1}(C)$ depends only on π_D's for $D <_c P_c^n(C)$, since this set is hereditarily $P_c^n(C)$-symmetric.

It follows that if $\leq_T$ restricted to $\{u : u \leq_T P^n(t)\}$ is isomorphic to $\leq_T$ restricted to $\{u : u \leq_T P^n(t')\}$, via an isomorphism sending t to t', then the first order theory of the natural model of the first n types of the simple theory of types with base type $\mathbb{C}(t)$ (as defined in the FM interpretation) is the same as the first order theory of the natural model of the first n types of the simple theory of types with base type $\mathbb{C}(t')$ (as defined in the FM interpretation) because these models are actually isomorphic (in a way not visible to the FM interpretation):

the isomorphism witnessing this is the action on sets of the framework map sending $\mathbb{C}(t)$ to $\mathbb{C}(t')$, which interacts correctly with other relevant framework maps because of the way the framework maps are set up and used in the definitions of maps π_D.

A subtle point is that we can assume without loss of generality that the given isomorphism from $\leq_T$ restricted to $\{u : u \leq_T P^n(t)\}$ to $\leq_T$ restricted to $\{u : u \leq_T P^n(t')\}$ is a composition of the ones used in guiding the construction, using the point which we noted in the description of the construction that which isomorphisms between downward closed subtrees of the clans are used does not affect the computation of the maps π_D. This should make it clearer that the structures of the two sequences of iterated power sets in the FM interpretation that we are considering are exactly parallel.

coding functions and orbit specifications: For any support S for an object x, we define $\chi_x^S(\rho(S))$ 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_x^S *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 $t \in \mathbb{T}$ such that $S_\alpha \in \mathbb{C}(t)$ or $S_\alpha \subseteq \mathbb{C}(t)$ We use the notation S_t^α for the restriction of S to $\{S_\beta : \beta < \alpha \wedge \iota(S_\beta) <_T t\}$.
2. If S_α is an atom, $\sigma(S)(\alpha) = (1, \iota(S_\alpha), \beta)$, where 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, $\sigma(S)(\alpha) = (2, \iota(S_\alpha), \gamma)$, where $\pi_{\mathbb{C}(\iota(S_\alpha))}([S_\alpha]) = S_\gamma \in P_c(\mathbb{C}(\iota(S_\alpha)))$.
4. If S_α is a near-litter with parent a set, $\sigma(S)(\alpha) = (3, \iota(S_\alpha), g)$, where $S_\alpha \in \bigcup K(\mathbb{C}(\iota(S_\alpha)))$, g is a coding function, and $g(S^\alpha_{\iota(S_\alpha)}) = \pi_{\mathbb{C}(\iota(S_\alpha))}([S_\alpha])$.
5. If S_α is a near-litter with parent an irregular atom, $\sigma(S)(\alpha) = (4, \iota(S_\alpha))$.

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 an 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 $\mathcal{P}_*^n(C)$, where $P_c^{n-1}(C)$ exists, are of size μ, by an analysis of the size and number of orbits in the allowable permutations.

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

For the case $n = 1$, we describe the family of functions F_1^C which we will use: the domain orbits in the strong supports are all the ones consisting entirely of atoms in C and litters included in C. 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^C 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 C not in $S \cup \bigcup S$ iff $f_\alpha = 1$. We have seen above that all symmetric subsets of C can be described in this way. It is then evident that there are no more than $2^\kappa < \mu$ elements in F_1^C for each C, and that the union of the ranges of the elements of F_1^C is $\mathcal{P}_*^1(C)$.

In showing that $\mathcal{P}_*^n(C)$ has cardinality μ, $n > 1$, we assume that $\mathcal{P}_*^m(D)$ has been shown to have cardinality μ in all cases where $P_c^{m-1}(D) <_c P_c^{n-1}(C)$, by construction of families of coding functions F_m^D.

Let $X \in \mathcal{P}_*^n(C)$. Let S be a strong support for X. For each $x \in X$, choose an $P^{n-2}(C)$ support T which is an end-extension of the maximum $P^{n-2}(C)$ support included in S (if $n = 2$ we further cut down to just atoms and litters in C included in S). We write $T \leq_{n-2}^C S$ to express that T is an end extension of the maximal $P^{n-2}(C)$-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 $F_x \in F_{n-1}^C$. 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^C_n: for any U with $\sigma(U) = \sigma(S)$, $F_X(U) = \{F_x(V) : F_x \in F^C_{n-1} \wedge (\exists x \in X : (\exists T : F_x(T) = x \wedge T \leq^C_{n-2} S \wedge \sigma(V) = \sigma(T) \wedge V \leq^C_{n-2} U))\}$. We define F^C_n as the set of all such functions F_X for $X \in \mathcal{P}^n_*(C)$.

Note that F_X is exactly determined by $\sigma(S)$, the orbit specification of S, and the set $Z = \{F_x \in F^C_{n-1} : (\exists x \in X : (\exists T : F_x(T) = x \wedge T \leq^C_{n-2} S))\}$: $F_X = \{F_x(V) : F_x \in Z \wedge V \in \mathtt{dom}(F_x) \wedge V \leq^C_{n-2} U\}$. The set of possible values for Z is of size $< \mu$ because $|F^C_{n-1}| < \mu$ by inductive hypothesis and μ is a strong limit cardinal. The collection of all orbit specifications of $P^{n-1}_c(C)$-strong supports is also of size $< \mu$, because each such specification is a small structure built from elements of $\mathbb{T}$, ordinals less than κ, and coding functions belonging to F^D_m's already known to be of size $< \mu$ by our inductive hypotheses. Thus the set F^C_n of all F_X's constructed as indicated is of size $< \mu$ as desired.

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 $F_X(S) \subseteq X$. A general element of $F_X(S)$ is of the form $F_x(T')$ where there are $x \in X, T \leq^C_{n-2} S$, such that $F_x(T) = x$, and $T' \leq^C_{n-2} S$ (of course T and T' have the same specification). Supports extending $S \cup T$ on the one hand and $S \cup T'$ on the other 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 the support extending $S \cup T$ (suitably ordered) to the support extending $S \cup T'$, 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.

Completion of FM construction verification: We have verified that the FM construction has the claimed properties.

We commence the promised application to consistency of New Foundations: We now present the proof of the consistency of Quine's NF using the FM construction just described.

Definition of NF: New Foundations (NF) is the unsorted first order theory with equality and membership as primitive relations which has as its axioms exactly those formulas for which some assignment of types to the variables appearing in the formula yields an axiom of TST. More generally, a formula ϕ of the language of unsorted set theory for which some assignment of types to the variables appearing in ϕ yields a well-formed formula of the sorted language of TST is called a *stratified* formula. Notice that the axiom of comprehension of NF asserts the existence of $\{x : \phi\}$ for each stratified formula ϕ, whence this scheme is generally called the scheme of stratified comprehension.

For any formula ϕ in the language of TST which contains no free variables, we define ϕ^+ as a formula obtained by replacing each variable in ϕ with its image under a bijection on variables raising sorts by one. All such formulas are equivalent by renaming of bound variables, so we can speak of *the* formula ϕ^+ without danger of confusion. The axiom scheme of Ambiguity asserts $\phi \leftrightarrow \phi^+$ for each formula ϕ (note that ϕ^+ is a theorem of TST if ϕ is a theorem of TST, since this is clearly true of axioms and the $\cdot^+$ construction commutes with proof rules in a suitable sense, but the scheme of Ambiguity says something stronger than this). Specker proved in 1962 that NF is consistent iff TST is consistent with the scheme of Ambiguity.

The instance of the FM construction that we use is described: Let λ be a limit ordinal. Let $\mathbb{T}$ be the collection of all nonempty finite subsets of λ. For $A, B \in \mathbb{T}$, define $A \leq_T B$ as holding iff A is a downward extension of B: $B \subseteq A$ and any element of $A \setminus B$ is less than all elements of B. Define $\leq_T^w$ as the restriction to nonempty sets of the unique well-ordering on finite subsets of λ in which the empty set is maximal, in which nonempty finite sets with distinct maximal elements appear in the order of their maximal elements, and in which the order of two nonempty finite sets with the same maximum element is the same as the order on the two sets with the maximum elements dropped (so that $\{\alpha\}$ appears last among all finite sets with maximum α, because the empty set is maximal among all finite sets). We define $P(A)$ as

$A \setminus \{\mathtt{min}(A)\}$, if this set is nonempty. We write A_n for $P^n(A)$ when this is defined.

We apply the FM construction to the indicated tree: Applying the abstract FM construction above to this situation, we obtain a map τ from nonempty finite subsets of λ to Scott cardinals in the FM interpretation such that

1. $2^{\tau(A)} = \tau(A_1)$ if $|A| \geq 2$.
2. The first order theory of a natural model of TST_n with base type implemented by a set with cardinality $\tau(A)$ will be the same as the first order theory of a natural theory of TST_n with base type implemented by a set with cardinality $\tau(B)$ if $A \setminus A_{n+1} = B \setminus B_{n+1}$ and both A and B have at least $n+1$ elements, because the natural models in question will be isomorphic via a composition of the isomorphisms constructed from framework maps, as discussed in the abstract context above: it should be evident that if $A \setminus A_{n+1} = B \setminus B_{n+1}$, that the portions of the tree $<_T$ at and below A_n and B_n respectively are isomorphic in a suitable way (map each set $A_n \cup X$ to $B_n \cup X$ for finite sets $X \subseteq \lambda$ dominated by $\mathtt{min}(A_n) = \mathtt{min}(B_n)$). That is, the theory of a natural model of TST_n with base type of cardinality $\tau(A)$ will be determined by the smallest $n+1$ elements of A (as long as A has at least $n+1$ elements).

A map τ with these properties has elsewhere been called a *tangled web* of index λ.

The existence of a tangled web establishes Con(NF): Let τ be as in the previous paragraph.

Let Σ be a finite collection of formulas of the language of TST with no free variables. There will be an n such that Σ is a finite collection of formulas of the language of TST_n: fix such an n.

Define a partition of all subsets A of λ of cardinality $n+1$ in which the compartment of the partition to which A belongs is determined by the truth-values of the formulas in Σ in any natural model of TST_n (in the FM interpretation) with type 0 implemented by a set of cardinality $\tau(A)$. This partition has no more than $2^{|\Sigma|}$ compartments: Ramsey's theorem applies.

Choose a homogeneous set H for this partition of size $n+2$. In a natural model of TST_{n+1} with type 0 of cardinality $\tau(H)$, each formula $\phi \in \Sigma$ has the same truth value as ϕ^+, since the theory of types 0 to $n-1$ of this model is determined by the base type cardinality $\tau(H)$, and the theory of types 1 to n of this model is determined by the base type cardinality $2^{\tau(H)} = \tau(H_1)$, and because H is homogeneous, $H \setminus \{\mathtt{max}(H)\}$ and $H \setminus \{\mathtt{min}(H)\}$ lie in the same compartment of the indicated partition. Thus each formula $\phi \leftrightarrow \phi^+$ for $\phi \in \Sigma$ holds in the indicated model of TST_{n+1}, whence the restriction of the scheme of Ambiguity to $\phi \in \Sigma$ is consistent with TST_{n+1}, whence with TST, whence the entire scheme of Ambiguity is consistent with TST by compactness. It follows that NF is consistent by the theorem of Specker cited above.

This argument is an adaptation of Jensen's 1969 consistency proof of NFU along lines proposed by us in our 1995 paper.