

New Foundations is consistent

Lavinia Randall Holmes* and Sky Wilshaw

September 24, 2026

Abstract

In this paper we present a proof that Quine’s set theory New Foundations of [13], 1937 is consistent. We do this by constructing a model of the “tangled type theory” proposed by the first author and shown to be equiconsistent with NF in [6]. It is worth noting that the fact that the structure we describe is a model of tangled type theory has been formally verified in Lean by the second author and others (in [26]); this paper alludes occasionally to the Lean formalization (and relevant sections are parallel to it in structure), but the paper version does not depend on the formal verification.

Contents

1	Introduction	3
1.1	Introductory note	3
1.2	Acknowledgements	3
2	Development of relevant theories	5
2.1	The simple theory of types TST and TSTU	5
2.1.1	Typical ambiguity	6
2.1.2	Historical remarks	6
2.1.3	Some mathematics in TST	8
2.2	New Foundations and NFU	10
2.3	Tangled type theory TTT and TTTU	13
2.3.1	There are α -models for every α	17
2.4	An axiomatization of TST with finitely many templates	20
3	The model description	25
3.1	The abstract supertype framework	25
3.2	Preliminaries of the construction	26
3.3	Permutations and supports	28
3.4	Position functions	29
3.5	Hypotheses for the recursion	30

*the same individual as the M. R. Holmes cited in the bibliography

3.6	Machinery for enforcing extensionality in the model	33
3.7	Allowable permutations and supports	36
3.8	Model elements defined	38
4	Verification that the structure defined is a model	42
4.1	The Freedom of Action theorem	42
4.2	Freedom of Action for incomplete orbits	48
4.3	Counting orbits of supports: coding functions and specifications .	51
4.4	Types are of size μ	55
4.5	The structure is a model of predicative TTT	59
4.6	Impredicativity: verifying the axiom of union	60
5	Conclusions, extended results, and questions	62
5.1	Ways in which this resolution of the NF consistency problem is unexciting	62
5.2	Questions and a result about global choice-like statements	62
5.3	Strong unstratified axioms for NF derived from strong partition properties	63
5.4	Some esoteric questions answered	65
5.5	The question of consistency strength	65
5.6	The wide open question: what do models of NF look like in general?	66
5.7	Postscript	66

1 Introduction

1.1 Introductory note

We are presenting an argument for the consistency of Quine’s set theory New Foundations (NF) (see [13]). The consistency of this theory relative to the usual systems of set theory has been an open question since the theory was proposed in 1937, with added urgency after Specker showed in 1954 that NF disproves the Axiom of Choice (see [21]). Jensen showed in 1969 that NFU (New Foundations with extensionality weakened to allow urelements) is consistent and in fact consistent with Choice (and also with Infinity and with further strong axioms of infinity) (see [11]).

The first author showed the equiconsistency of NF with the quite bizarre system TTT (Tangled Type Theory) in 1995 [see [6], but this paper is not really recommended; the presentation of its results here is better], which gave a possible approach to a consistency proof. Since 2010, the first author has been attempting to write out proofs, first of the existence of a “tangled web of cardinals” in Mac Lane set theory [not defined in this paper] and then directly of the existence of a model of tangled type theory. These proofs have been difficult to read, insanely involved, and involve the sort of elaborate bookkeeping which makes it easy to introduce errors every time a new draft is prepared. We hope that the present version is more pleasant, though we know it is still forbidding. The second author (with the assistance of others initially) has formally verified the proof of the first author of the existence of a model of TTT (see [26]), and so of the consistency of New Foundations, in the Lean proof verification system, and in the process has suggested minor corrections and considerable formal improvements to the argument originally proposed by the first author. The second author reports that the formalized proof is still difficult to read and insanely involved with nasty bookkeeping. Both authors feel that there ought to be a simpler approach, but the existing argument at least strongly resists attempts at simplification.

All the theories mentioned here are discussed (and referenced) in the next section.

Any remarks in the paper in the first person singular may be attributed to the first author.

1.2 Acknowledgements

The first author thanks Robert Solovay, who read a number of early versions of a related argument for $\text{Con}(\text{NF})$ and offered just criticisms which I have tried to take to heart. I also thank Thomas Forster and Asaf Karagila, who have endured attempts of mine to present various relatives of this argument at length. Further, I thank the members of the group of students at Cambridge who attempted formalization of an earlier version of this proof in Summer 2022. We would like to thank Peter Lumsdaine for his work with the students in the summer project. Cambridge students who participated in the summer project

(including the second author) were funded partly by DPMMS, and Queens' College, but principally by the Wes and Margaret foundation, who had also supported work on New Foundations in the past, and we are grateful.

We appreciate extensive help from the anonymous referee.

2 Development of relevant theories

2.1 The simple theory of types TST and TSTU

We introduce a theory which we call the simple typed theory of sets or TST, a name favored by the school of Belgian logicians¹ who studied NF (*théorie simple des types*). This is not the same as the simple type theory of Ramsey and it is most certainly not Russell's type theory (see historical remarks below).

TST is a first order multi-sorted theory with sorts (types) indexed by the nonnegative integers. The primitive predicates of TST are equality and membership.

The type of a variable x is written $\mathbf{type}(x)$: this will be a nonnegative integer. A countably infinite supply of variables of each type is supposed. We provide a bijection $(x \mapsto x^+)$ from variables to variables of positive type satisfying $\mathbf{type}(x^+) = \mathbf{type}(x) + 1$.²

An atomic equality sentence $x = y$ is well-formed iff $\mathbf{type}(x) = \mathbf{type}(y)$. An atomic membership sentence $x \in y$ is well-formed iff $\mathbf{type}(x) + 1 = \mathbf{type}(y)$.

The axioms of TST are extensionality axioms and comprehension axioms.

The extensionality axioms are all the well-formed assertions of the shape $(\forall xy : x = y \leftrightarrow (\forall z : z \in x \leftrightarrow z \in y))$. For this to be well typed, the variables x and y must be of the same type, one type higher than the type of z .

The comprehension axioms are all the well-formed assertions of the shape $(\exists A : (\forall x : x \in A \leftrightarrow \phi))$, where ϕ is any formula in which A does not occur free.

The witness to $(\exists A : (\forall x : x \in A \leftrightarrow \phi))$ is unique by extensionality, and we introduce the notation $\{x : \phi\}$ for this object. Of course, $\{x : \phi\}$ is to be assigned type one higher than that of x ; in general, complex terms will have types as variables do.

The modification which gives TSTU (the simple type theory of sets with urelements) replaces the extensionality axioms with the formulas of the shape

$$(\forall xyw : w \in x \rightarrow (x = y \leftrightarrow (\forall z : z \in x \leftrightarrow z \in y))),$$

allowing many objects with no elements (called atoms or urelements) in each positive type. A technically useful refinement adds a constant $\emptyset^i$ of each positive type i with no elements: we can then address the problem that $\{x^i : \phi\}$ is not necessarily uniquely defined when ϕ is uniformly false by defining $\{x^i : \phi\}$ as $\emptyset^{i+1}$ in this case. The superscript is sometimes omitted from $\emptyset^i$ (and other constants parametrized by types) when the type can be deduced from context.

¹led by Maurice Boffa, to whom the first author owes a profound debt, for being harbored in his group at the beginning of my career, which made it possible for me to launch my research program without distractions attendant on looking for a job.

²We do *not* furnish our variables with type superscripts. One could follow the convention that if a variable has a natural number superscript, this determines its type, and that the effect of $+$ on a superscripted variable is to increment the superscript, but we do not expect variables to have superscripts.

2.1.1 Typical ambiguity

TST(U) exhibits a symmetry, often called typical ambiguity, which is important in the sequel.

If ϕ is a formula, define ϕ^+ as the result of replacing every variable x (free and bound) in ϕ with x^+ (and occurrences of $\emptyset^i$ with $\emptyset^{i+1}$ if this is in use). It should be evident that if ϕ is well-formed, so is ϕ^+ , and that if ϕ is a theorem, so is ϕ^+ (the converse is not the case). Further, if we define a mathematical object as a set abstract $\{x : \phi\}$ we have an analogous object $\{x^+ : \phi^+\}$ of the next higher type (this process can be iterated).

The axiom scheme asserting $\phi \leftrightarrow \phi^+$ for each closed formula ϕ is called the Ambiguity Scheme. Notice that this is a stronger assertion than is warranted by the symmetry of proofs described above.

2.1.2 Historical remarks

TST is not the type theory of the *Principia Mathematica* of Russell and Whitehead ([18]), though a description of TST is a common careless description of Russell's theory of types.

Russell described something like TST informally in his 1904 *Principles of Mathematics* ([17]). The obstruction to giving such an account in *Principia Mathematica* was that Russell and Whitehead did not know how to describe ordered pairs as sets. As a result, the system of *Principia Mathematica* has an elaborate system of complex types inhabited by n -ary relations with arguments of specified previously defined types, further complicated by predicativity restrictions (which are in effect cancelled by an axiom of reducibility). The simple theory of types of Ramsey eliminates the predicativity restrictions and the axiom of reducibility, but is still a theory with complex types inhabited by n -ary relations.

Russell noticed a phenomenon like the typical ambiguity of TST in the more complex system of *Principia Mathematica*, which he refers to as “systematic ambiguity”.

In 1914 ([25]), Norbert Wiener gave a definition of the ordered pair as a set (not the one now in use) and seems to have recognized that the type theory of *Principia Mathematica* could be simplified to something like TST, but he did not give a formal description. The theory we call TST was apparently first described by Tarski in the 1930s³.

It is worth observing that the axioms of TST look exactly like those of “naive set theory”, the restriction preventing paradox being embodied in the restriction of the language by the type system. For example, the Russell paradox is averted because one cannot have $\{x : x \notin x\}$ because $x \in x$ (and so its negation $\neg x \in x$) cannot be a well-formed formula.

³We regard it as an important feature of TST that the nature of type 0 is left completely undetermined, so we do not regard descriptions of arithmetic of order ω appearing about the same time as part of this history. Arithmetic of order ω is a similar typed theory but lacks the ambiguity of interest here.

It was shown around 1950 (in [10]) that Zermelo set theory proves the consistency of TST with the axiom of infinity; TST + Infinity has the same consistency strength as Zermelo set theory with separation restricted to bounded formulas.

2.1.3 Some mathematics in TST

We briefly discuss some mathematics in TST.

We indicate how to define the natural numbers. We use the definition of Frege (n is the set of all sets with n elements). 0 is $\{\emptyset\}$ (notice that we get a natural number 0 in each type $i + 2$; we will be deliberately ambiguous in this discussion, but we are aware that anything we define is actually not unique, but reduplicated in each type above the lowest one in which it can be defined). For any set A at all we define $\sigma(A)$ as $\{a \cup \{x\} : a \in A \wedge x \notin a\}$. This is definable for any A of type $i + 2$ (a being of type $i + 1$ and x of type i). Define 1 as $\sigma(0)$, 2 as $\sigma(1)$, 3 as $\sigma(2)$, and so forth. Clearly we have successfully defined 3 as the set of all sets with three elements, without circularity. But further, we can define $\mathbb{N}$ as $\{n : (\forall I : 0 \in I \wedge (\forall x \in I : \sigma(x) \in I) \rightarrow n \in I)\}$, that is, as the intersection of all inductive sets. $\mathbb{N}$ is again a typically ambiguous notation: there is an object defined in this way in each type $i + 3$.

The collection of all finite sets can be defined as $\bigcup \mathbb{N}$. The axiom of infinity can be stated as $V \notin \bigcup \mathbb{N}$ (where $V = \{x : x = x\}$ is the typically ambiguous symbol for the type $i + 1$ set of all type i objects). It is straightforward to show that the natural numbers in each type of a model of TST with Infinity are isomorphic in a way representable in the theory.

Ordered pairs can be defined following Kuratowski and a quite standard theory of functions and relations can be developed. Cardinal and ordinal numbers can be defined as Frege or Russell would have defined them, as isomorphism classes of sets under equinumerosity and isomorphism classes of well-orderings under similarity.

The Kuratowski pair $(x, y) = \{\{x\}, \{x, y\}\}$ is of course two types higher than its projections, which must be of the same type. There is an alternative definition (due to Quine in [12]) of an ordered pair $\langle x, y \rangle$ in TST + Infinity which is of the same type as its projections x, y . This is a considerable technical convenience but we will not need to define it here. Note for example that if we use the Kuratowski pair, the cartesian product $A \times B$ is two types higher than A, B , so we cannot define $|A| \cdot |B|$ as $|A \times B|$ if we want multiplication of cardinals to be a sensible operation. Let ι be the singleton operation and define $T(|A|)$ as $|\iota " A|$ (this is a very useful operation sending cardinals of a given type to cardinals in the next higher type which seem intuitively to be the same; also, it is clearly injective, so has a (partial) inverse operation T^{-1}). The definition of cardinal multiplication if we use the Kuratowski pair is then $|A| \cdot |B| = T^{-2}(|A \times B|)$. If we use the Quine pair this becomes the usual definition $|A| \cdot |B| = |A \times B|$. Use of the Quine pair simplifies matters in this case, but it should be noted that the T operation remains quite important (for example it provides the internally representable isomorphism between the systems of natural numbers in each sufficiently high type).

Note that the form of Cantor's Theorem in TST is not $|A| < |\mathcal{P}(A)|$, which would be ill-typed, but $|\iota " A| < |\mathcal{P}(A)|$: a set has fewer unit subsets than subsets. The exponential map $\exp(|A|) = 2^{|A|}$ is not defined as $|\mathcal{P}(A)|$, which would be one type too high, but as $T^{-1}(|\mathcal{P}(A)|)$, the cardinality of a set X such that

$|\iota^{\text{“}}X| = |\mathcal{P}(A)|$; notice that this is partial. For example $2^{|V|}$ is not defined (where $V = \{x : x = x\}$, an entire type), because there is no X with $|\iota^{\text{“}}X| = |\mathcal{P}(V)|$, because $|\iota^{\text{“}}V| < |\mathcal{P}(V)| \leq |V|$, and of course there is no set larger than V in its type.

2.2 New Foundations and NFU

In [13], 1937, Willard van Orman Quine proposed a set theory motivated by the typical ambiguity of TST described above. The paper in which he did this was titled “New foundations for mathematical logic”, and the set theory it introduces is called “New Foundations” or NF, after the title of the paper.⁴

Quine’s observation is that since any theorem ϕ of TST is accompanied by theorems $\phi^+, \phi^{++}, \phi^{+++}, \dots$ and every defined object $\{x : \phi\}$ is accompanied by $\{x^+ : \phi^+\}, \{x^{++} : \phi^{++}\}, \{x^{+++} : \phi^{+++}\}$, so the picture of what we can prove and construct in TST looks rather like a hall of mirrors⁵, we might reasonably (?) suppose that the types are all the same.

The concrete implementation follows. NF is the first order unsorted theory with equality and membership as primitive with an axiom of extensionality ($\forall xy : x = y \leftrightarrow (\forall z : z \in x \leftrightarrow z \in y)$) and an axiom of comprehension ($\exists A : (\forall x : x \in A \leftrightarrow \phi)$) for each formula ϕ in which A is not free which can be obtained from a formula of TST by dropping all distinctions of type. We give a precise formalization of this idea: provide a bijective map ($x \mapsto x^*$) from the countable supply of variables (of all types) of TST onto the countable supply of variables of the language of NF. Where ϕ is a formula of the language of TST, let ϕ^* be the formula obtained by replacing every variable x , free and bound, in ϕ with x^* . For each formula ϕ of the language of TST in which A is not free in ϕ^* and each variable x^* , an axiom of comprehension of NF asserts ($\exists A : (\forall x^* : x^* \in A \leftrightarrow \phi^*)$).

In the original paper, this is expressed in a way which avoids explicit dependence on the language of another theory. We do this in a way which is usual now but not necessarily identical to what is done in the original paper. Let ϕ be a formula of the language of NF. A function σ is a stratification of ϕ if it is a (possibly partial) map from variables to non-negative integers such that for each atomic subformula ‘ $x = y$ ’ of ϕ we have $\sigma(x) = \sigma(y)$ and for each atomic subformula ‘ $x \in y$ ’ of ϕ we have $\sigma(x) + 1 = \sigma(y)$. A formula ϕ is said to be stratified iff there is a stratification of ϕ . Then for each stratified formula ϕ of the language of NF and variable x we have an axiom ($\exists A : (\forall x : x \in A \leftrightarrow \phi)$). The stratified formulas are exactly the formulas ϕ^* up to renaming of variables.

For each stratified formula ϕ , there is a unique witness to

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

(uniqueness follows by extensionality) which we denote by $\{x : \phi\}$.

⁴There is a persistent rumor that Quine’s set theory in its original version was inconsistent. This is not the case (and in fact we show here that the system of the 1937 paper in which the system was introduced is consistent). The truth behind the rumor is that the system of the first edition of Quine’s book *Mathematical Logic* ([14]), in which proper classes were added to NF, was inconsistent; the inconsistency was corrected in the second edition of 1951, and the results of this paper show that the system of the 1951 book is consistent. We do not discuss the systems with proper classes further here.

⁵I have been asked whether the turn of phrase “hall of mirrors” here is my own or due to Quine himself or others in the interim: I actually do not know!

NF has been dismissed as a “syntactical trick” because of the way it is defined. It might go some way toward dispelling this impression to note that the stratified comprehension scheme is equivalent to a finite collection of its instances, so the theory can be presented in a way which makes no reference to types at all. This is a result of Hailperin ([4]), refined by others, and discussed in some detail in a subsequent subsection of this section, because it is relevant to the Lean formalization. One obtains a finite axiomatization of NF by analogy with the method of finitely axiomatizing von Neumann–Gödel–Bernays predicative class theory. It should further be noted that the first thing one does with any finite axiomatization is prove stratified comprehension as a meta-theorem, in practice, but it remains significant that the theory can be axiomatized with no reference to types at all. It is also worth noting that the collection of all typed instances of the Hailperin axioms is an axiomatization of TST, not a finite axiomatization, but an axiomatization with finitely many templates.

Jensen in [11], 1969 proposed the theory NFU which replaces the extensionality axiom of NF with

$$(\forall xyw : w \in x \rightarrow (x = y \leftrightarrow (\forall z : z \in x \leftrightarrow z \in y))),$$

allowing many atoms or urelements. One can reasonably add an elementless constant $\emptyset$, and define $\{x : \phi\}$ as $\emptyset$ when ϕ is false for all x .

Jensen showed that NFU is consistent and moreover NFU + Infinity + Choice is consistent. We will give an argument similar in spirit though not the same in detail for the consistency of NFU in the next section.

An important theorem of Specker ([22], 1962) is that NF is consistent if and only if TST + the Ambiguity Scheme is consistent. His method of proof adapts to show that NFU is consistent if and only if TSTU + the Ambiguity Scheme is consistent. Jensen used this fact in his proof of the consistency of NFU. We prove a version of Specker’s result using concepts from this paper below.

In [21], 1954, Specker had shown that NF disproves Choice, and so proves Infinity. At this point if not before it was clear that there is a serious issue of showing that NF is consistent relative to some set theory in which we have confidence. There is no evidence that NF is any stronger than TST + Infinity, the lower bound established by Specker’s result.

Note that NF or NFU supports the implementation of mathematics in the same style as TST, but with the representations of mathematical concepts losing their ambiguous character. The number 3 really is realized as the unique set of all sets with three elements, for example. The universe is a set and sets make up a Boolean algebra. Cardinal and ordinal numbers can be defined in the manner of Russell and Whitehead.

The apparent vulnerability to the paradox of Cantor is an illusion. Applying Cantor’s theorem to the cardinality of the universe in NFU gives $|\iota V| < |\mathcal{P}(V)| \leq |V|$ (the last inequality would be an equation in NF), from which we conclude that there are fewer singletons of objects than objects in the universe. The operation $(x \mapsto \{x\})$ is not a set function, and there is every reason to expect it not to be, as its definition is unstratified. The resolution of the Burali-

Forti paradox is also weird and wonderful in $NF(U)$, but would take us too far afield.

2.3 Tangled type theory TTT and TTTU

In [6], 1995, the first author described a reduction of the NF consistency problem to consistency of a typed theory, motivated by reverse engineering from Jensen's method of proving the consistency of NFU.

Let λ be a limit ordinal. It can be ω but it does not have to be.

In the theory TTT (tangled type theory) which we develop, each variable x is supplied with a type $\mathbf{type}(x) < \lambda$; we are provided with countably many distinct variables of each type.

For any formula ϕ of the language of TST and any strictly increasing sequence $\{s_i\}_{i \in \mathbb{N}}$ in λ , let ϕ^s be the formula obtained by replacing each variable of type i with a variable of type $s(i)$. To make this work rigorously, we suppose that we have a bijection from type i variables of the language of TST to type α variables of the language of TTT for each natural number i and ordinal $\alpha < \lambda$.

TTT is then the first order theory with types indexed by the ordinals below λ whose well-formed atomic sentences ' $x = y$ ' have $\mathbf{type}(x) = \mathbf{type}(y)$ and whose well-formed atomic sentences ' $x \in y$ ' satisfy $\mathbf{type}(x) < \mathbf{type}(y)$, and whose axioms are the sentences ϕ^s for each axiom ϕ of TST and each strictly increasing sequence s in λ . TTTU has the same relation to TSTU (with the addition of constants $\emptyset^{\alpha, \beta}$ for each $\alpha < \beta < \lambda$ such that $(\forall x^\alpha : x^\alpha \notin \emptyset^{\alpha, \beta})$ is an axiom).

It is important to notice how weird a theory TTT is. This is not cumulative type theory. Each type β is being interpreted as a power set of *each* lower type α . Cantor's theorem in the metatheory makes it clear that most of these power set interpretations cannot be honest.

A formulation of TTT which makes no reference to TST is possible.

The variable graph of ϕ we define as the graph whose nodes are variables in ϕ and whose edges are pairs of variables which occur together in atomic subformulas of ϕ . We say that variables u and v are connected (in ϕ) if there is a path from u to v in this graph. We say that the formula ϕ is connected iff all pairs of variables in ϕ are connected.

It is easy to show that every instance of TTT comprehension is equivalent to a connected instance of TTT comprehension; if $(\exists A : (\forall x : x \in A \leftrightarrow \phi^s))$ is an instance of TTT comprehension, add enough conjuncts to ϕ^s of the form $u_{s_i} \in v_{s_{i+1}} \vee \neg u_{s_i} \in v_{s_{i+1}}$ to get an equivalent formula which is connected, being sure to add one involving x if necessary (if new variables have to be added, they can be treated as parameters or a quantifier over each new variable can be applied to the entirety of the transformed ϕ^s).

TTT can be described as the first order theory with types indexed by the ordinals below λ whose well-formed atomic sentences ' $x = y$ ' have $\mathbf{type}(x) = \mathbf{type}(y)$ and whose well-formed atomic sentences ' $x \in y$ ' satisfy $\mathbf{type}(x) < \mathbf{type}(y)$, and whose axioms are the well-formed sentences

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

(extensionality axioms) and the connected⁶ well formed sentences

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

in which A does not occur free in ϕ and which satisfy the condition that for each pair of atomic subformulas $a \in b$ and $c \in d$ in the axiom, $\mathbf{type}(a) = \mathbf{type}(c) \leftrightarrow \mathbf{type}(b) = \mathbf{type}(d)$. We refer to these sentences as alternative comprehension axioms.

It should be evident that any alternative comprehension axiom is an instance of the comprehension axiom of TTT as defined above. Any connected comprehension axiom of TTT as formulated originally is an alternative comprehension axiom, and the connected instances of the comprehension scheme of TTT as originally defined imply all the other instances. So the alternative comprehension scheme is equivalent to the original one.

TTTU is obtained by weakening the extensionality axioms to allow atoms and adding empty set constants.

The alternative formulation has the merit of being simpler (though probably even more mysterious) and making no reference to the syntax or axioms of another theory. But the original formulation is useful to us in proving things about the theory, so we stick with it as the official formulation.

There is now a striking

Theorem 2.1 (Holmes). $\mathbf{TTT}(\mathbf{U})$ is consistent iff $\mathbf{NF}(\mathbf{U})$ is consistent.

Proof. Suppose $\mathbf{NF}(\mathbf{U})$ is consistent. Let (M, E) be a model of $\mathbf{NF}(\mathbf{U})$ (a set M with a membership relation E). Implement type α as $M \times \{\alpha\}$ for each $\alpha < \lambda$. Define $E_{\alpha, \beta}$ for $\alpha < \beta$ as $\{((x, \alpha), (y, \beta)) : xEy\}$. This gives a model of $\mathbf{TTT}(\mathbf{U})$. Empty sets in $\mathbf{TTTU}$ present no essential additional difficulties.

Suppose $\mathbf{TTT}(\mathbf{U})$ is consistent, and so we can assume we are working with a fixed model of $\mathbf{TTT}(\mathbf{U})$. Let Σ be a finite set of sentences in the language of $\mathbf{TST}(\mathbf{U})$. Our first aim is to construct a model of $\mathbf{TST}(\mathbf{U})$ in which $\phi \leftrightarrow \phi^+$ holds for each sentence $\phi \in \Sigma$.

1. Let n be the smallest natural number such that all variables which occur in Σ have type less than n . We define a partition of the n -element subsets of λ . Each $A \in [\lambda]^n$ is put in a compartment determined by the truth values of the sentences ϕ^s in our model of $\mathbf{TTT}(\mathbf{U})$, where $\phi \in \Sigma$ and $\mathbf{rng}(s \upharpoonright \{0, \dots, n-1\}) = A$: there are no more than $2^{|\Sigma|}$ compartments. By Ramsey's theorem, there is an infinite set $H \subseteq \lambda$ homogeneous for this partition which includes the range of a strictly increasing sequence h .
2. Any strictly increasing sequence s determines a model of $\mathbf{TST}(\mathbf{U})$ in which each type i is implemented as type $s(i)$ of our fixed model of $\mathbf{TTT}(\mathbf{U})$ and membership in the derived model of $\mathbf{TST}(\mathbf{U})$ of type i objects in type $i+1$ objects is implemented as membership in the fixed model of $\mathbf{TTT}(\mathbf{U})$ of type $s(i)$ objects in type $s(i+1)$ objects.

⁶The condition that ϕ be connected is actually not needed, but verifying that the resulting scheme is equivalent to TTT comprehension involves moderately complex reasoning about reduction of formulas to connected forms which we have no need to give here.

3. This model determined in this way by h satisfies $\phi \leftrightarrow \phi^+$ for each $\phi \in \Sigma$. This achieves our first aim.
4. But this implies by compactness that the full Ambiguity Scheme $\phi \leftrightarrow \phi^+$ is consistent with TST(U). We could at this point say that NF(U) is consistent by the 1962 result of Specker ([22]), but we will continue and indicate how to prove this directly (our second aim).
5. We pause to add a useful notion to our logic. The Hilbert symbol is a primitive term construction $(\epsilon x : \phi)$ (same type as x) with axiom schemes $(\exists x : \phi) \rightarrow \phi[(\epsilon x : \phi)/x]$ and $(\forall x : \phi \leftrightarrow \psi) \rightarrow (\epsilon x : \phi) = (\epsilon x : \psi)$ which cannot appear in instances of comprehension. We do not define the quantifiers in terms of the Hilbert symbol, because the quantifiers do need to appear in instances of comprehension.
6. If we have a model of TST(U) augmented with a Hilbert symbol which satisfies Ambiguity (for all formulas, including those which mention the Hilbert symbol) then we can readily get a model of NF(U), by constructing a term model using the Hilbert symbol in the natural way, then identifying all terms with their type-raised versions (an element of the term model is a class of terms $(\epsilon x : \phi)$ in the language of TST(U) which is obtained from a single such term by all possible upward or downward shifts of type by a constant amount: we write this $[(\epsilon x : \phi)]$).
7. All statements in the resulting type-free theory can be decided: the truth value of a closed atomic formula $[(\epsilon x : \phi)] R [(\epsilon y : \psi)]$ in the model of NF(U) is determined by replacing each side of the formula with one of its elements in such a way as to get a well-typed formula, which will always be possible if types are taken high enough, and reading the truth value of the resulting formula from the model of TST(U) with Ambiguity and Hilbert symbol; R is either $=$ or $\in$ [note that the relation interpreting equality in the term model is a nontrivial equivalence relation]. Truth values of more complex statements are determined in standard ways. That this makes axioms of NF(U) true should be clear.
8. Now observe that a model of TTT(U) can readily be equipped with a Hilbert symbol if this creates no obligation to add instances of comprehension containing the Hilbert symbol (use a well-ordering of the set implementing each type to interpret a Hilbert symbol $(\epsilon x : \phi)$ in that type as the first x such that ϕ or as a default object of the appropriate type if there is no such x), and the argument in paragraphs (1)–(4) adapts to show consistency of TST(U) plus Ambiguity with the Hilbert symbol.
9. We remark that our method of proving (a version of) the Specker ambiguity result is distinctly different from the original method. Strictly speaking, we have only proved Specker's result for TST(U) + Ambiguity + existence of a Hilbert symbol, but that is all we need for our result.

10. From this it follows that $\text{NF}(\mathbf{U})$ is consistent under the hypothesis that $\text{TTT}(\mathbf{U})$ is consistent, which is what was to be shown.

□

Theorem 2.2 (essentially due to Jensen). NFU is consistent.

Proof. It is enough to exhibit a model of TTTU . Represent type α as $V_{\omega+\alpha} \times \{\alpha\}$ for each $\alpha < \lambda$ ($V_{\omega+\alpha}$ being the rank indexed by $\omega + \alpha$ in the usual cumulative hierarchy). Define $\in_{\alpha,\beta}$ for $\alpha < \beta < \lambda$ as

$$\{((x, \alpha), (y, \beta)) : x \in V_{\omega+\alpha} \wedge y \in V_{\omega+\alpha+1} \wedge x \in y\}.$$

This gives a model of TTTU in which the membership of type α in type β is defined in such a way that each (y, β) with $y \in V_{\omega+\beta} \setminus V_{\omega+\alpha+1}$ is interpreted as an urelement over type α .

Our use of $V_{\omega+\alpha}$ enforces Infinity in the resulting models of NFU (note that we did not have to do this: if we set $\lambda = \omega$ and interpret type α using V_α we prove the consistency of NFU with the negation of Infinity). It should be clear that Choice (stated in a stratified way) holds in the models of NFU eventually obtained if it holds in the ambient set theory.

This shows in fact that mathematics in NFU is quite ordinary (with respect to stratified sentences), because mathematics in the models of TSTU embedded in the indicated model of TTTU is quite ordinary. The notorious ways in which NF evades the paradoxes of Russell, Cantor and Burali-Forti can be examined in actual models and we can see that they work and how they work (since they work in NFU in the same way they work in NF). □

Of course Jensen did not phrase his argument in terms of tangled type theory. Our contribution here was to reverse engineer from Jensen's original argument for the consistency of NFU an argument for the consistency of NF itself, which requires additional input which we did not know how to supply (a proof of the consistency of TTT itself). An intuitive way to say what is happening here is that Jensen noticed that it is possible to skip types in a certain sense in TSTU in a way which is not obviously possible in TST itself; to suppose that TTT might be consistent is to suppose that such type skipping is also possible in TST , and the description of TTT is an explicit realization of what is meant by this. We leave it to the historically inclined reader to examine Jensen's original proof and determine the exact relationship between the approaches.

2.3.1 There are α -models for every α

This section is devoted to adapting a result of Jensen to our context. Jensen proved in [11] that for any infinite ordinal α of the metatheory, there is a model of $\text{NFU} + \text{Infinity}$ in which there is a well-ordering of type α . This is of particular interest in the case $\alpha = \omega$, which establishes that the natural numbers of a model can correspond precisely to the natural numbers of the metatheory.

Nothing in our construction of a model of TTT below depends on this section: the fact that an existence of an ω -model of NF follows from our main result is significant for extensions of our results and uses this theorem.

A relation R of one of the theories under discussion is a set of ordered pairs⁷ of the theory in question; note that it corresponds to a set of ordered pairs of the metatheory with the same field. If this relation corresponding to R is a well-ordering, we can talk about its order type both in terms of the theory under consideration and in terms of the metatheory, though ordinals of the theory under consideration and the metatheory may not be the same.

In any of the theories under discussion, an ordinal is taken to be an equivalence class of well-orderings under similarity. The relations which belong to an ordinal of one of these theories may fail to be well-orderings in terms of the metatheory. If α is an ordinal of the metatheory, we say that α is *implemented* in a type of one of the theories under discussion exactly if there is an ordinal in that type (in the sense of the theory) whose elements in fact correspond to well-orderings of order type α in the sense of the metatheory, in the sense outlined in the previous paragraph. It is straightforward to show that if α is implemented in a type, the ordinal implementing it is uniquely determined.

Theorem 2.3. For any infinite ordinal α of the metatheory, if there is a model of $\text{TTT}(\text{U})$ in which the ordinal α has an implementation α^s for each strictly increasing sequence s of types, by which we mean an ordinal in type s_4 whose elements of type s_3 correspond to well-orderings of order type α of subsets of type s_0 , using types s_1 and s_2 for the intermediate pairing machinery, and which is such that the ordinal λ indexing the types is a strong limit cardinal of cofinality $> |\alpha|^{|\alpha|}$, then there is a model of $\text{NF}(\text{U})$ in which the ordinal α is implemented, which further satisfies [after neglect of types] any sentence in the language of $\text{TST}(\text{U})$ whose translations via all increasing sequences of types are satisfied in the model of $\text{TTT}(\text{U})$.

Proof. The proof is a modified version of the proof given above for consistency of $\text{NF}(\text{U})$, using a more powerful partition principle. We do not claim originality, referring the reader to [11], but our presentation is not the same, and it is adapted to talk of tangled type theory.

The Erdős-Rado partition principle as usually formulated says $\exp_r(\kappa)^+ \rightarrow (\kappa^+)_\kappa^{r+1}$: a partition of the collection of $r + 1$ element subsets of $\exp_r(\kappa)^+$ into κ compartments has a homogeneous set of size κ^+ . It follows from this that for any $r \geq 1$, a partition of the r element subsets of $\exp_{r-1}(\kappa)^+$ into κ

⁷The discussion here presupposes use of the Kuratowski ordered pair; it could be adapted to the use of another pair definition.

compartments has a homogeneous set of size κ^+ . And further it follows that for any $r \geq 1$, a partition of the r element subsets of $\exp_r(\kappa)$ into κ compartments has a homogeneous set of size κ^+ . It is this modified partition principle that we find convenient to use in the argument which follows.

We extend the language of $\text{TTT}(\mathbf{U})$ to include constant terms α^s and β^s for each ordinal $\beta < \alpha$ of the metatheory for each sequence s as described above. We similarly extend the language of $\text{TST}(\mathbf{U})$ (in this case only needing constants α^i and β^i for $\beta < \alpha$ for each $i \geq 4$) and note that the translation of a sentence of $\text{TST}(\mathbf{U})$ with language augmented with these ordinal constants into terms of a strictly increasing sequence of types in $\text{TTT}(\mathbf{U})$ is straightforward, and we allow use of these constants in the comprehension axiom of $\text{TTT}(\mathbf{U})$ as usual. We further include the Hilbert symbol in our language, with the same restriction stated above that the Hilbert symbol may not appear in instances of comprehension, noting our intention to construct a term model.

We consider partitions of $[\lambda]^n$ determined by the set Σ_n of *all* closed Hilbert symbols (“closed” meaning the symbol contains no parameters) of the form $(\epsilon\beta^i : \beta^i < \alpha^i \wedge \phi)$ in the language of $\text{TST}(\mathbf{U})$ augmented as indicated above and using only types less than n , with the compartment in the partition containing $A \in [\lambda]^n$ determined by the ordinal values $< \alpha$ [or default value, where we assume the default value is not an ordinal] assigned to all of these terms in our model of $\text{TTT}(\mathbf{U})$ when types $0 - (n - 1)$ are replaced by the types in A in increasing order in each of the Hilbert symbols in Σ_n . Here we are speaking loosely in that we say an ordinal $\gamma < \alpha$ is assigned as value to $(\epsilon\beta^i : \beta^i < \alpha^i \wedge \phi)$ when its value is in fact γ^i : it should be clear that this is unproblematic and easier to say.

There are no more than $|\alpha|$ such Hilbert symbols, and no more than $|\alpha|^{|\alpha|}$ compartments in each of these partitions.

Note that assignment of values to all such Hilbert symbols induces a complete assignment of truth values to sentences of the extended language:

$$(\epsilon\beta^i : \beta^i < \alpha \wedge \beta^i < 2 \wedge (\beta = 1 \leftrightarrow \phi))$$

codes a sentence ϕ into Σ_n if ϕ mentions only types below $n > i$, because the assignment of ordinal value to this term corresponds to the assignment of truth value to ϕ .

The Erdős–Rado theorem (as modified) allows us to conclude that for any infinite $\mu < \lambda$, we have a homogeneous set of size μ^+ for the given partition in $\exp^n(\mu) < \lambda$. Further, we can find a single assignment of values to Hilbert symbols as described above which is associated with arbitrarily large homogeneous sets below λ , because the cofinality of λ is greater than $|\alpha|^{|\alpha|}$, the cardinality of the set of compartments in the partition, corresponding precisely to the cardinality of the set of such assignments of values. If we have chosen such an assignment of values for a specific value of n , we can further choose such an assignment of values for $n + 1$ which extends it: for any μ , choose a homogeneous set for the partition of $[\lambda]^n$ realizing the fixed assignment of values and of cardinality $\nu > \exp^{n+1}(\mu)$, then find a homogeneous set of size μ^+ for

the partition of $[\lambda]^{n+1}$ inside this homogeneous set, which will determine an assignment of values to Hilbert symbols using $n + 1$ types extending the fixed one chosen for Hilbert symbols using n types. Thus there are arbitrarily large homogeneous sets for assignments of values to Hilbert symbols using $n + 1$ types extending the fixed one for Hilbert symbols using n types. Thus there is a fixed assignment of values to Hilbert symbols using $n + 1$ types which extends the given fixed assignment for n types and which is associated with arbitrarily large homogeneous sets, by the already stated considerations of cofinality. Thus we can proceed through ω steps and find an assignment of values to all Hilbert symbols with the property that its restriction to any fixed number n of types is realized in arbitrarily large homogeneous sets below λ .

Thus we have constructed a complete assignment of values to such Hilbert symbols in the extended TST(U) language, and thus a complete assignment of truth values to all sentences [including membership and equality statements between Hilbert symbols of completely arbitrary form], which is ambiguous [invariant under displacement of types] and consistent with any statement holding at all sequences of types in the model of TTT(U). Thus for example we will get extensionality if the model of TTTU is actually a model of TTT. In fact, for any fixed n all these assignments for Hilbert symbols mentioning only types less than n will be realized at some sequence of types in the model of TTT(U) [verifying that this assignment of values is consistent and in effect a complete description of a term model of NF(U)].

Further, the model of NFU thus described in effect (in the previous paragraph) – a term model determined by our assignment of values to Hilbert symbols – contains no ordinal less than what it calls α other than the ordinals which it calls β for $\beta < \alpha$ in the metatheory, and the natural well-ordering on these ordinals is realized in the model of NF(U) and corresponds to a well-ordering of order type α in the metatheory. The reason for this is that the model of NFU is a term model, and any term denoting an ordinal less than α can be slightly modified to the exact form used in defining the partitions⁸ and the process will assign that term a concrete value corresponding to an ordinal $\beta < \alpha$ of the metatheory, so the order on the ordinals less than the ordinal called α in the model of NF(U) has order type α in the metatheory.

□

⁸subject to making a technical requirement on the Hilbert symbol: we require (as is clearly possible) that if $(\epsilon x : \phi)$ is an ordinal constant β^s for β less than α , then $(\epsilon x : \phi) = (\epsilon x : x < \alpha^s \wedge \phi)$

2.4 An axiomatization of TST with finitely many templates

We discuss a finite axiomatization of NF derived from that of Hailperin (the statement of the axioms is taken from an implementation of Hailperin's axiom set of [4] in Metamath ([2]), and there are minor changes from the original formulation; the proofs are our own), making the important observation that it actually provides us with an axiomatization of TST with finitely many axiom templates (in the sense that each axiom is a type shifted version of one of a finite set of axioms). Notation introduced in this section is not used in the rest of the paper, and nothing in subsequent sections except a brief remark in the last paragraphs of section 4 depends on anything here. This finite axiomatization is however used in the Lean formalization, which is why we review it in detail here.

The finite axiomatization of NF takes this form (the definitions inserted are ours, and we have modified the order of the axioms to make the definitions work sensibly). We also present this as an axiomatization of TST with finitely many templates, with the proviso that each typed form of each axiom is asserted:

extensionality axiom: $(\forall x : (\forall y : (\forall z : z \in x \leftrightarrow z \in y) \rightarrow x = y))$

anti-intersection axiom: $(\forall xy : (\exists z : (\forall w : w \in z \leftrightarrow \neg(w \in x \wedge w \in y))))$

singleton axiom: $(\forall x : (\exists y : (\forall z : z \in y \leftrightarrow z = x)))$

definition: $\{x\}$ denotes for each x the set whose only element is x , whose existence is provided by the previous axiom. We define $\iota(x)$ as $\{x\}$ and define $\iota^1(x)$ as $\iota(x)$ and $\iota^{n+1}(x)$ as $\{\iota^n(x)\}$, for each concrete natural number n .

cardinal one axiom: $(\exists x : (\forall y : y \in x \leftrightarrow (\exists z : (\forall w : w \in y \leftrightarrow w = z))))$

definition: We define 1 as the set of all singletons, provided by the previous axiom.

definition: $x|y$ denotes the set z whose existence is provided by the anti-intersection axiom: $z \in x|y \leftrightarrow \neg(z \in x \wedge z \in y)$. We define x^c as $x|x$. We define $x \cap y$ as $(x|y)^c$. We define $x \cup y$ as $x^c|y^c$. We define V as $1|1^c$ (noting that it is straightforward to prove $x|1^c = V$ for any x , since this is the universal set). We define $\{x, y\}$ as $\{x\} \cup \{y\}$. We define (x, y) as $\{\{x\}, \{x, y\}\}$. We define (x, y, z) as $(\{\{x\}\}, (y, z))$. More generally, we define $(x_1, \dots, x_n)$ as $(\iota^{2n-4}(x_1), (x_2, \dots, x_n))$ for $n > 2$ a concrete natural number. [The treatment of n -tuples is what makes this axiomatization singularly awful].

cross product axiom: $(\forall x : (\exists y : (\forall z : z \in y \leftrightarrow (\exists wt : z = (w, t) \wedge t \in x))))$

definition: We define $V \times x$ as the set introduced by the previous axiom: $z \in V \times x \leftrightarrow (\exists wt : z = (w, t) \wedge t \in x)$. Note that $V \times V$ is the set of all ordered pairs.

converse axiom: $(\forall x : (\exists y : (\forall zw : (z, w) \in y \leftrightarrow (w, z) \in x)))$

definition: For any set R , we define R^{-1} as the intersection of $V \times V$ with a set introduced by the previous axiom:

$$(\forall zw : (z, w) \in R^{-1} \leftrightarrow (w, z) \in R) \wedge (\forall u : u \in R^{-1} \rightarrow (\exists zw : (z, w) = u)).$$

definition: We define $x \times V$ as $(V \times x)^{-1}$ and $x \times y$ as $(x \times V) \cap (V \times y)$.

definitions: We define $\iota^2 x$ as $(x \times x) \cap 1$. This is the image of x under the double singleton operation.

Note that $\iota^2 V$ is the equality relation. Define $\iota^{2(n+1)} x$ as $\iota^2 (\iota^{2n} x)$.

We define $x_1 \times x_2 \dots \times x_n$ as $\iota^{2(n-2)} x \times (x_2 \times \dots \times x_n)$.

We define V^1 as V and V^{n+1} as $\iota^{2(n-1)} V \times V^n$, for each concrete n . V^n is the set of all n -tuples.

singleton image axiom: $(\forall x : (\exists y : (\forall zw : (\{z\}, \{w\}) \in y \leftrightarrow (z, w) \in x)))$.

definition: We define R^ι for any set R as the intersection of a set provided by the previous axiom with 1×1 . R^ι is the set whose members are exactly the ordered pairs $(\{z\}, \{w\})$ such that $(z, w) \in R$. Let R^{ι^1} be defined as R^ι and $R^{\iota^{n+1}}$ be defined as $(R^{\iota^n})^\iota$.

insertion two axiom: $(\forall x : (\exists y : (\forall zwt : (z, w, t) \in y \leftrightarrow (z, t) \in x)))$

We define $I_2(x)$ as the intersection of a set provided by the previous axiom with $V \times V \times V$.

insertion three axiom: $(\forall x : (\exists y : (\forall zwt : (z, w, t) \in y \leftrightarrow (z, w) \in x)))$

We define $I_3(x)$ as the intersection of a set provided by the previous axiom with $V \times V \times V$.

definition: It seems natural to define $I_1(R)$ as $\iota^2 V \times R$; this requires no new axiom.

definition: Define $I_{1,n}(R)$ as $\iota^{2(n-1)} V \times R$: this is correct for prepending all possible initial projections to an n -tuple. Then define $I_{1,n}^1(R)$ as $I_{1,n}(R)$ and define $I_{1,n}^{m+1}(R)$ as $I_{1,n+m}(I_{1,n}^m(R))$: this takes into account the fact that the tuples get longer. The intention here is that $I_{1,n}^m$ represents the result of applying the insertion one operation to an n -ary relation m times, yielding an $(n+m)$ -ary relation.

Define $I_{2,n}(R)$ and $I_{2,n}^1(R)$ as $I_2(R) \cap V^{n+1}$ [one effect of the intersection with V^{n+1} is to restrict second projections of tuples in $I_2(R)$ to appropriately iterated singletons], and define $I_{2,n}^{m+1}(R)$ as $I_{2,n+m}(I_{2,n}^m(R))$: this takes into account the fact that the tuples get longer. The intention here is that $I_{2,n}^m$ represents the result of applying the insertion two operation to an n -ary relation m times, yielding an $(n+m)$ -ary relation.

type lowering axiom: $(\forall x : (\exists y : (\forall z : z \in y \leftrightarrow (\forall w : (w, \{z\}) \in x))))$

definition: We define $\text{TL}(x)$ by $(\forall z : z \in \text{TL}(x) \leftrightarrow (\forall w : (w, \{z\}) \in x))$. This is a very strange operation!

We define $\iota^{-1}x$ as $\text{TL}(V \times x)$. This is the set of all elements of singletons belonging to x . We can then define ιx , the elementwise image of x under the singleton operation, as $\iota^{-1}(\iota^2(x))$.

Further, we define $\iota^{-(n+1)}(x)$ as $\iota^{-1}(\iota^{-n}(x))$ for each concrete natural number n , and $\iota^{n+1}(x)$ as $\iota(\iota^n(x))$.

We develop an important operation step by step.

$$\text{TL}(R) = \{z : (\forall w : (w, \{z\}) \in R)\}$$

$$\text{Dually, } (\text{TL}(R^c))^c = \{z : (\exists w : (w, \{z\}) \in R)\}$$

$$\text{Now } (\text{TL}((R^\iota)^c))^c = \{z : (\exists w : (w, \{z\}) \in R^\iota)\},$$

which is the same as $(\text{TL}((R^\iota)^c))^c = \{z : (\exists w : (\{w\}, \{z\}) \in R^\iota)\}$ because all elements of the domain of R^ι are singletons,

$$\text{which is the same as } (\text{TL}((R^\iota)^c))^c = \{z : (\exists w : (w, z) \in R)\}$$

so we define $\text{rng}(R)$ as $(\text{TL}((R^\iota)^c))^c$, and define $\text{dom}(R)$ as $\text{rng}(R^{-1})$.

subset axiom: $(\exists x : (\forall yz : (y, z) \in x \leftrightarrow (\forall w : w \in y \rightarrow w \in z)))$

We define $[\subseteq]$ as the intersection of a set provided by the previous axiom with $V \times V$: $[\subseteq]$ is the set of all ordered pairs (x, y) such that $x \subseteq y$.

We define $[\in]$ as $[\subseteq] \cap (1 \times V)$.

This is not our favorite finite axiomatization of NF (or our favorite finite template axiomatization of TST) but it is the one verified in the Lean formalization and also basically the oldest one, so we present a verification of it in outline at least.

What we need to do is verify that $\{x : \phi\}$ exists for each formula ϕ of the language of TST, to ensure that comprehension holds. It is enough to do this for all formulas satisfying a convenient requirement: in each subformula $(\forall t_i : \psi)$ or $(\exists t_i : \psi)$ of ϕ we require that all occurrences of t_i in ψ be free and the index of t_i be less than the index of any other variable appearing free in ψ . It should be clear that we can do this without loss of generality. We do this by induction on the structure of formulas.

$$\{x : \neg\phi\} \text{ is } \{x : \phi\}^c.$$

$$\{x : \phi \wedge \psi\} \text{ is } \{x : \phi\} \cap \{x : \psi\}.$$

Now we have the much more complex task of analyzing

$$\{t_n : (\forall t_i : \phi(t_1, \dots, t_n))\}.$$

Choose a type τ' higher than the type of each variable, bound or free, in ϕ . Where the type of t_i is τ_i , we define x_i as $\iota^{\tau' - \tau_i}(t_i)$: we construct

$$\{t_n : (\forall t_i : \phi(t_1, \dots, t_n))\}$$

by defining manipulations which allow us to build sets $\{(x_1, \dots, x_n) : \phi^*(x_1, \dots, x_n)\}$ in which all the variables are of the same type. We write ϕ^* to suggest that the formula ϕ must be transformed to effect our change of variables: $t_i = t_j$ is equivalent to $x_i = x_j$, and $t_i \in t_j$ is equivalent to $(x_i, x_j) \in [\in]_{\iota^{\tau' - \tau_j}}$ (the reader will see that we use this representation below, though embedded in larger tuples). A quantifier over t_i is equivalent to a quantifier over x_i restricted to $\iota^{\tau' - \tau_i} V$.

For any such representation, we have a type signature

$$\iota^{\tau' - \tau_1} V \times \dots \times \iota^{\tau' - \tau_n} V.$$

We abbreviate this as τ^* .

The set $\{(x_1, \dots, x_n) : x_1 = x_n\}$ is obtained as $I_{2,2}^{n-2}(\iota^2 V) \cap \tau^*$.

The set $\{(x_1, \dots, x_n) : x_1 = x_i\}$ ($i < n$) is obtained as $I_{2,3}^{i-2}(I_3(\iota^2 V)) \cap \tau^*$.

We then can represent the set $\{(x_1, \dots, x_n) : x_i = x_j\}$ (wlog $i < j$) as $I_{1,n-i+1}^{i-1}(\{(x_1, \dots, x_{j-i+1}) : x_1 = x_{j-i+1}\}) \cap \tau^*$. *NOTE to referee: typo corrected here, addressing their objection*

The set $\{(\{x_1\}, x_2) : x_1 \in x_2\}$ has already been defined above as $[\in]$.

The set $\{(x_1, \dots, x_n) : (\exists uv : x_i = \iota^{k+1}(u) \wedge x_j = \iota^k(v) \wedge u \in v)\}$ is

$$\mathbf{rng}^2(I_3([\in]_{\iota^{k+2n}})) \cap \{(x_1, \dots, x_{n+2}) : x_1 = x_{i+2}\} \cap \{(x_1, \dots, x_{n+2}) : x_2 = x_{j+2}\}.$$

Note to referee: correction gratefully accepted in formula just above. This handles all representations of membership statements between x_i 's in the framework we are using.

The set $\{(x_1, \dots, x_n) : (\exists x_1 : (x_1, \dots, x_n) \in R)\}$ is representable as $I_{1,n-1}(\mathbf{rng}(R))$.

This only allows us to quantify over the lowest numbered variable. This is actually sufficient. We have renamed bound variables so that all bound variables with different binders are distinct and in any subformula $(\exists y : \psi)$ y will have lower index than any variable free in ψ . Then the quantified variable might not be x_1 , but if it is x_k , we can use the representation $I_{1,n-k}^k(\mathbf{rng}^k(R))$: strip off the first $k - 1$ variables, which are dummies already quantified over, quantify over x_k (stripping off one more variable) and put k dummies back, as it were.

This gives sufficient machinery to handle representations of all sentences in the language of TST (or stratified sentences in the language of NF) in the format we are using. However, the final representation of $\{t_n : \phi(t_1, \dots, t_n)\}$ obtained in this way will be of the form $R = \{(\iota^{\tau' - \tau_1}(t_1), \dots, \iota^{\tau' - \tau_n}(t_n)) : \phi(t_1, \dots, t_n)\}$ where some of the t_i 's are bound variables (all values of which will appear), and some are parameters. Let P be the set of i such that t_i is a parameter.

The final computation of $\{t_n : \phi(t_1, \dots, t_n)\}$ will be as

$$\iota^{-(\tau' - \tau_n)} \left(\mathbf{rng}^{n-1} \left(R \cap \bigcap_{i \in P} \{(x_1, \dots, x_n) : x_i = \iota^{\tau' - \tau_i}(t_i)\} \right) \right)$$

where a final bit of computation must be exhibited: $\{(x_1, \dots, x_n) : x_i = t\}$ is realized as $I_{1,2}^{i-1}(\{\iota^{2(n-i-1)}(t)\} \times V) \cap \tau^*$.

We state for the record that we think this is a bad finite axiomatization of NF, or finite template axiomatization of TST: we think the treatment of n -tuples is terribly difficult to work with. But it does work, and the second author chose to verify the consistency of TTT by verifying that the finite (template) axiomatization is satisfied in the structure for the language of TTT that we define, so it merits discussion here. It should be noted that any finite template axiomatization of TST could be used: this one is historically prior, but (as we remark), awkward to use, and there is an oddity, because the one impredicative axiom (the axiom of type lowering) does a lot of extra work, since it is also essential for constructing domains and ranges of functions. In the formalization, the proof of type lowering was in effect divided into the proof of the existence of $\iota^3\text{“TL}(x)$, which is predicative, and the proof of the existence of set unions of sets of singletons.

The second author reports that she chose this particular finite axiomatization for verification purposes despite its awkwardness because it uses very few defined concepts, and it is therefore clearer to a reader of the Lean formalization that the proofs indeed say what is required.

3 The model description

In this section we give a complete description of what we claim is a model of tangled type theory. Our metatheory is some fragment of ZFC.

3.1 The abstract supertype framework

We first discuss some abstract properties of the types we will construct, and explain the system of “supertypes”. Types are indexed by a well-ordering $\leq_\tau$ (from which we define a strict well-ordering $<_\tau$ in the obvious way). We refer to elements of the domain of $\leq_\tau$ as “type indices”.

We first define a system of “supertypes” (using the same type labels).

For each element t of $\text{dom}(\leq_\tau)$ we will define a set τ_t^* , called supertype t .

If m is the minimal element of the domain of $\leq_\tau$, we choose a set τ_m^* as supertype m .

$\leq_\tau$ and τ_m^* are the only parameters of the system of supertypes (which is not a model of TTT, but a sort of maximal structure for the language of TTT).

We describe the construction of τ_t^* , assuming that $t \in \text{dom}(\leq_\tau)$ and $t \neq m$ and for all $u <_\tau t$, we have defined τ_u^* .

We define τ_t^* as

$$\left\{ X \cup \{\{\tau_u^* : u <_\tau t\}\} : X \subseteq \bigcup_{u <_\tau t} \tau_u^* \right\}.$$

Foundation in the metatheory ensures a clean construction here. An element x of supertype $t >_\tau m$ is a subset of the union of all lower types, with $t^+ = \{\tau_u^* : u <_\tau t\}$ added as an element, so such an x is always nonempty. The set t^+ has supertype u as an element for each $u <_\tau t$, so t^+ and so x cannot belong to any supertype u with $u <_\tau t$, by foundation. The labelling element t^+ cannot belong to supertype t by foundation, because an element of supertype t must be nonempty and have t^+ as an element. Further, t^+ cannot belong to any supertype v with $t <_\tau v$, because any element of v contains v^+ as an element which contains supertype t as an element and any element of supertype t contains t^+ as an element, so $t^+ \in v$ would violate foundation in the metatheory.

The membership relations of this structure are transparent: $x \in_{t,u} y$ ($t <_\tau u$) is defined as $x \in \tau_t^* \wedge y \in \tau_u^* \wedge x \in y$. The considerations above show that there are no unintended memberships caused by the labelling elements t^+ , because the labelling elements cannot themselves belong to any supertype. Note the presence of $\emptyset_t = \{t^+\}$ in supertype t , which has no elements of any type $u <_\tau t$ (and is distinct from $\emptyset_v$ for $v \neq t$).

The system of supertypes is certainly not a model of TTT, because it does not satisfy extensionality. It is easy to construct many sets in a higher type with the same extension over a given lower type, by modifying the other extensions of the object of higher type.

The system of supertypes does satisfy the comprehension scheme of TTT. One can use Jensen’s method to construct a model of stratified comprehension

with no extensionality axiom from the system of supertypes, and stratified comprehension with no extensionality axiom interprets NFU in a manner described by Marcel Crabbé in [1].

Proposition 3.1 (the generality of the system of supertypes). Any model of TTT (assuming there are any) is isomorphic to a substructure of a system of supertypes.

Proof. Let M be a model of TTT (more generally, any structure for the language of TTT in which each object outside the base type is determined given all of its extensions). Let $\leq_M$ be the well-ordering on the types of M and let m be the minimal type of M . We will assume as above that $\leq_\tau$ is a well-ordering of type labels t with corresponding actual types τ_t of M : of course, we could use the actual types of M as type indices, but we preserve generality this way. We also assume that the sets implementing the types of M are disjoint (it is straightforward to transform a model in which the sets implementing the types are not disjoint to one in which they are, without disturbing its theory, by replacing each $x \in \tau_t$ with (x, t)).

We consider the supertype structure generated by $\leq_\tau := \leq_M$ and $\tau_m^* := \tau_m$. We indicate how to define an embedding from M into this supertype structure.

Define $I(x) = x$ for $x \in \tau_m = \tau_m^*$.

If we have defined I on each type $u <_\tau t$, we define $I(x)$, for $x \in \tau_t$, as

$$\bigcup_{u <_\tau t} \{I(y) : y \in_{u,t}^M x\} \cup \{\{\tau_u^* : u <_\tau t\}\}.$$

It should be clear that as long as M satisfies the condition that an element of any type other than the base type is uniquely determined given all of its extensions in lower types, I is an isomorphism from M to a substructure of the stated system of supertypes. A model of TTT, in which any one extension of an element of any higher type in a lower type exactly determines the object of higher type, certainly satisfies this condition. So the problem of constructing a model of TTT is equivalent to the problem of constructing a model of TTT which is a substructure of a supertype system. $\square$

Some advantages of this framework are that the membership relations in TTT are interpreted as subrelations of the membership relation of the metatheory, while the types are sensibly disjoint.

Remark 3.2. Notice that if $\alpha >_\tau \beta$ are type indices, and $x \in \tau_\alpha^*$, $x \cap \tau_\beta^*$ is the extension of x over the supertype τ_β^* . This will be inherited by a scheme of types τ_γ with each $\tau_\gamma \subseteq \tau_\gamma^*$ if an additional condition holds: for $\alpha >_\tau \beta$, we will have for $x \in \tau_\alpha$ that $x \cap \tau_\beta^*$ is the extension of x over the supertype τ_β^* as already noted: for it to be the extension over the type τ_β we need the general condition $x \cap \tau_\delta^* \subseteq \tau_\delta$ for all $\gamma > \delta$ and $x \in \tau_\gamma$.

3.2 Preliminaries of the construction

Now we introduce the notions of our particular construction in this framework.

Definition 3.3 (model parameters). Let λ be a limit ordinal. Let $\leq_\tau$ be the order on $\lambda \cup \{-1\}$ which has -1 as its minimal element and agrees otherwise with the usual order on λ .

Let κ be an uncountable regular cardinal (that is, an uncountable regular initial ordinal). Sets of cardinality $< \kappa$ we call *small*. Sets which are not small we may call *large*.

We make no assumption about the relative sizes of κ and λ .

Let μ be a strong limit cardinal with $\kappa < \mu$, $\lambda \leq \mu$ and with the cofinality of μ at least $\max(\kappa, \lambda)$.

Remark 3.4. The minimal model parameters are $\lambda = \omega$, $\kappa = \omega_1$, $\mu = \beth_{\omega_1}$.

Definition 3.5 (supertypes). Let $\tau_{-1}^* = \tau_{-1}$ be

$$\{(\nu, \beta, \gamma, \alpha) : \nu < \mu \wedge \beta \in \lambda \cup \{-1\} \wedge \gamma \in \lambda \setminus \{\beta\} \wedge \alpha < \kappa\}.$$

Note that this completes the definition of the supertype structure we are working in as defined in section 3.1: we now have a definite reference for τ_α^* for $\alpha \in \lambda$. The important part of this definition is that τ_{-1}^* has size μ ; the precise form of τ_{-1}^* is chosen to aid with definition 3.18. This type, while important for the model construction, will not be part of our final model of TTT; the types of the final model will be indexed by λ .

Notice that if α, β are type indices, $\alpha \in \beta$ is a convenient short way to say $-1 <_\tau \alpha <_\tau \beta$.

The levels of the structure we will define will be denoted by τ_α for $\alpha \in \lambda \cup \{-1\}$. As we have already noted, $\tau_{-1} = \tau_{-1}^*$, and for $-1 < \alpha$, we will have $\tau_\alpha \subset \tau_\alpha^*$.

Definition 3.6 (extended type index). A nonempty finite subset of $\lambda \cup \{-1\}$ with minimum element -1 may be termed an *extended type index*. If A is a finite subset of $\lambda \cup \{-1\}$ with at least two elements, A_1 is defined as $A \setminus \{\min(A)\}$, and further A_{n+1} is defined as $(A_n)_1$ where this makes sense: the notation A_2 for $(A_1)_1$ sees use.

Definition 3.7 (atoms, litters and near-litters). We may refer to elements of τ_{-1} as *atoms* from time to time, though they are certainly not atomic in terms of the metatheory.

A *litter* is a subset of τ_{-1} of the form $L_{\nu, \beta, \gamma} = \{(\nu, \beta, \gamma, \alpha) : \alpha < \kappa\}$. The litters make up a partition of type -1 (which is of size μ) into size κ sets.

On each litter $L = L_{\nu, \beta, \gamma}$ define a well-ordering $\leq_L$: $(\nu, \beta, \gamma, \alpha) \leq_L (\nu, \beta, \gamma, \alpha')$ iff $\alpha \leq \alpha'$. The strict well-ordering $<_L$ is defined in the obvious way. This well-ordering is used in only one place in the paper (definition 4.2), and its use could easily be avoided, but we find its concreteness appealing.

A *near-litter* is a subset of τ_{-1} with small symmetric difference from a litter. We define $M \sim N$ as $|M \Delta N| < \kappa$, for M, N near-litters: in English, we say that M is *near* N iff $M \sim N$. Note that nearness is an equivalence relation on near-litters.

We define N° , for N a near-litter, as the (necessarily unique) litter L such that $L \sim N$.

In our model, any near-litter will be the extension of κ -amorphous sets in each type above the base type, and we will use these amorphous sets to limit how much the model knows about the relationship between the different extensions of a model element.

Proposition 3.8. There are exactly μ near-litters.

Proof. A near-litter is determined as the symmetric difference of a litter L (there are μ litters) and a small subset of τ_{-1} . So it is sufficient to show that a set of size μ has no more than μ small subsets. As μ has cofinality at least κ , each small subset of μ has a strict upper bound $\nu < \mu$, and so is an element of $\bigcup_{\nu < \mu} \mathcal{P}(\nu)$. But as μ is a strong limit cardinal, each $\mathcal{P}(\nu)$ has size less than μ , so there are only μ bounded subsets of μ ; in particular, there are only μ small subsets of μ . $\square$

3.3 Permutations and supports

Our argument is essentially a Fraenkel–Mostowski construction, so we make use of a group of permutations. More precisely, each type in our model will have a group of permutations acting on it. To make this concrete, we first define groups of permutations that act on the entire supertype structure.

Definition 3.9 (structural permutation). A -1 -*structural permutation* is a permutation π of $\tau_{-1}^* = \tau_{-1}$ such that if N is a near-litter, its image $\pi''N$ is a near-litter.

An α -*structural permutation* ($-1 < \alpha$) is a permutation π of τ_α^* such that for each type $\beta < \alpha$ there is a β -structural permutation π_β such that $\pi(x) \cap \tau_\beta^* = \pi_\beta''(x \cap \tau_\beta^*)$ for any $x \in \tau_\alpha^*$. The maps π_β are referred to as *derivatives* of π .

Each type α in the model will come equipped with a designated subgroup of α -structural permutations, called the α -*allowable permutations*. The action of this group will preserve model elements: if π is α -allowable and $x \in \tau_\alpha$, then $\pi(x) \in \tau_\alpha$.

Our proofs will rely on the fact that the action of an α -allowable permutation on an α -model element is completely determined by its action on an α -*support*. Supports are combinatorial structures that can be used to track how a permutation acts along different derivatives on a small part of our model.

Definition 3.10 (support condition). An α -*support condition* ($-1 \leq \alpha$) is defined as a pair (x, A) , where

1. A is an extended type index with maximum α (recall that an extended type index has minimum -1); and
2. $x \subseteq \tau_{-1}$ is either a singleton or a near-litter, and must be a singleton in the case $\alpha = -1$.

Definition 3.11 (support). An α -*support* is defined as a function S from a small ordinal to α -support conditions. We write S_δ instead of $S(\delta)$.

We define the action of an α -structural permutation π on an α -support S : if $S_\delta = (x, A)$, then $\pi[S]_\delta = (\pi_A x, A)$. This means that $\pi[S] = S$ iff $\pi_A x = x$ for all $(x, A) \in \text{rng}(S)$.

Remark 3.12. -1 -supports behave somewhat differently from α -supports for $-1 < \alpha$, and we consider the notion of -1 -support to be merely a technical convenience, often making the statements of definitions and lemmas more uniform. The condition that -1 -support conditions cannot contain near-litters is used exactly once, in the proof of proposition 4.50.

Remark 3.13 (counting supports). It is straightforward to observe that there are μ α -supports for any α : there are μ atoms, μ near-litters, and $< \mu$ finite subsets of $\alpha + 1 < \lambda \leq \mu$ (all type indices involved in an α -support are $\leq \alpha$); thus the set of α -support conditions (which we will call **SC** for the moment) is of size μ ; note that the set $\kappa \times \text{SC}$ is of cardinality μ and each α -support is a small subset of $\kappa \times \text{SC}$, and so, as we have already seen than sets of size μ have μ small subsets, it follows that there are no more than μ α -supports.

3.4 Position functions

At various points throughout this paper, we will need to establish results that depend on the precise way that the types in our model are tangled together. Such results cannot all be proven by induction along λ , because the mechanism that we use to enforce extensionality introduces constraints between model elements across arbitrary types. To address this, we are going to introduce μ -valued position functions to let us prove these results by induction along μ instead, simultaneously across all types. We will describe these position functions in the next section.

In order to anchor our different position functions to a common reference, we will first define a position function for atoms and near-litters.

Lemma 3.14 (position function for near-litters and singletons of atoms). There is a bijection ι_* from the set of all near-litters and singletons of atoms to μ with the following properties:

1. $\iota_*(L) < \iota_*(\{a\})$ if $a \in L$ and L is a litter;
2. $\iota_*(N^\circ) < \iota_*(N)$ if N is a near-litter which is not a litter; and
3. $\iota_*(\{a\}) < \iota_*(N)$ if $a \in N \Delta N^\circ$.

Throughout this paper, we fix such a function ι_* , treating this lemma as a definition.

The function may be constructed directly: first choose an ordering of type μ on litters, then put singletons of atoms directly after the litter they are inside (and before all later litters), then put near-litters after the corresponding litter and all singletons of atoms in the symmetric difference, then map each near-litter or singleton of an atom to its position in the resulting order. We do not run out of room before finishing the construction because μ has cofinality at least κ . We prove more formally that there are such functions just below.

Lemma 3.15 (position function generator). If X is a set of size at most μ and $f : X \rightarrow \mathcal{P}(\mu)$ is such that each $f(x)$ is bounded, then there is an injection $\iota : X \rightarrow \mu$ such that $\iota(x) \notin f(x)$.

Proof. Choose a well-ordering $<_X$ of X of type at most μ . At each stage x , assume we have already constructed $\iota(y)$ for $y <_X x$. We then define

$$\iota(x) = \min(\mu \setminus f(x) \setminus \{\iota(y) : y <_X x\}).$$

Note that as $f(x)$ is bounded in μ , and μ has initial order type, the set $\mu \setminus f(x)$ is of cardinality μ . Since $\{\iota(y) : y <_X x\}$ has cardinality less than μ , the difference must be nonempty, so the definition given always succeeds. $\square$

We now show how to construct a position function for near-litters and singletons of atoms satisfying the given constraints.

Proof of lemma 3.14. First, let ι_0 be any injection from the set of litters into μ . Next, use lemma 3.15, with X taken to be the collection of all singletons of atoms, setting $f(\{a\}) = \{\nu : \nu \leq \iota_0(L)\}$ where L is the unique litter such that $a \in L$, to create ι_1 .

Next, use lemma 3.15 again, with X taken to be the collection of all near-litters, setting

$$f(N) = \{\nu : \nu \leq \iota_0(N^\circ)\} \cup \bigcup_{a \in N \Delta N^\circ} \{\nu : \nu \leq \iota_1(\{a\})\},$$

to obtain ι_2 .

Finally, we define a map ι from near-litters and singletons of atoms to $3 \times \mu$ by

$$g(x) = \begin{cases} (0, \iota_0(x)) & \text{if } x \text{ is a litter} \\ (1, \iota_1(x)) & \text{if } x \text{ is a singleton} \\ (2, \iota_2(x)) & \text{otherwise} \end{cases}$$

We can then define $\iota_*(x)$ to be the position of $g(x)$ in the restriction of the lexicographic order on $3 \times \mu$ to the image of g (which is of order type μ).

The resulting function ι_* is then clearly a bijection, and has the required properties by construction, noting that near-litters that happen to be litters will automatically satisfy the third condition. $\square$

3.5 Hypotheses for the recursion

The construction of the types τ_α is by a recursion. The initial type τ_{-1} has already been defined. For an ordinal $\alpha < \lambda$ we are assuming that τ_β for $-1 \leq \beta < \alpha$ have already been constructed and satisfy various hypotheses of the recursion. In this section, we will describe the hypotheses that we assume while we are constructing our model at a fixed type $-1 < \alpha$. In these hypotheses, β is any type index such that $-1 \leq \beta < \alpha$.

First, we define useful notation for a concept that will appear frequently.

Definition 3.16 (typed objects). If X is a subset of τ_γ and $\gamma < \beta$, we define X_β as the unique element Y of τ_β such that $Y \cap \tau_\gamma = X$ (if this exists). Of course, this notation is only usable to the extent that we suppose that extensionality holds. If N is a near-litter we refer to N_β as a *typed near-litter* (of type β).

- (I1) There is a designated subset $\tau_\beta \subseteq \tau_\beta^*$, whose elements are called *type β model elements*. At the base type, we have $\tau_{-1} = \tau_{-1}^*$. Elements of model elements are model elements: if $x \in y \in \tau_\beta$ and $x \in \tau_\gamma^*$, then $x \in \tau_\gamma$.
- (I2) There is a designated subgroup of β -structural permutations called the *β -allowable* permutations, which map model elements to model elements. All -1 -structural permutations are -1 -allowable.
- (I3) For every model element $x \in \tau_\beta$, there is a β -support S such that if π is β -allowable, $\pi[S] = S$ implies $\pi(x) = x$. Such an S is said to *support*, or *be a support for*, x .

Note that we can easily prove (I3) for $\beta = -1$: every $a \in \tau_{-1}$ has a support S with domain 1, given by $S_0 = (\{a\}, \{-1\})$.

A key fact that we will rely upon in the proof is that the size of each type is exactly μ . It will turn out (in definition 3.18) that the most convenient way to phrase this in the inductive hypotheses is as follows:

- (I4) Let τ_β^+ be the collection of pairs (x, S) where x is a β -model element and S is a β -support for x . Then, there is an injection $\iota_\beta^+ : \tau_\beta^+ \rightarrow \mu$ called the *position function*.

Since there are exactly μ β -supports (remark 3.13), this hypothesis implies that $|\tau_\beta| \leq \mu$ for $\beta < \alpha$.

When we define the types in our model, we will need a way to hide certain data from the model's internal knowledge (in definition 3.18). We will do this by embedding that data in some near-litters, and then turning these near-litters into model elements.

- (I5) If $-1 < \beta$, each near-litter N gives rise to a unique model element $x \in \tau_\beta$ such that $x \cap \tau_{-1} = N$. Such an element is called a *typed near-litter*, and is denoted by N_β as stated above.

In particular, as there are μ near-litters, we must have $|\tau_\beta| \geq \mu$, and so by combining (I4) and (I5), we must have the equality $|\tau_\beta| = \mu$.

These five inductive hypotheses on the types below α are enough to define the type τ_α and its group of α -allowable permutations. Indeed, we will easily be able to satisfy (I1), (I2), (I3), and (I5) at level α . However, we will see that proving the equation $|\tau_\beta| = \mu$, and thus obtaining (I4), is a significant technical challenge. To accomplish this, we add some extra inductive hypotheses that describe the precise way that the layers of our type hierarchy interact with each other. These extra inductive hypotheses use the fact that τ_α and the α -allowable permutations will have already been defined. This will then allow us to prove the bound $|\tau_\beta| = \mu$ by a symmetry argument (section 4).

- (I6) Extensionality holds at *each* pair of type indices $-1 < \gamma < \beta \leq \alpha$: if $x, y \in \tau_\beta$ are such that $x \cap \tau_\gamma = y \cap \tau_\gamma$, then $x = y$.
- (I7) When $-1 < \gamma < \beta \leq \alpha$, for each $x \in \tau_\gamma$, the singleton $\{x\}_\beta$ exists: there is a (unique) element $y \in \tau_\beta$ such that $y \cap \tau_\gamma = \{x\}$.

We now make an assumption on the position functions ι_β^+ . Recall that ι_β^+ is a function defined on pairs (x, S) where S is a β -support for x . Such a support is a function, and its image is a small set of support conditions (t, A) , where t is either a near-litter or a singleton of an atom. Our condition on ι_β^+ is that the position of (x, S) lies after the positions of all atoms and near-litters that appear within S . We will use this condition to prove results about our entire structure by induction along μ .

- (I8) Let $-1 \leq \beta < \alpha$:

- For any $(x, S) \in \tau_\beta^+$, $\iota_*(t) < \iota_\beta^+(x, S)$ if (t, A) is in the range of S .
- $\iota_*(N) \leq \iota_\beta^+(N_\beta, S)$ for any near-litter N and support S of N .

Note that this hypothesis does not apply at level α ; this is because we do not yet know that $|\tau_\alpha| = \mu$, so the position function ι_α^+ has not yet been constructed. We can supply such a position function at type -1 now.

Lemma 3.17 (position function at type -1). We can satisfy (I4) and (I8) at type -1 . Concretely, there is an injection $\iota_{-1}^+ : \tau_{-1}^+ \rightarrow \mu$ such that when $(x, S) \in \tau_{-1}^+$ and (t, A) is in the range of S , we have $\iota_*(t) < \iota_{-1}^+(x, S)$.

Proof. Invoke lemma 3.15 with $X = \tau_{-1}^+$ and

$$f(x, S) = \{\nu : \exists \delta : S_\delta = (t, A) \wedge \nu \leq \iota_*(t)\} \quad \square$$

Finally, we stipulate a concrete definition for the allowable permutations. To give this hypothesis, we first define a gadget called the *f maps*. These maps provide a way to turn a model element with support into a litter, which can then be injected back into the model at a different type (I5).

Definition 3.18 (*f maps*, crucially important).⁹ We define for each type index β less than α and each ordinal $\gamma < \lambda$ distinct from β a function $f_{\beta, \gamma}$ (whose definition does not actually depend on α : it will be the same at every stage). $f_{\beta, \gamma}$ is an injection from τ_β^+ into the set of litters, with range included in $\{L_{\nu, \beta, \gamma} : \nu < \mu\}$ to ensure that distinct *f maps* have disjoint ranges.

When we define $f_{\beta, \gamma}(x)$, we presume that we have already defined it for y where $\iota_\beta^+(y) < \iota_\beta^+(x)$. We define $f_{\beta, \gamma}(x)$ as L , where $\iota_*(L)$ is minimal such that

⁹The use of model elements with support rather than simply model elements as domain elements of the *f maps* is a substantial contribution of the second author to the mathematics of the paper, above simply verifying the work of the first author. The proof could be carried out without this, but it is much easier to present with this refinement. There are other ways in which the second author has contributed to the mathematics, but this one is especially worthy of note.

1. $L \in \{L_{\nu, \beta, \gamma} : \nu < \mu\}$,
2. $\iota_{\beta}^{+}(x) < \iota_{*}(L)$ [and so for any $N \sim L$, $\iota_{\beta}^{+}(x) < \iota_{*}(N)$, and for any $z \in L$, $\iota_{\beta}^{+}(x) < \iota_{*}(\{z\})$],
3. and for any y with $\iota_{\beta}^{+}(y) < \iota_{\beta}^{+}(x)$, $f_{\beta, \gamma}(y) \neq L$.

Note that the ranges of distinct f maps are disjoint sets.

We may now provide a definition for allowable permutations.

(I9) Let $\beta \leq \alpha$. A β -structural permutation π is β -allowable if and only if:

- (a) each of its derivatives π_{γ} are allowable; and
- (b) for all $\gamma \neq \delta$ such that $-1 \leq \gamma < \beta$ and $-1 < \delta < \beta$, and for all $(x, S) \in \tau_{\gamma}^{+}$, we have

$$f_{\gamma, \delta}(\pi_{\gamma}(x), \pi_{\gamma}[S]) \sim (\pi_{\delta})_{-1} "f_{\gamma, \delta}(x, S)$$

Note that this condition is vacuously true at type $\beta = -1$.

This definition relies on the f maps, which in turn rely on supports, which themselves rely on already having definitions for allowable permutations at lower levels. For this reason, we could not have given a definition of allowable permutations at the start, and must instead recover the correct definition (via inductive hypothesis) once notions such as the f maps have been defined.

With these latter four hypotheses, we will prove in section 4 that $|\tau_{\alpha}| = \mu$, which will allow us to complete the model construction at level α .

3.6 Machinery for enforcing extensionality in the model

This section begins the main inductive step. We assume hypotheses (I1)–(I5) above for $\beta < \alpha$, and construct data for our model at level α .

We describe the mechanism which enforces extensionality in the substructure of this supertype structure that we will build. Each model element will have a particular extension, called its *distinguished extension*, from which all *alternate* extensions can be calculated using the f maps.

Definition 3.19 (pre-extensional). We define the notion of *pre-extensional* element of τ_{α}^{*} . An element x of τ_{α}^{*} is pre-extensional iff there is a $\beta < \alpha$ such that

1. $x \cap \tau_{\beta}^{*} \subseteq \tau_{\beta}$;
2. $\beta = -1$ if $x \cap \tau_{-1}$ is nonempty or if any $x \cap \tau_{\gamma}$ ($\gamma \in \alpha$) is empty; and
3. for each $\gamma \in \alpha \setminus \{\beta\}$,

$$x \cap \tau_{\gamma} = \{N_{\gamma} : (\exists a \in x \cap \tau_{\beta} : \exists S : N \sim f_{\beta, \gamma}(a, S))\}.$$

We say for any $x \in \tau_\alpha^*$ and β with this property that $x \cap \tau_\beta$ is a *distinguished extension* of x .

We will ensure later that all elements of τ_α are pre-extensional.

Remark 3.20 (typed near-litters). We now know how to compute all other extensions of typed near-litters N_α , because the -1 -extension of N_α must be its distinguished extension by the second condition, and this indicates how to compute all the other extensions.

In fact, an even stronger condition holds, although the current inductive hypotheses are not strong enough to state this entirely formally. As soon as we compute all maps $f_{-1,\gamma}$ for $\gamma < \lambda$, which happens in the first stage of the construction, we can actually determine for each near-litter N and each $\eta < \lambda$ the element of τ_η^* which will be N_η . Suppose we have already done this for all $\chi < \eta$. The -1 -extension of N_η is given, and since all maps $f_{-1,\gamma}$ are already known we can identify the near-litters M_γ ($\gamma < \eta$) which should belong to alternate extensions of N_η , and further we know exactly which element of τ_γ^* should be M_γ by inductive hypothesis.

Definition 3.21 (alternate extension map). For any $-1 < \gamma < \alpha$, $\beta < \alpha$ distinct from γ , and nonempty subset a of τ_β , we define $A_\gamma(a)$ as

$$\{N_\gamma : (\exists x \in a : \exists S : N \sim f_{\beta,\gamma}(x, S))\}.$$

For any nonempty subset x of τ_γ there is at most one subset y of any type such that $A_\gamma(y) = x$. There cannot be more than one such y in any given type because the f maps are injective. There cannot be more than one such y in different types because the ranges of f maps with distinct index pairs are disjoint. We use the notation $A^{-1}(x)$ for this set if it exists, defining a very partial function A^{-1} on nonempty subsets of types. We note that the function A^{-1} defined at each stage of the construction depends on the value of α : the function A^{-1} defined at a later stage extends the one defined at an earlier stage. In the contexts in which we use the function A^{-1} , α is fixed, so we do not see the need to adorn the notation with an additional subscript: it could be written $A_{(\alpha)}^{-1}$ but this does not seem necessary.

Note that any distinguished extension of any type element x is the image under A^{-1} of its other extensions if it is nonempty.

Using properties of the f maps, we can prove the following useful well-foundedness constraint.

Proposition 3.22. No nonempty subset of a type has infinitely many iterated images under A^{-1} , and there are no cycles in A^{-1} .

Proof. Let a be a nonempty subset of type γ for which $A^{-1}(a)$ exists.

Since $A^{-1}(a)$ exists, every element of a is of the form N_γ where N is a near-litter. Choose $N_\gamma \in a$ such that $\iota_*(N)$ is minimal. Note that N is in fact a litter by the constraints in lemma 3.14.

Let $b = A^{-1}(a)$: we have $a = A_\gamma(b)$, where b is a subset of some τ_δ with $\delta \neq \gamma$. In particular, $N = f_{\delta,\gamma}(u, U)$ for some $u \in \tau_\delta$ and U a support of u . If $u = M_\delta$ for some near-litter M , we can further state that $\iota_*(N) > \iota_*^+(M_\delta, U) \geq \iota_*(M)$: if b itself has an image under A^{-1} , the minimum value of ordinals $\iota_*(M)$ for $M_\delta \in b$ will be less than the minimum value of ordinals $\iota_*(N)$ for $N_\gamma \in a$, which establishes that there is an ordinal parameter determined by a nonempty subset of a type which decreases strictly when A^{-1} is applied (if it is applicable), and so no nonempty subset of a type can have infinitely many iterated images under A^{-1} , nor can there be any cycles in A^{-1} . $\square$

We now verify that the order conditions in the definition of $f_{\beta,\gamma}$ ensure that distinguished extensions are unique.

Proposition 3.23. For any pre-extensional $x \in \tau_\alpha^*$ and $\beta < \alpha$ there is only one set $x \cap \tau_\beta$ which is a distinguished extension of x .

Proof. If any $x \cap \tau_\beta$ ($\beta \in \alpha$) is empty or if $x \cap \tau_{-1}$ is not empty, $x \cap \tau_{-1}$ is the unique distinguished extension (if it is empty of course it coincides with all the other extensions).

So, what remains is the case of x with $x \cap \tau_{-1}$ empty and each $x \cap \tau_\beta$ nonempty for $\beta < \alpha$. Suppose that $x \cap \tau_\beta$ and $x \cap \tau_\delta$ were both distinguished. Then we would have

$$x \cap \tau_\beta = A^{-1}(x \cap \tau_\delta) = A^{-1}(A^{-1}(x \cap \tau_\beta)),$$

which contradicts proposition 3.22 as this would imply that A^{-1} has a cycle of length 2. $\square$

Definition 3.24 (extensional). We say that an element of τ_α^* is *extensional* iff it is pre-extensional and its distinguished extension has an even number of iterated images under A^{-1} .

Note that if x is an extensional element of τ_α^* then each of its alternate extensions has an odd number of iterated images under A^{-1} .¹⁰

Proposition 3.25 (extensionality). Two extensional elements of τ_α^* with any common extension (over a type other than τ_{-1}) are equal, thus verifying (I6) at level α .

Proof. If two extensional elements of τ_α^* have an empty extension (over a type other than τ_{-1}) in common, they both have all extensions empty and are equal. If two extensional elements of τ_α^* have a nonempty extension in common, it will be the distinguished extension of both, or a non-distinguished extension of both, since distinguished and non-distinguished extensions are taken from disjoint classes of subsets of types (when nonempty). In either case we deduce that the two elements have the same distinguished extension and thus have all extensions the same and are equal. $\square$

¹⁰We do know that we are carefully, explicitly, spelling out a construction which looks very much like the construction of the bijection in the Cantor–Schröder–Bernstein theorem. But the details of the maps involved are used, so everything must be spelled out.

Note that this gives weak extensionality over τ_{-1} (many objects have empty extension over type -1) but it gives full extensionality over any other type.

This completes the mechanism for enforcement of extensionality in the structure we are defining.

3.7 Allowable permutations and supports

A crucial aspect of this is that we will need to define τ_α so that it has cardinality μ for the process to continue (I4). It is certainly not a sufficient restriction to require elements of τ_α to be extensional: we will require a further symmetry condition.

Definition 3.26 (iterated derivatives). For any finite subset A of $\lambda \cup \{-1\}$ with maximum β and for any β -structural permutation π , define π_A as $(\pi_{A \setminus \{\min(A)\}})_{\min(A)}$ and $\pi_{\{\beta\}} = \pi$. The maps π_A may be referred to as iterated derivatives of π .¹¹ It should be clear that a structural permutation is exactly determined by its iterated derivatives which are -1 -structural.

Where π is a β -structural permutation with $\beta > -1$, we define π^+ as π_{-1} . This is occasionally useful to reduce notational clutter.

Definition 3.27 (operations on supports). We define various operations to manipulate supports.

1. For any support condition (x, B) we define $(x, B)^{\uparrow A}$ as $(x, B \cup A)$ if all elements of the set A are strictly larger than all elements of the set B . Further, if S is a support, we define $S^{\uparrow A}$ so that $(S^{\uparrow A})_\epsilon = (S_\epsilon)^{\uparrow A}$. By an abuse of notation we may write $(x, B)^{\uparrow \beta}$ or $S^{\uparrow \beta}$ where β is an ordinal for $(x, B)^{\uparrow \{\beta\}}$ or $S^{\uparrow \{\beta\}}$.
2. For any supports S and T we denote by $S + T$ a support which consists of S , followed by T : what this means is that $(S + T)_\epsilon = S_\epsilon$ for ϵ in the domain of S , $(S + T)_{\text{dom}(S)+\epsilon} = T_\epsilon$ for ϵ in the domain of T .
3. Recall that we say that an element x of τ_β^* has β -support S iff for every β -allowable permutation π , if $\pi[S] = S$ then $\pi(x) = x$. We say that an element x of τ_β^* which has a β -support is β -symmetric. We allow ourselves to say briefly “ x has support S ” and “ x is symmetric” if $x \in \tau_\beta$ and S is a β -support of x .

It is important to note that if S is a support of $x \in \tau_\beta$, $\pi[S]$ is a support of $\pi(x)$ for any β -allowable permutation π .

Structural permutations are defined on the supertype structure generally. We need a subclass of structural permutations which respects our extensionality requirements.

¹¹There is a silly notational point here: we might want to suppose π_A and π_β to be essentially distinguished in some way we do not actually implement (for example by type face) in order to prevent confusion of $\pi_{\{0\}}$ with π_1 . A similar problem exists due to the identification of finite ordinals with finite sets of ordinals. However, it can also be noted that $\pi_{\{0, \dots, n\}}$ and π_{n+1} cannot make sense for the same allowable permutation π , so we think this is harmless.

Definition 3.28 (allowable permutation). An α -allowable permutation is an α -structural permutation, each of whose derivatives π_β is a β -allowable permutation (as defined in (I2)), and which satisfies a coherence condition relating the f maps and derivatives of the permutation: for any $\beta, \gamma < \alpha$ for which $f_{\beta, \gamma}$ is defined,

$$f_{\beta, \gamma}(\pi_\beta(x), \pi_\beta[S]) \sim \pi_\gamma^+ f_{\beta, \gamma}(x, S).$$

This coherence condition is motivated in remark 3.29.

Any allowable permutation is actually defined on the entire supertype structure, though what interests us about it is its actions on objects in our purported TTT model. Note that (I9) stipulates that definition 3.28 correctly defines allowable permutations at all levels $-1 < \beta \leq \alpha$ (although we will not need this hypothesis yet).

Remark 3.29 (motivation of the coherence condition in definition 3.28). Recall the coherence condition for α -allowable permutations: for any $(x, S) \in \tau_\beta^+$ and any $\gamma \in \alpha \setminus \{\beta\}$,

$$f_{\beta, \gamma}(\pi_\beta(x), \pi_\beta[S]) \sim \pi_\gamma^+ f_{\beta, \gamma}(x, S).$$

The motivation for this is that we need α -allowable permutations to send extensional elements of supertypes to extensional elements. Suppose $x \in \tau_\alpha$ and $x \cap \tau_\beta = \{b\}$. If x is extensional, this has to be the distinguished extension of x . For any $\gamma \in \alpha \setminus \{\beta\}$, it follows that $x \cap \tau_\gamma$ is the set of all N_γ such that $N \sim f_{\beta, \gamma}(b, S)$ for some support S of b . Now consider an α -allowable permutation π . Then $\pi(x)$ has β -extension $\{\pi_\beta(b)\}$, and must have γ -extension equal to

$$\pi_\gamma^+ \{N_\gamma : \exists S : N \sim f_{\beta, \gamma}(b, S)\},$$

but for $\pi(x)$ to be extensional it must also have its γ -extension equal to

$$\{N_\gamma : \exists S : N \sim f_{\beta, \gamma}(\pi_\beta(b), S)\}.$$

This tells us that

$$\pi_\gamma(f_{\beta, \gamma}(b, S)_\gamma) \in \{N_\gamma : (\exists T : N \sim f_{\beta, \gamma}(\pi_\beta(b), T))\}$$

for each support S of b . The coherence condition enforces this neatly, showing that it is motivated by considerations required to get extensionality to work: the action of π_β conveniently correlates supports of b with supports of $\pi_\beta(b)$.

Proposition 3.30 (allowable permutations preserve A -maps). Let $-1 \leq \beta < \alpha$ and $-1 < \gamma < \alpha$ be distinct. If π is α -allowable and $a \subseteq \tau_\beta$, then $\pi_\gamma^+ A_\gamma(a) = A_\gamma(\pi_\beta^+ a)$.

Proof. Recall that for nonempty $a \subseteq \tau_\beta$ we defined $A_\gamma(a)$ in definition 3.21 as

$$\{N_\gamma : (\exists x \in a : (\exists S : N \sim f_{\beta, \gamma}(x, S)))\}.$$

We show that if π is α -allowable, then $\pi_\gamma^+ A_\gamma(a) = A_\gamma(\pi_\beta^+ a)$ follows from the coherence condition.

Each element of $\pi_\gamma "A_\gamma(a)$ is of the form $\pi_\gamma(N_\gamma)$ where N is a near-litter and there is some $x \in a$ and some β -support S of x such that $N \sim f_{\beta,\gamma}(x, S)$. Then

$$\pi_\gamma(N_\gamma) \cap \tau_{-1} = (\pi_\gamma)_{-1} "N \sim (\pi_\gamma)_{-1} "f_{\beta,\gamma}(x, S) \sim f_{\beta,\gamma}(\pi_\beta(x), \pi_\beta[S]).$$

So by the uniqueness condition in **(I5)**, any element of $\pi_\gamma "A_\gamma(a)$ is in $A_\gamma(\pi_\beta "a)$.

Conversely, each element of $A_\gamma(\pi_\beta "a)$ is of the form N_γ , where N is a near-litter and there is some $x \in a$ and some β -support S of $\pi_\beta(x)$ such that $N \sim f_{\beta,\gamma}(\pi_\beta(x), S)$. We then have $N \sim (\pi_\gamma)_{-1} "f_{\beta,\gamma}(x, \pi_\beta^{-1}[S])$. We want to show that $\pi_\gamma^{-1}(N_\gamma) \in A_\gamma(a)$. We have

$$\pi_\gamma^{-1}(N_\gamma) \cap \tau_{-1} = (\pi_\gamma)_{-1}^{-1} "N \sim (\pi_\gamma)_{-1}^{-1} "((\pi_\gamma)_{-1} "f_{\beta,\gamma}(x, \pi_\beta^{-1}[S])) = f_{\beta,\gamma}(x, \pi_\beta^{-1}[S]),$$

establishing what we need. $\square$

Notice that this shows that the coherence condition implies that the image under an allowable permutation of a pre-extensional element of our structure is pre-extensional.

Proposition 3.31 (allowable permutations preserve extensionality). Allowable permutations map extensional elements of τ_α^* to extensional elements.

Proof. Let $-1 \leq \beta < \alpha$ and $-1 < \gamma < \alpha$ be such that $\beta \neq \gamma$, and let π be an α -allowable permutation. By proposition 3.30, if $a \subseteq \tau_\beta$, then $A^{-1}(a)$ exists and is in τ_γ exactly if $A^{-1}(\pi_\beta "a)$ exists and is in τ_γ , and moreover $A^{-1}(\pi_\beta "a)$ is equal to $\pi_\gamma "A^{-1}(a)$ if it exists under these conditions. This verifies that the coherence condition implies that allowable permutations preserve full extensionality, not just pre-extensionality: the number of iterated images under A^{-1} of an extension that exist is not affected by application of an allowable permutation in a suitable sense. $\square$

3.8 Model elements defined

Definition 3.32 (model elements). We define τ_α as the set of elements x of τ_α^* such that $x \cap \tau_\beta^* \subseteq \tau_\beta$ for each $\beta < \alpha$, x is extensional, and x has an α -support.

Note that an image of an element of τ_α under an α -allowable permutation will belong to τ_α , because supportedness and extensionality are preserved by allowable permutations (proposition 3.31).

This definition explicitly enforces **(I1)**, **(I2)**, **(I3)**, **(I6)**, and **(I9)** at level α . We will now prove **(I5)** and **(I7)**: that typed near-litters and singletons exist at level α .

Proposition 3.33. For any near-litter N , there is a (necessarily unique) element $x \in \tau_\alpha$ such that $x \cap \tau_{-1} = N$, satisfying **(I5)** at level α .

Proof. Let $x \in \tau_\alpha^*$ be the supertype element such that $x \cap \tau_{-1} = N$, and $x \cap \tau_\beta = A_\beta(N)$ for $\beta \in \alpha$. By construction, its -1 -extension is distinguished. It has support S with domain 1 and $S_0 = (N, \{-1, \alpha\})$: if π fixes S , then $\pi_{-1} "N = N$, so the -1 -extension of x is fixed by π . But by extensionality over -1 (which holds since $x \cap \tau_{-1}$ is nonempty), π must fix all extensions of x . $\square$

Proposition 3.34 (symmetric comprehension). Let X be a subset of a type $-1 < \beta < \alpha$. Suppose that there is an α -support S with the property that for any α -allowable π , when $\pi[S] = S$, we have $\pi_\beta "X = X$. Then $X = x \cap \tau_\beta$ for some (necessarily unique) model element $x \in \tau_\alpha$.

Proof. If X is empty, then X is simply the β -extension of the empty model element. Suppose X is nonempty, and let n be the number of iterated images that X has under A^{-1} .

If n is even, consider the element $x \in \tau_\alpha^*$ such that $x \cap \tau_\beta^* = X$, $x \cap \tau_{-1}^* = \emptyset$, and whose other extensions at levels $-1 < \gamma < \alpha$, $\gamma \neq \beta$ are given by $x \cap \tau_\gamma^* = A_\gamma(X)$. We show this is a model element. Clearly each extension satisfies $x \cap \tau_\gamma^* \subseteq \tau_\gamma$. It is pre-extensional because $x \cap \tau_\beta$ is its distinguished extension, and it is extensional as n is even. It has support S : if π is α -allowable and fixes S , then $\pi_\beta "(x \cap \tau_\beta) = x \cap \tau_\beta$ by assumption, so by extensionality, π must fix x . So $x \in \tau_\alpha$ and $x \cap \tau_\beta = X$.

If n is odd, let $Y = A^{-1}(X) \subseteq \tau_\gamma$. We claim that S has the property that if π is α -allowable and $\pi[S] = S$, then $\pi_\gamma "Y = Y$. Indeed, as allowable permutations commute with A maps by proposition 3.30,

$$\pi_\gamma "Y = \pi_\gamma "A^{-1}(X) = A^{-1}(\pi_\beta "X) = A^{-1}(X) = Y.$$

This allows us to repeat the argument of the previous paragraph using Y to obtain a model element $x \in \tau_\alpha$ satisfying $x \cap \tau_\gamma = Y$, where $x \cap \tau_\gamma$ is distinguished. Then, by pre-extensionality, we must have $x \cap \tau_\beta = A_\beta(x \cap \tau_\gamma) = X$, as required. $\square$

Remark 3.35. All model elements for $\alpha > 0$ are of this form: if $x \in \tau_\alpha$ and $\beta \in \alpha$, then any support S for x satisfies the hypothesis that $\pi[S] = S$ implies $\pi_\beta "(x \cap \tau_\beta) = x \cap \tau_\beta$.

Corollary 3.36. The singleton $\{x\}_\alpha$ exists for all $x \in \tau_\beta$, satisfying (I7) at level α .

Proof. If S is a support for x , then $S^{\uparrow\alpha}$ satisfies the hypothesis of symmetric comprehension as required. $\square$

Corollary 3.37 (κ -completeness of the structure). Let $\beta \in \alpha$. For any small subset X of type α , there is a model element $y \in \tau_\alpha$ such that

$$y \cap \tau_\beta = \bigcup_{x \in X} x \cap \tau_\beta.$$

Proof. Let S be a support whose range is obtained from the unions of ranges of supports for each $x \in X$. Then, if $\pi[S] = S$, we have

$$\pi_\beta \left(\bigcup_{x \in X} x \cap \tau_\beta \right) = \bigcup_{x \in X} \pi_\beta (x \cap \tau_\beta) = \bigcup_{x \in X} \pi(x) \cap \tau_\beta = \bigcup_{x \in X} x \cap \tau_\beta.$$

The conclusion then follows directly from symmetric comprehension. $\square$

Corollary 3.38. As small unions exist and singletons exist, X_α must also exist for any small set $X \subseteq \tau_\beta$.

We still have to prove that the cardinality of τ_α , and so of τ_α^+ , is μ , to show that the construction works; we will verify this in the next section. However, we can show now that given $|\tau_\alpha| = \mu$, we can satisfy the remaining hypotheses (I4), (I8).

Proposition 3.39 (position functions). Given that τ_α^+ has size μ , there exist a choice of injections ι_α^+ enforcing (I4) and (I8). More precisely,

- $\iota_*(t) < \iota_\alpha^+(x, S)$ if (t, A) is in the range of S .
- $\iota_*(N) \leq \iota_\alpha^+(N_\alpha, S)$ for any near-litter N and support S of N

Recall that we then define $x \leq_\alpha^+ y$ as $\iota_\alpha^+(x) \leq \iota_\alpha^+(y)$.

Proof. We use the position function generator (proposition 3.15): recall that this states that if X is a set of size at most μ and $f : X \rightarrow \mathcal{P}(\mu)$ is such that each $f(x)$ is bounded, then there is an injection $\iota : X \rightarrow \mu$ such that $\iota(x) \notin f(x)$. Setting $X = \tau_\beta^+$, we define

$$g(x, S) = \bigcup_{N \text{ s.t. } x=N_\delta} \iota_*(N) \cup \bigcup_{(y,A) \in \text{rng}(S)} \iota_*(y)$$

and let $f(x, S) = \{\nu : (\exists \xi \in g(x, S) : \nu \leq \xi)\}$. As μ has cofinality at least κ and each $g(x, S)$ is small, each $f(x, S)$ is bounded as required. $\square$

At this point we have a complete description of the structure which we claim is a model of TTT.

It should be noted that type 0 has a very simple description (which we do not use in any of our later results):

Proposition 3.40. Suppose that $\alpha = 0$, and let $x \in \tau_0$. Then $x \cap \tau_{-1}$ has small symmetric difference from a union of a small or co-small set of litters.

Proof. Let S be a 0-support for x . The elements of the range of S are of the form $(t, \{-1, 0\})$ where t is either a near-litter or a singleton of an atom. These subsets $t \subseteq \tau_{-1}$ give rise to a partition P of τ_{-1} , where atoms $a, b \in \tau_{-1}$ lie in the same compartment if and only if whenever $S_\delta = (t, A)$, we have $a \in t \iff b \in t$.

We claim that if $a, b \in \tau_{-1}$ lie in the same compartment, then $a \in x \iff b \in x$. To show this, consider the 0-structural permutation π that swaps a with b along the path $\{-1, 0\}$ and acts as the identity elsewhere:

$$\pi_{-1}(a) = b; \quad \pi_{-1}(b) = a; \quad \pi_{-1}(c) = c \text{ otherwise}$$

All -1 -structural permutations are allowable, so π_{-1} is allowable, and the coherence condition for $\alpha = 0$ is vacuous, so π itself must be allowable. Now, as a and b lie in the same compartment, it is easy to see that $\pi[S] = S$, which implies that $\pi(x) = x$. Therefore, as allowable permutations preserve membership,

$$a \in x \iff \pi_{-1}(a) \in \pi(x) \iff b \in x$$

In particular, for any compartment $C \in P$, either $C \cap x = \emptyset$ or $C \subseteq x$.

Note that there are a small number of compartments in P , each of which is either small, a near-litter, or the unique large set C of size μ corresponding to atoms that lie outside all t appearing in the range of S . Let us first consider the case $C \cap x = \emptyset$. Define

$$X = \{L : |x \cap L| = \kappa\}$$

We claim that X is a small set of litters such that $x \cap \tau_{-1}$ is near $\bigcup X$. To show this, consider the set

$$Y = \{L : (\exists \delta : S_\delta = (t, A) \wedge t \cap L \neq \emptyset)\}$$

As Y is clearly small, we can prove that X is small by noticing that $X \subseteq Y$: indeed, if $L \notin Y$, then L is entirely contained inside C , and so $x \cap L$ must be the empty set, giving $L \notin X$.

Now we prove that $(x \cap \tau_{-1}) \Delta \bigcup X$ is small. First, it is a subset of $\bigcup Y$: indeed, if $a \notin \bigcup Y$, then $a \in C$ so $a \notin x$ by disjointness. So as Y is small and κ is regular, it suffices to notice that for each litter $L \in Y$,

$$\left| \left((x \cap \tau_{-1}) \Delta \bigcup X \right) \cap L \right| < \kappa$$

Finally, consider the case where $C \subseteq x$. By symmetric comprehension (proposition 3.34), there is a model element $y \in \tau_0$ such that $y \cap \tau_{-1} = \tau_{-1} \setminus x$. Since y also has support S , the argument above shows that $y \cap \tau_{-1}$ is near the union of a small set of litters. From this, it follows that $x \cap \tau_{-1}$ is near the union of the complement of that set of litters, which itself is a co-small set of litters as required. $\square$

It turns out that the same holds in higher types: the -1 -extension of any type α model element is near the union of a small or co-small set of litters. However, we have not yet established enough theory about allowable permutations at higher types to prove this result for $\alpha > 0$.

4 Verification that the structure defined is a model

From this point, we assume all of the inductive hypotheses of section 3.5.

4.1 The Freedom of Action theorem

In this section, we will prove a theorem that gives us a way to construct allowable permutations (theorem 4.19). To do this, we will define a certain class of objects for each type, called *approximations*, that behave like partial structural permutations. We will say that an approximation is called *coherent* if it satisfies a particular version of the coherence condition on allowable permutations (definition 3.28). We will then show that every coherent approximation can be identified with the restriction of some allowable permutation.

Definition 4.1 (-1 -approximation). A -1 -approximation is a function ψ such that:

1. The domain and image of ψ are the same and ψ is injective.
2. Each domain element of ψ is either an atom or a litter, and moreover, ψ maps atoms to atoms and litters to litters.
3. For each litter L , ψ and ψ^{-1} are defined on only a small collection of atoms $a \in L$.

We will associate a partial function ψ^* on atoms to each -1 -approximation ψ . The action of ψ on atoms will agree with the action of ψ^* , and the action of ψ on litters will agree with the pointwise action of ψ^* only up to nearness.

Definition 4.2. If ψ is a -1 -approximation, we define the partial function ψ^* by:

1. If a is an atom and $a \in \text{dom}(\psi)$, then $\psi^*(a) = \psi(a)$.
2. If a is an atom with $a \notin \text{dom}(\psi)$ but $a \in L$ and $L \in \text{dom}(\psi)$, then

$$\psi^*(a) = \pi_{M,N}(a), \text{ where } M = L \setminus \text{dom}(\psi) \text{ and } N = \psi(L) \setminus \text{dom}(\psi)$$

where for any co-small subsets of litters M, N , the map $\pi_{M,N}$ is the unique map ρ from M to N such that for any $x, y \in M$,

$$x <_{M \circ} y \leftrightarrow \rho(x) <_{N \circ} \rho(y) :$$

$\pi_{M,N}$ is the unique map from M onto N which is strictly increasing in the order determined by fourth projections of atoms. Notice that $\pi_{M,N} \circ \pi_{L,M} = \pi_{L,N}$ will hold if L, M, N are co-small subsets of litters, and $\pi_{L,M} \circ \pi_{M,L} = \pi_{L,L}$ which is the identity map on L , under the same conditions.¹²

¹²The choice of these maps does not need to be so concrete, but the fact that it can be indicates for example that there is no use of choice here. We like the concreteness of this approach.

Remark 4.3. If N is a near-litter and $N \subseteq \text{dom}(\psi^*)$, then $N^\circ \in \text{dom}(\psi)$, and additionally $\psi^* \text{“} N \sim \psi(N^\circ)$. The converse is not true: $N^\circ \in \text{dom}(\psi)$ does not imply that $N \subseteq \text{dom}(\psi^*)$ [if this were true in general then ψ^* would be total so ψ would be total on litters], but it does imply that $N^\circ \subseteq \text{dom}(\psi^*)$. Thus, if L is a litter then $L \in \text{dom}(\psi)$ iff $L \subseteq \text{dom}(\psi^*)$ and under these circumstances, for any other litter M , we have $M = \psi(L)$ iff $M \sim \psi^*(L)$.

Note that if n is any integer, ψ^n is also a -1 -approximation with the same domain, where we take the convention that ψ^0 is the identity map on $\text{dom}(\psi)$. Using the condition $\pi_{M,N} \circ \pi_{L,M} = \pi_{L,N}$, we obtain the equation $(\psi^n)^* = (\psi^*)^n$.

From now on, we avoid using the action of ψ directly where possible, and instead use ψ^* , since the behavior of such functions is closer to the behavior of the various classes of permutations that we use throughout the paper.

Remark 4.4. ψ^* is a permutation of atoms. If it is defined on all of τ_{-1} , it is a -1 -allowable permutation.

Definition 4.5 (extension). We define a partial order on -1 -approximations by setting $\psi \preceq_{-1} \chi$ if $\psi \subseteq \chi$ and $\text{dom}(\psi) \cap \tau_{-1} = \text{dom}(\chi) \cap \tau_{-1}$. That is, χ may define images for more litters than ψ , but may not define images for any new atoms. If $\psi \preceq_{-1} \chi$, we may call χ an *extension* of ψ . Note that if $\psi \preceq_{-1} \chi$, then $\psi^* \subseteq \chi^*$.

Definition 4.6 (approximation, in general). If $-1 < \beta$, a β -approximation is defined as a function with $\{-1\} \cup \beta$ as domain such that $\psi(\gamma)$ is a γ -approximation for each γ in the domain. We write ψ_γ instead of $\psi(\gamma)$.

We define some operations on approximations.

1. If A is a finite subset of $\lambda \cup \{-1\}$ with maximum element β , we define the derivative of ψ along A in the following way: $\psi_A = (\psi_{A \setminus \{\min(A)\}})_{\min(A)}$ if $|A| \geq 2$, and $\psi_{\{\beta\}} = \psi$.
If $\beta = -1$, construe ψ_{-1} as ψ .
2. A β -approximation ψ acts on a support S by

$$(\psi^*[S])_\delta = (\psi_A^* \text{“} x, A) \text{ where } S_\delta = (x, A)$$

whenever this is defined for each $\delta \in \text{dom}(S)$.

3. We define the partial order $\preceq_\beta$ on β -approximations by defining $\psi \preceq_\beta \chi$ whenever $\psi_\gamma \preceq_\gamma \chi_\gamma$ for all $\gamma < \beta$, and define an *extension* of a β -approximation ψ as an approximation $\chi \succeq_\beta \psi$.

Definition 4.7 (flexibility). A near-litter N is *A-flexible*, where A is an extended type index, if $|A| \leq 2$ or N° is not in the range of any $f_{\gamma, \min(A_1)}$ for $-1 \leq \gamma < \min(A_2)$.

Definition 4.8 (coherent). Let $-1 \leq \beta \leq \alpha$. We say that a β -approximation ψ is *coherent* (c.f. the coherence condition on allowable permutations from definition 3.28) if:

1. If L is A -flexible for some A , then $\psi_A^* L$ is also A -flexible whenever this is defined.
2. If A is an extended type index with maximum β and $\min(A_1) = \delta < \beta$, and $\gamma < \min(A_2)$ and $(x, S) \in \tau_\gamma^+$ are such that $f_{\gamma, \delta}(x, S) \subseteq \text{dom}(\psi_A^*)$, then $(\psi_{A_2})_\gamma^*$ is defined on S , and there is some γ -allowable permutation π such that

$$(\psi_{A_2})_\gamma^*[S] = \pi[S]$$

and additionally that

$$\psi_A^* f_{\gamma, \delta}(x, S) \sim f_{\gamma, \delta}(\pi(x), \pi[S]).$$

The idea behind the definition is to ensure that ψ^* acts like the action of an allowable permutation by mirroring the coherence condition along all derivatives. We could have given this definition in a recursive fashion instead, defining an approximation as coherent if it satisfies the coherence condition at the top level and all of its derivatives are coherent. However, this formulation allows us to simplify some of the proofs in this section by avoiding recursion along λ .

Note that in the second condition above, if π' is any other γ -allowable permutation satisfying

$$(\psi_{A_2})_\gamma^*[S] = \pi'[S]$$

then we can conclude

$$\psi_A^* f_{\gamma, \delta}(x, S) \sim f_{\gamma, \delta}(\pi'(x), \pi'[S])$$

as π and π' have actions that coincide on S and therefore at x .

Remark 4.9. If ψ is coherent, then ψ_ϵ is coherent for any $\epsilon < \beta$, and every -1 -approximation is coherent.

We now show that composites and inverses of coherent approximations are coherent. This is analogous to the fact that allowable permutations form a subgroup of structural permutations.

Proposition 4.10. If ψ and ψ' are coherent such that all ψ_A and ψ'_A are defined on the same domain, then the approximation ψ'' defined by $\psi''_A = \psi_A \circ \psi'_A$ is also coherent.

Proof. Suppose that L is A -flexible. Then $(\psi'_A)^* L$ is A -flexible, so $\psi_A^* (\psi'_A)^* L$ is also flexible. Now suppose that $\gamma, \delta, (x, S)$ are as in the second condition for coherence, so $f_{\gamma, \delta}(x, S) \subseteq \text{dom}((\psi'_A)^{-1})^*$. Then there is π' such that

$$(\psi'_{A_2})_\gamma^*[S] = \pi'[S]; \quad (\psi'_A)^* f_{\gamma, \delta}(x, S) \sim f_{\gamma, \delta}(\pi'(x), \pi'[S]).$$

But then $f_{\gamma, \delta}(\pi'(x), \pi'[S])$ satisfies the same hypothesis for ψ , so we obtain π such that

$$(\psi_{A_2})_\gamma^*[\pi'[S]] = \pi[\pi'[S]]; \quad \psi_A^* f_{\gamma, \delta}(\pi(x), \pi[S]) \sim f_{\gamma, \delta}(\pi(\pi'(x)), \pi[\pi'[S]]).$$

Then, the permutation $\pi \circ \pi'$ satisfies the requirements:

$$(\psi''_{A_2})^*_\gamma[S] = (\pi \circ \pi')[S];$$

$$(\psi''_A)^* \text{“} f_{\gamma,\delta}(x, S) \sim \psi_A^* \text{“} f_{\gamma,\delta}(\pi'(x), \pi'[S]) \sim f_{\gamma,\delta}(\pi(\pi'(x)), \pi[\pi'[S]]). \quad \square$$

Proposition 4.11. If ψ is coherent, then ψ^{-1} is coherent.

Proof. Suppose that L is A -flexible, that $(\psi_A^{-1})^* \text{“} L$ is defined, but that it is not A -flexible. Then there is some $\gamma \neq \min(A_1) = \delta$ such that $f_{\gamma,\delta}(x, S) = (\psi_A^{-1})^* \text{“} L$. Now, by the second condition of coherence, there is a γ -allowable π such that

$$(\psi_{A_2})^*_\gamma[S] = \pi[S]; \quad \psi_A^* \text{“} f_{\gamma,\delta}(x, S) \sim f_{\gamma,\delta}(\pi(x), \pi[S]).$$

From this, we infer that

$$(\psi_{A_2}^{-1})^*_\gamma[\pi[S]] = S; \quad (\psi_A^{-1})^* \text{“} f_{\gamma,\delta}(\pi(x), \pi[S]) \sim f_{\gamma,\delta}(x, S) = (\psi^{-1})^* \text{“} L.$$

Therefore, $f_{\gamma,\delta}(\pi(x), \pi[S]) \sim L$, so L is indeed A -inflexible.

Now suppose that A , γ , δ , and (x, S) are as in the second condition, so $f_{\gamma,\delta}(x, S) \subseteq \text{dom}((\psi_A^{-1})^*)$. Consider $L = (\psi_A^{-1})^* \text{“} f_{\gamma,\delta}(x, S)$. If L is A -flexible, then $\psi_A^* \text{“} L$ would be A -flexible, a contradiction. Therefore, there are γ' , $(y, T) \in \tau_{\gamma'}^+$, and a γ' -allowable π' such that

$$L = f_{\gamma',\delta}(y, T); \quad (\psi_{A_2})^*_{\gamma'}[T] = \pi'[T]; \quad \psi_A^* \text{“} f_{\gamma',\delta}(y, T) \sim f_{\gamma',\delta}(\pi'(y), \pi'[T]).$$

Then

$$f_{\gamma,\delta}(x, S) = f_{\gamma',\delta}(\pi'(y), \pi'[T])$$

so by injectivity, $\gamma = \gamma'$, and $x = \pi'(y)$, and $S = \pi'[T]$. We thus have

$$(\psi_{A_2}^{-1})^*_\gamma[S] = T = (\pi')^{-1}[S].$$

So set $\pi = (\pi')^{-1}$, and verify:

$$(\psi_A^{-1})^* \text{“} f_{\gamma,\delta}(x, S) \sim f_{\gamma,\delta}(y, T) = f_{\gamma,\delta}((\pi')^{-1}(x), (\pi')^{-1}[T])$$

as required. $\square$

Definition 4.12 (approximate). A -1 -approximation ψ is said to *approximate* a -1 -allowable permutation π if

1. $\psi^* \subseteq \pi$; and
2. if $a \notin \text{dom}(\psi)$ and L is a litter, we have $a \in L$ if and only if $\pi(a) \in (\pi \text{“} L)^\circ$.

If $-1 < \beta$, a β -approximation ψ is said to *approximate* a β -allowable permutation π if for each $\gamma < \beta$, ψ_γ approximates π_γ .

Remark 4.13. If $\psi \preceq_\beta \chi$ and χ approximates π , then ψ approximates π .

Definition 4.14 (freedom of action). We say that *freedom of action* holds at some type index β if every coherent β -approximation ψ approximates some β -allowable permutation π .

Remark 4.15. If ψ is a coherent approximation such that ψ_A is defined on all litters (or equivalently, ψ_A^* is defined on all atoms) for all A containing -1 , then (I9) implies that ψ approximates a unique allowable permutation π , and π is given by $\pi_A = \psi_A^*$. So to prove that freedom of action holds at level β , it suffices by remark 4.13 to show that every coherent β -approximation has a coherent extension χ such that χ_A is defined on all litters for all A containing -1 .

Lemma 4.16. Let ψ be a coherent β -approximation. Then ψ admits a coherent extension χ such that for each extended type index A with maximum β , χ_A is defined on all A -flexible litters.

Proof. The construction

$$\chi_A = \psi_A \cup \{(L, L) : L \notin \text{dom}(\psi_A^*), L \text{ is } A\text{-flexible}\}$$

suffices. □

Remark 4.17. This lemma shows in particular that freedom of action holds at type -1 . Indeed, suppose that ψ is a (coherent) -1 -approximation and χ is an extension as above. By remark 4.4, χ^* is a -1 -allowable permutation as all near-litters are $\{-1\}$ -flexible, and χ clearly approximates it. But $\psi \preceq_{-1} \chi$, so ψ also approximates χ^* (remark 4.13).

Lemma 4.18. Let ψ be a coherent β -approximation. Let A be an extended type index with maximum β , and let $\delta = \min(A_1)$ and $\gamma < \min(A_2)$. Let $(x, S) \in \tau_\gamma^+$ be such that $(\psi_{A_2})_\gamma^*[S]$ is defined.

Then if freedom of action holds at type level γ , there is a coherent extension $\chi \succeq_\beta \psi$ such that $\chi_A^* f_{\gamma, \delta}(x, S)$ is defined.

Proof. The γ -approximation $(\psi_{A_2})_\gamma$ is coherent (remark 4.9), so it approximates some γ -allowable permutation π by freedom of action. In particular, $((\psi_{A_2}^n)_\gamma)^*[S] = \pi^n[S]$ for every integer n . We intend to add the orbit

$$f_{\gamma, \delta}(\pi^n(x), \pi^n[S]) \mapsto f_{\gamma, \delta}(\pi^{n+1}(x), \pi^{n+1}[S])$$

to ψ_A^* , at least up to nearness.

Suppose that there is some n such that ψ_A^* is defined on $f_{\gamma, \delta}(\pi^n(x), \pi^n[S])$. We have $((\psi_{A_2}^n)_\gamma)^*[S] = \pi^n[S]$, so $((\psi_{A_2}^{-n})_\gamma)^*[\pi^n[S]] = S$. Therefore, as ψ^{-n} is coherent (by propositions 4.10 and 4.11), we obtain

$$(\psi^{-n})_A^* f_{\gamma, \delta}(\pi^n(x), \pi^n[S]) \sim f_{\gamma, \delta}(x, S).$$

Therefore,

$$f_{\gamma, \delta}(x, S) \subseteq \text{dom}(\psi_A^*).$$

So we do not need to extend ψ ; we are already done.

Otherwise, we can extend ψ to an approximation χ given by

$$\chi_A = \psi_A \cup \{(f_{\gamma,\delta}(\pi^n(x), \pi^n[S]), f_{\gamma,\delta}(\pi^{n+1}(x), \pi^{n+1}[S])) : n \in \mathbb{Z}\}$$

and $\chi_B = \psi_B$ for all $B \neq A$. It is easy to see that χ is an extension of ψ , and

$$\chi_A^* f_{\gamma,\delta}(\pi^n(x), \pi^n[S]) \sim f_{\gamma,\delta}(\pi^{n+1}(x), \pi^{n+1}[S])$$

for each integer n . Finally, we must check that this χ is coherent, but this holds by construction. $\square$

Theorem 4.19 (Freedom of Action). Freedom of action holds at all type indices $\beta \leq \alpha$.¹³

Proof. By induction, we may assume freedom of action holds at all levels $\delta < \beta$. Let ψ be a coherent β -approximation, and use Zorn's lemma to extend ψ to a maximal coherent extension χ ; this step uses the fact that coherence is preserved under suprema of chains of approximations.

Suppose that χ_A^* is not defined on all litters for some A . Let L be the litter with minimal position $\iota_*(L)$ such that there is an extended type index A with maximum element β such that $L \not\subseteq \text{dom}(\chi_A^*)$.

Suppose that L is A -flexible. By lemma 4.16, χ admits a coherent extension φ such that $L \subseteq \text{dom}(\varphi_A^*)$. This contradicts maximality of χ .

Now suppose that L is not A -flexible. Writing $\gamma = \min(A_1)$, there are $\delta < \min(A_2)$ and $(x, S) \in \tau_{\delta}^+$ such that $L = f_{\delta,\gamma}(x, S)$.

We claim that $(\chi_{A_2})_{\delta}^*[S]$ is defined; this will then give a contradiction by lemma 4.18. To show this, we must prove that if $(t, B) \in \text{rng}(S)$, then $(\chi_{A_2})_B^* t$ is defined. By definition 3.18, we have $\iota_*^+(x, S) < \iota_*(L)$, and by (I8), $\iota_*(t) \leq \iota_*^+(x, S)$. If t is a singleton of an atom and $t \subset M$ for a litter M , then by lemma 3.14, we obtain $\iota_*(M) < \iota_*(t)$, and therefore that $(\chi_{A_2})_B^* t$ is defined by minimality of the position of L . Alternatively, if t is a near-litter, then if M is any litter such that $M \cap t \neq \emptyset$, lemma 3.14 again implies that $\iota_*(M) \leq \iota_*(t)$, so $(\chi_{A_2})_B^* M$ is defined. Combining all such litters, we conclude that $(\chi_{A_2})_B^* t$ is defined.

Therefore χ_A^* must be defined on all litters for all A . By remark 4.15, this concludes the proof: ψ approximates the allowable permutation π given by $\pi_A = \chi_A^*$. $\square$

Remark 4.20. The use of Zorn's lemma in the previous proof is merely a technical convenience. Its use can be excised by instead computing the value of χ_A^* at each atom directly, under the inductive hypothesis that its value at each atom at an earlier position was already computed for all A .

In fact, all of the proofs of this subsection can be phrased in such a way that the axiom of choice is not invoked. In particular, the coherent extension

¹³It should be noted that the second author has entirely reorganized and rewritten this proof, and indeed the entire section on Freedom of Action. The underlying mathematics is the same but the presentation is much cleaner.

defined in 4.18 is uniquely determined: all choices of π yield the same extension χ . This means that every coherent approximation defined on all flexible litters has a unique coherent extension defined on all litters (but we will not use this fact).

4.2 Freedom of Action for incomplete orbits

In this subsection, we give an alternative form of the Freedom of Action Theorem (4.19) that is more useful when we do not already have full orbits of atoms or litters for use in approximations.

Definition 4.21 (interference). Let $x, y \subseteq \tau_{-1}$. Their *interference* is defined to be the union of the small elements of $\{x \Delta y, x \cap y\}$, which is a small subset of τ_{-1} .

Thus the interference between two near-litters M and N is either $M \Delta N$ or $M \cap N$, whichever is small.

Definition 4.22 (actions). A -1 -action is a function ξ such that:

1. The domain and image of ξ are small, and ξ is injective.
2. Each domain element of ξ is either an atom or a near-litter, and moreover, ξ maps atoms to atoms and near-litters to near-litters.
3. If ξ is defined at near-litters N_1 and N_2 , then $N_1^\circ = N_2^\circ$ if and only if $\xi(N_1)^\circ = \xi(N_2)^\circ$. This ensures that ξ gives rise to an injective partial function on litters.
4. Whenever $a, N \in \text{dom}(\xi)$, we have $a \in N$ if and only if $\xi(a) \in \xi(N)$.
5. The interference of near-litters in the domain of ξ is included in the domain of ξ , and the interference of near-litters in the range of ξ is included in the range of ξ .

Note that this definition ensures that the inverse of a -1 -action is a -1 -action.

Definition 4.23. For $-1 < \beta$, a β -action is a function with $\{-1\} \cup \beta$ as domain such that $\xi(\gamma)$ is a γ -action for each $\gamma < \beta$.

Definition 4.24. We define derivatives for actions in the same way as we do for approximations. For convenience, if a is an atom we define $\xi_A(\{a\})$ as $\{\xi_A(a)\}$. Actions can be applied to supports as follows:

$$(\xi[S])_\delta = (\xi_A(x), A) \text{ where } S_\delta = (x, A)$$

As with approximations, this operation is partial. An action is called *coherent* if:

1. If N is A -flexible for some A , then $\xi_A(N)$ is also A -flexible whenever this is defined.

2. If A is an extended type index with maximum β and $\min(A_1) = \delta < \beta$, and $\gamma < \min(A_2)$ and $(x, S) \in \tau_\gamma^+$ are such that $f_{\gamma, \delta}(x, S) \sim N \in \text{dom}(\xi_A)$, then there is some γ -allowable permutation π such that

$$(\xi_{A_2})_\gamma[S] = \pi[S] \text{ and } \xi_A(N) \sim f_{\gamma, \delta}(\pi(x), \pi[S]).$$

A β -action ξ is said to approximate a β -allowable permutation π if $\xi_A(x) = \pi_A x$ for all x in the domain of ξ_A .

Theorem 4.25 (Freedom of Action for incomplete orbits). Every coherent β -action approximates some β -allowable permutation π .

Proof. Let ξ be a coherent β -action. First, note that we may assume without loss of generality that whenever $N \in \text{dom}(\xi_A)$, we have

$$N \Delta N^\circ \subseteq \text{dom}(\xi_A) \text{ and } \xi_A(N) \Delta \xi_A(N)^\circ \subseteq \text{rng}(\xi_A). \quad (*)$$

We can extend any ξ to an action ξ' with this property. If there is an atom in $N \setminus N^\circ \setminus \text{dom}(\xi_A)$ then N° is not in the domain of ξ and all $N' \sim N$ which are in the domain of ξ contain the atom, and we can map each such atom to an unused atom which is in $\xi_A(N)^\circ$ and also inside all ξ_A -images of near-litters $N' \sim N$. Doing this for all such atoms, we can define ξ on each $N \setminus N^\circ \setminus \text{dom}(\xi_A)$. Next, if there is an atom in $N^\circ \setminus N \setminus \text{dom}(\xi_A)$, then no near-litter containing the atom can be in the domain of ξ since in that case we must have already defined ξ at this atom, and so we can map each such atom into an unused litter. This handles near-litters N in the domain of ξ ; near-litters in the range of ξ are handled dually. This modification preserves coherence, as it only adds images of atoms.

Now, we will turn ξ into a β -approximation ψ . To do this, we need to fill in orbits of litters and atoms.

We first show how to define orbits of flexible litters; the orbits for inflexible litters will appear automatically by coherence. For each A , let g_A be the partial function given by $g_A(N^\circ) = \xi_A(N)^\circ$ whenever N is A -flexible and $\xi_A(N)$ is defined; the third condition on the definition of actions ensures that this map is well-defined and injective. We define the restriction of ψ_A to litters to be a small bijective extension of g_A , where any new litters in $\text{dom}(\psi_A)$ are A -flexible and disjoint from domain and range elements of ξ_A .

Next, we show how to define orbits of atoms. For each A , we define the restriction of ψ_A to atoms to be a small bijective extension of the restriction of ξ_A to atoms, with the property that $a \in N$ iff $\psi_A(a) \in \xi_A(N)$ whenever the right-hand side is defined. We now explain how to do this. Our construction proceeds in ω stages, where at each stage, we define ψ_A and ψ_A^{-1} at any atoms contained in the domain or range of the function defined so far. We describe the construction of images under ψ_A ; the inverse is dual.

Suppose that we need to define $\psi_A(a)$, so in particular $a \notin \text{dom}(\xi_A)$, and that $\xi_A(N)$ is defined for some $N \ni a$. By $(*)$, we know that $a \in N^\circ$. Additionally, we know that $a \in N'$ for every $N' \sim N$ with $N' \in \text{dom}(\xi_A)$, and $a \notin N''$ for

every $N'' \approx N$ with $N'' \in \text{dom}(\xi_A)$, because of the interference conditions. So we define $\psi_A(a)$ to be any unused atom in $\xi_A(N)^\circ$ that is inside all $\xi_A(N')$ for $N' \approx N$, but outside all $\xi_A(N'')$ for $N'' \approx N$. Such an atom will always exist because there are only a small amount of atoms excluded by each near-litter N' , and there are only a small number of such near-litters.

If $\xi_A(N)$ is not defined for any $N \ni a$, then we are free to choose any unused atom for $\psi_A(a)$.

We perform the dual construction to define ψ_A^{-1} on required atoms. When the construction concludes, the function ψ_A will still be small because κ is uncountable and regular.

The β -approximation ψ defined as above is obviously coherent because all of the litters L in the domain of ψ_A are A -flexible, so by the Freedom of Action Theorem (4.19), ψ approximates some β -allowable permutation π .

We claim that ξ approximates this π . The case for atoms is obvious, so it suffices to show that for each near-litter N such that $\xi_A(N)$ is defined, we have $\xi_A(N) = \pi_A \text{``} N$ . By induction we may assume that  $\xi_B(M) = \pi_B \text{``} M$ holds for all extended type indices B and near-litters M with $\iota_*(M) < \iota_*(N)$.

First, we show that $\xi_A(N)^\circ = (\pi_A \text{``} N)^\circ$ . If  $N$  is  $A$ -flexible, we have  $\xi_A(N)^\circ = \psi_A(N^\circ) = (\pi_A \text{``} N)^\circ$ as required. Suppose instead that $\gamma = \min(A_1)$ and $\delta < \min(A_2)$ are such that $N \sim f_{\delta,\gamma}(x, S)$ for $(x, S) \in \tau_\delta^+$. For any (t, B) in the range of S , by coherence of ξ we know that $(\xi_{A_2})_B(t)$ is defined, and moreover by inductive hypothesis, that $(\xi_{A_2})_B(t) = (\pi_{A_2})_B \text{``} t$. But this shows that ξ_{A_2} agrees with π_{A_2} on S , and so again by coherence, we can conclude that

$$\xi_A(N) \sim f_{\delta,\gamma}((\pi_{A_2})_\delta(x), (\pi_{A_2})_\delta[S]) \sim \pi_A \text{``} f_{\delta,\gamma}(x, S) \sim \pi_A \text{``} N.$$

We now conclude the proof by using $(*)$ to show that $\xi_A(N) = \pi_A \text{``} N$.

Suppose that $a \in \xi_A(N)$. If $a \in \text{rng}(\psi_A)$, then $\psi_A^{-1}(a) \in N$ by construction of ψ_A , but $\psi_A^{-1}(a) = \pi_A^{-1}(a)$, so $\pi_A^{-1}(a) \in N$. Otherwise, $a \in \xi_A(N)^\circ$ by $(*)$, so by the second condition in definition 4.12 we have $\pi_A^{-1}(a) \in (\pi_A^{-1} \text{``} (\xi_A(N)^\circ))^\circ = N^\circ$, so $\pi_A^{-1}(a) \in N$ by $(*)$.

Conversely, suppose that $a \in \pi_A \text{``} N$ . If  $a \in \text{rng}(\psi_A)$ , then  $\pi_A^{-1}(a) = \psi_A^{-1}(a)$ , so  $a \in \xi_A(N)$  by construction of  $\psi_A$ . Otherwise,  $\pi_A^{-1}(a) \in N^\circ$  by  $(*)$ , so by the same condition, we have  $a \in (\pi_A \text{``} N^\circ)^\circ = \xi_A(N)^\circ$, so by $(*)$ we conclude that $a \in \xi_A(N)$. $\square$

4.3 Counting orbits of supports: coding functions and specifications

Now we argue that type α is of size μ , which ensures that the construction actually succeeds at every type; verification of the remaining inductive hypotheses for stage α of the construction is thus completed.

Definition 4.26 (strong support). A β -support S is called *strong* if it satisfies the additional properties that

1. if (x, A) and (y, A) are in the range of S , then $(\{z\}, A)$ is in the range of S for all atoms z in the interference of x and y (defined in definition 4.21),
2. and if (x, A) is in the range of S , $\delta = \min(A_1)$, $\gamma < \min(A_2)$, and $x^\circ = f_{\gamma, \delta}(y, T)$, then the range of $T^{\uparrow A_2}$ is a subset of the range of S : supports appearing in inverse images under f of litters which are near the first projections of an element of the support have a type-raised copy (mod reindexing) appearing in the support.

Remark 4.27. It should be evident that if π is a β -allowable permutation and S is a strong β -support, $\pi[S]$ is also a strong β -support.

Proposition 4.28. Every β -support S is an initial segment of some strong support.

Proof. First, for each (x, A) in the range of S where $\delta = \min(A_1)$, $\gamma < \min(A_2)$ and $x^\circ = f_{\gamma, \delta}(y, T)$, add the range of $T^{\uparrow A_2}$ to the end of S . Repeating this process ω times, we obtain a new support $S' \supseteq S$ that satisfies the second property of being a strong support; the domain of S' is still a small ordinal as κ is uncountable and regular. Finally, we append to S all pairs $(\{z\}, A)$ where z is in the interference of near-litters M, N where $(M, A), (N, A)$ are in the range of S . This produces a support $S'' \supseteq S'$ which is clearly strong. $\square$

Incidentally, we can now revisit an earlier result (proposition 3.40) that we could not prove in full generality:

Proposition 4.29. Let $x \in \tau_\beta$ for $\beta < \alpha$. Then $x \cap \tau_{-1}$ has small symmetric difference from a union of a small or co-small set of litters.

Proof. Let S be a strong support for x . Following the strategy from proposition 3.40, there is a partition P of τ_{-1} defined by the entries of S with second component $\{-1, \beta\}$. Suppose that a and b are atoms that lie in the same compartment of P . Define a β -action ξ such that $\xi_A(x) = x$ whenever (x, A) appears in the range of S , but also $\xi_{-1}(a) = b$ and $\xi_{-1}(b) = a$. This is a well-defined and coherent action as S is strong. So by Freedom of Action for incomplete orbits, there is such a β -allowable π that fixes S (and hence x) but has $\pi_{-1}(a) = b$ and $\pi_{-1}(b) = a$. This implies that $a \in x \cap \tau_{-1}$ if and only if $b \in x \cap \tau_{-1}$, and we can conclude in the same way as in proposition 3.40. $\square$

Definition 4.30 (coding function). For any $(x, S) \in \tau_\beta^+$, we can define a function $\chi_{x,S}$ which sends $T = \pi[S]$ to $\pi(x)$ for every T in the orbit of S under the action of allowable permutations. We call such functions *β -coding functions*. Note that if $\pi[S] = \pi'[S]$ then $(\pi^{-1} \circ \pi')[S] = S$, so $(\pi^{-1} \circ \pi')(x) = x$, so $\pi(x) = \pi'(x)$, ensuring that the map $\chi_{x,S}$ for which we gave an implicit definition is well defined.

We will now define a certain combinatorial object that will characterize the orbit of a strong support under the action of allowable permutations (although we can define it even when the support in question is not strong).

Definition 4.31 (specification). The *specification* of a β -support S is the function S^* with the same domain as S given by the following criteria. We use the notation S_ϵ^* for $S^*(\epsilon)$.

1. If S_ϵ is $(\{x\}, A)$, then S_ϵ^* is given by

$$S_\epsilon^* = (0, A, \{\zeta : x \in \pi_1(S_\zeta) \wedge A = \pi_2(S_\zeta)\}).$$

2. If S_ϵ is (N, A) where N is A -flexible, then S_ϵ^* is given by

$$S_\epsilon^* = (1, A, \{\zeta : (\exists N' : S_\zeta = (N', A) \wedge N \sim N')\}).$$

3. If S_ϵ is (N, A) where N is not A -flexible, so $|A| > 2$, $\delta = \min(A_1)$, and $N \sim f_{\gamma, \delta}(x, T)$ with $-1 \leq \gamma < \min(A_2)$, then S_ϵ^* is given by

$$S_\epsilon^* = (2, A, \chi_{x,T}, \{(\zeta, \eta) : S_\zeta = (T^{\uparrow A_2})_\eta\}).$$

Remark 4.32. It should be evident that for any β -allowable permutation π and β -support S , $(\pi[S])^* = S^*$. What is less evident and our first target result here is that if S is a strong support then any strong T with $T^* = S^*$ is the image of S under the action of an allowable permutation: the specifications precisely code the orbits in the strong supports under the allowable permutations. We first introduce some auxiliary notions.

Definition 4.33. Let S and T be β -supports. Define the relation $\mapsto_A^{S,T}$ between atoms and near-litters by:

1. $a \mapsto_A^{S,T} a'$ whenever $\exists \delta : S_\delta = (\{a\}, A) \wedge T_\delta = (\{a'\}, A)$; and
2. $N \mapsto_A^{S,T} N'$ whenever $\exists \delta : S_\delta = (N, A) \wedge T_\delta = (N', A)$.

For strong β -supports S and T with the same specification, we are going to define a β -action ξ by $\xi_A(x) = y$ whenever $x \mapsto_A^{S,T} y$. Once we show that this is well-defined and coherent, we will be able to invoke the Freedom of Action Theorem (4.25) to produce a β -allowable permutation π such that $\pi[S] = T$.

For the rest of this subsection, fix strong β -supports S and T such that $S^* = T^*$, and abbreviate $\mapsto_A^{S,T}$ as $\mapsto_A$.

Lemma 4.34. If $S_\delta = (a, A)$, then there is at least one a' such that $a \mapsto_A a'$ (and similarly for near-litters).

Proof. As S and T have the same specification, if $S_\delta = (x, A)$, then T_δ is well-defined and equal to (y, A) , and moreover x is a singleton iff y is a singleton. In particular, the positions δ of S containing (singletons of) atoms along a given path A are exactly the same as the positions of T containing atoms along A , and the same holds for near-litters. $\square$

Lemma 4.35. If $a \mapsto_A a'$ and $b \mapsto_A b'$, then $a = b$ if and only if $a' = b'$.

Proof. Suppose that a is an atom such that both $a \mapsto_A b$ and $a \mapsto_A c$. This would mean that $S_\delta = (\{a\}, A)$ and $T_\delta = (\{b\}, A)$, but additionally that $S_\epsilon = (\{a\}, A)$ and $T_\epsilon = (\{c\}, A)$. But the specification for S contains the information that the atoms at index δ and ϵ are equal (condition 1 of definition 4.31), so as S and T share a specification, the atoms b and c must also be equal. Running the same argument with the roles of S and T swapped shows the converse. $\square$

Lemma 4.36. Suppose that $N_1 \mapsto_A N'_1$ and $N_2 \mapsto_A N'_2$. Then $N_1 \sim N_2$ if and only if $N'_1 \sim N'_2$.

Proof. We only prove the forward direction; the converse follows as above by interchanging S and T .

If N_1 is A -flexible, the result is easy to deduce: S^* directly contains the information that the near-litters at indices for N_1 and N_2 are near, and so T^* contains the information that N'_1 is near N'_2 , as required.

Now suppose that N_1 is A -inflexible, so $N_1 \sim f_{\gamma, \delta}(x, U)$ where $\delta = \min(A_1)$ and $-1 \leq \gamma < \min(A_2)$. By condition 3 of definition 4.31, S^* contains the information that N_1 is A -inflexible, as well as the coding function $\chi_{x, U}$. We can extract γ from this coding function, so we can conclude that N'_1 is also A -inflexible, with $N'_1 \sim f_{\gamma, \delta}(x', U')$ for $\chi_{x, U} = \chi_{x', U'}$. Hence, $x' = \chi_{x', U'}(U') = \chi_{x, U}(U')$, so there is some γ -allowable π_1 such that $U' = \pi_1[U]$ and $x' = \pi_1(x)$. Now, N_2 is also near $f_{\gamma, \delta}(x, U)$, so by the same argument, there must be some π_2 such that $N'_2 \sim f_{\gamma, \delta}(\pi_2(x), \pi_2[U])$. So it remains to prove that $\pi_1[U] = \pi_2[U]$.

Suppose that $U_\eta = (y, B)$. Since S is strong, there is some ζ such that $S_\zeta = (U^{\uparrow A_2})_\eta$. Since S^* stores this pair (ζ, η) and $S^* = T^*$, we must have $T_\zeta = ((\pi_1[U])^{\uparrow A_2})_\eta$. But by the same logic, we must have $T_\zeta = ((\pi_2[U])^{\uparrow A_2})_\eta$, so $\pi_1[U]_\eta = \pi_2[U]_\eta$, as required. $\square$

Corollary 4.37. If $N_1 \mapsto_A N'_1$ and $N_2 \mapsto_A N'_2$, then $N_1 = N_2$ if and only if $N'_1 = N'_2$.

Proof. If $N_1 = N_2$, then by the previous lemma, we have $N'_1 \sim N'_2$. Suppose that $N'_1 \neq N'_2$, so without loss of generality, assume there is some atom $a \in N'_1 \setminus N'_2$. Since T is strong, this atom must appear in T : there is δ such that $T_\delta = (\{a\}, A)$. The specification T^* contains the information that the atom at position δ is contained in the near-litter at the position of N'_1 but not the near-litter at the position of N'_2 . Since $S^* = T^*$, the same holds of S : the unique atom b such

that $b \mapsto_A a$ must be contained in N_1 but not N_2 , a contradiction. The other direction follows by symmetry. $\square$

Lemma 4.38. The relations $\mapsto_A$ assemble into a β -action ξ defined by $\xi_A(x) = y$ if and only if $x \mapsto_A y$.

Proof. The only condition of definition 4.22 left to check is that $a \in N$ if and only if $\xi_A(a) \in \xi_A(N)$, but this follows directly from the data contained in condition 1 of specifications. $\square$

Lemma 4.39. The action ξ as defined is coherent.

Proof. Suppose that N is A -flexible. Then S^* stores this information, so $\xi_A(N)$ must also be A -flexible as $S^* = T^*$.

Suppose instead that $N \sim f_{\gamma,\delta}(x, U)$ and $N \in \text{dom}(\xi_A)$ where $\min(A_1) = \delta$. Then as S and T have the same specification, we know that $\xi_A(N) \sim f_{\gamma,\delta}(y, V)$ where $\chi_{x,U} = \chi_{y,V}$. So there must be some γ -allowable π such that $y = \pi(x)$ and $V = \pi[U]$, and this is precisely what is needed for the definition of coherence for actions. $\square$

We are now ready to prove:

Proposition 4.40. The specification of a strong β -support exactly determines the orbit in the action of β -allowable permutations on supports to which it belongs: if two strong β -supports have the same specification, they are in the same orbit.

Proof. Let S and T be strong β -supports. We have already seen that $\pi[S] = T$ implies that $S^* = T^*$. For the converse, suppose that $S^* = T^*$. By the results of this subsection, there is a coherent β -action ξ such that $\xi_A(x) = y$ whenever $S_\delta = (x, A)$ and $T_\delta = (y, A)$. Then, by the Freedom of Action Theorem (4.25), there is a β -allowable permutation π such that $\pi[S] = T$. $\square$

Proposition 4.41. Let $\beta \leq \alpha$. Suppose that we already know that there are $< \mu$ γ -coding functions for each $\gamma < \beta$ (which we will be able to assume by inductive hypothesis). Then there are $< \mu$ specifications of β -supports.

Proof. The elements of the range of a β -specification are built from finite sets of type indices below β , sets of (pairs of) elements of κ (support domain elements), and γ -coding functions for $\gamma < \beta$. As $\beta < \lambda \leq \mu$, there are less than μ possible finite sets of type indices below β . Since $\beta < \lambda \leq \text{cf}(\mu)$, the sum of cardinals each $< \mu$ indexed by ordinals $< \beta$ will be less than μ , so there are less than μ possible coding functions that can be used here. Hence the elements of the range of such a specification are taken from a set of size less than μ , and so the cardinality of the set of specifications is bounded by $\nu^{<\kappa}$ where $\nu < \mu$. Now, $\nu^{<\kappa}$ is the cardinality of the set of functions from some small ordinal to ν , which is less than or equal to the cardinality of the power set of $\nu \times \kappa$, which in turn is less than μ because μ is strong limit. $\square$

Proposition 4.42. Suppose that we already know that there are $< \mu$ γ -coding functions for each $\gamma < \beta$. Then there are less than μ orbits in supports under β -allowable permutations.

Proof. For any support S that is not necessarily strong, there is a suitable weaker notion of specification: give the specification for a strong support T of which S is an initial segment, together with the domain of S . These weak specifications are not unique, but they do determine the orbit in which S lies in the allowable permutations. Since there are less than μ specifications by proposition 4.41, and $\kappa < \mu$, we conclude that there are less than μ orbits in the β -supports under the allowable permutations. $\square$

4.4 Types are of size μ

The strategy of our argument for the size of the types is to show that there are $< \mu$ coding functions for each type, which implies that there are no more than μ (and so exactly μ) elements of each type, since every element of a type is obtainable by applying a coding function (of which there are $< \mu$) to a support (of which there are μ).

We will prove this cardinality bound on coding functions by induction on $\beta \leq \alpha$. It is important to note that this inductive proof occurs internally to a single instance of the main inductive step; that is, α remains fixed, and we do not need to record the fact that the bound on coding functions holds in our main list of inductive hypotheses (I1)–(I9).

This argument will proceed by showing that all coding functions can be written in a particular form, in terms of coding functions at some lower level $-1 \neq \gamma < \beta$. In particular, this result will depend on the existence of some ordinal $\gamma \in \beta$, so we prove the result for types 0 and -1 explicitly here first.

Lemma 4.43. There are less than μ coding functions for type 0 and for type -1 .

Proof. We describe all coding functions for type 0. The orbit of a strong 0-support in the allowable permutations is determined by the positions in the support occupied by near-litters, and for each position in the support occupied by a singleton, the positions, if any, of the near-litters in the support which include it, and for each position in the support occupied by a near-litter, the positions, if any, occupied by near-litters which are near it (this is exactly the information contained in the 0-specification of the support). There are no more than 2^κ ways to specify an orbit. Now for each such support, there is a natural partition of type -1 into near-litters, singletons, and a large complement set; we define this partition explicitly in the proof of proposition 3.40. The partition has $\nu < \kappa$ elements, and there are $2^\nu \leq 2^\kappa$ objects of type 0 with this support, each determined by specifying for each compartment in the partition whether it is to be included or excluded from the set. So there are $2^\nu \leq 2^\kappa$ coding functions whose domain is the orbit of this support, and so for any support (that need not be strong), there are at most 2^κ coding functions whose domain contains

it, by first extending it to a strong support. So there are no more than $2^\kappa < \mu$ coding functions over type 0.

Any type -1 coding function is associated with a uniquely determined type 0 coding function of a singleton, so there are no more of the former than of the latter. $\square$

With those two cases handled, we return to focus on the case where there is some ordinal $\gamma \in \beta$.

Definition 4.44 (combination). Let $\gamma \in \beta$ be ordinals at most α . Let s be a set of β -coding functions. We define the *combination* of s along γ to be the function $C_{s,\gamma}$ from β -supports to τ_β given by

$$C_{s,\gamma}(S) = \left(\bigcup_{\chi \in s} \bigcup_{T \supseteq S} \chi(T) \cap \tau_\gamma \right)_\beta.$$

That is, for a support S , the γ -extension of the model element $C_{s,\gamma}(S)$ is precisely the union of the γ -extensions of model elements obtained by applying coding functions from s to supports extending S .

That $C_{s,\gamma}(S)$ is always well-defined requires proof: by symmetric comprehension (proposition 3.34), it suffices to show that when a β -allowable permutation π fixes S , π_γ fixes the set

$$\bigcup_{\chi \in s} \bigcup_{T \supseteq S} \chi(T) \cap \tau_\gamma,$$

but this follows from the fact that for any β -allowable π ,

$$\begin{aligned} \pi_\gamma \left(\bigcup_{\chi \in s} \bigcup_{T \supseteq S} \chi(T) \cap \tau_\gamma \right) &= \bigcup_{\chi \in s} \bigcup_{T \supseteq S} \pi_\gamma (\chi(T) \cap \tau_\gamma) \\ &= \bigcup_{\chi \in s} \bigcup_{T \supseteq S} \pi(\chi(T)) \cap \tau_\gamma \\ &= \bigcup_{\chi \in s} \bigcup_{T \supseteq S} \chi(\pi[T]) \cap \tau_\gamma \\ &= \bigcup_{\chi \in s} \bigcup_{T \supseteq \pi[S]} \chi(T) \cap \tau_\gamma. \end{aligned}$$

The same calculation shows that for any orbit of β -supports o under the action of allowable permutations, the restriction $C_{s,\gamma} \upharpoonright o$ is a coding function.

Our claim is that every β -coding function can be written in this form, where the elements of the set s are particular coding functions of singletons of objects of type γ . We will make the reference of “particular” here more precise.

Definition 4.45 (raised singleton). Let $\gamma \in \beta$. A β -coding function χ is called a *raised γ -singleton* if there is a β -support S and $(x, U) \in \tau_\gamma^+$ such that $\chi = \chi_{\{x\}_\beta, S+U^{\uparrow\beta}}$.

Lemma 4.46. Let $\gamma \in \beta$ be ordinals at most α , and let χ be a β -coding function. Then there is a set s of raised γ -singletons such that $\chi = C_{s,\gamma} \upharpoonright \text{dom}(\chi)$.

Proof. Let $\chi = \chi_{x,S}$. The set we claim works is

$$s = \{\chi_{\{y\}_\beta, S+U^{\uparrow\beta}} : y \in x \cap \tau_\gamma \wedge U \text{ supports } y\}.$$

Since $C_{s,\gamma} \upharpoonright \text{dom}(\chi)$ is a coding function, it suffices to check that $C_{s,\gamma}(S) = x$, which we show by γ -extensionality.

If $y \in C_{s,\gamma}(S) \cap \tau_\gamma$, then there is some $z \in x \cap \tau_\gamma$, some support U of z , and $T \supseteq S$ such that

$$y \in \chi_{\{z\}_\beta, S+U^{\uparrow\beta}}(T).$$

Hence there is a β -allowable π such that $\pi[S+U^{\uparrow\beta}] = T$; in particular, $\pi[S] = S$ as T is an extension of S , so $\pi(x) = x$. Therefore,

$$\pi_\gamma^{-1}(y) \in \chi_{\{z\}_\beta, S+U^{\uparrow\beta}}(\pi^{-1}[T]) = \chi_{\{z\}_\beta, S+U^{\uparrow\beta}}(S+U^{\uparrow\beta}) = \{z\}_\beta,$$

giving $\pi_\gamma^{-1}(y) \in x$ and so $y \in x$ as required.

Conversely, if $y \in x \cap \tau_\gamma$, then for any support U of y , we immediately obtain $y \in C_{s,\gamma}(S)$ from the fact that $\chi_{\{y\}_\beta, S+U^{\uparrow\beta}} \in s$. $\square$

Lemma 4.47. There are less than μ coding functions for every type index $\beta \leq \alpha$.

Proof. We may assume by induction that there are less than μ coding functions of every type $\gamma < \beta$. Since the cases for $\beta = 0, -1$ were proven in lemma 4.43, we may additionally assume that there exists some ordinal $\gamma \in \beta$. By lemma 4.46, every β -coding function can be written in the form $C_{s,\gamma} \upharpoonright d$, where s is a set of raised γ -singletons and d is an orbit of β -supports under the action of allowable permutations. By proposition 4.42, there are less than μ such orbits, so it suffices to show there are less than μ possible sets of raised γ -singletons; indeed, as μ is a strong limit, it further suffices to simply show that there are less than μ raised γ -singletons.

We show that any raised singleton $\chi_{\{x\}_\beta, S+U^{\uparrow\beta}}$ is determined by $\text{dom}(S)$, the orbit of $S+U^{\uparrow\beta}$ under the allowable permutations, and the γ -coding function $\chi_{x,U}$. Indeed, suppose that S and T have the same domain, $S+U^{\uparrow\beta}$ and $T+V^{\uparrow\beta}$ share the same orbit under the β -allowable permutations, and the coding functions $\chi_{x,U}$ and $\chi_{y,V}$ coincide. Then there is a β -allowable π such that $\pi[S] = T$ and $\pi_\gamma[U] = V$, and additionally there is a γ -allowable π' such that $\pi'[U] = V$ and $\pi'(x) = y$. Therefore, we additionally have that $\pi_\gamma(x) = y$ as U supports x and both π_γ and π' have the same action on U . Hence, the two raised singletons must coincide.

It now simply remains to note that there are less than μ such orbits of β -supports under the allowable permutations by proposition 4.42, and that we assume by inductive hypothesis that there are less than μ coding functions of type γ , in order to conclude the desired result. $\square$

Thus, we conclude

Theorem 4.48. Each type τ_α has exactly μ elements.

Proof. Any element of a type is determined by a support (of which there are μ by remark 3.13) and a coding function (there are $< \mu$ of these by lemma 4.47), so a type has no more than μ elements (and obviously has at least μ elements). $\square$

4.5 The structure is a model of predicative TTT

There is then a very direct proof of the following:

Proposition 4.49. The structure presented is a model of predicative TTT (in which the definition of a set at a particular type may not mention any higher type).

Proof. Use E for the membership relation $\in_{TTT}$ of the structure defined above (in which the membership of type β objects in type α objects is actually a subrelation of the membership relation of the metatheory, a fact inherited from the scheme of supertypes). It should be evident that $xEy \leftrightarrow \pi_\beta(x)E\pi(y)$, where x is of type β , y is of type α , and π is an α -allowable permutation.

Suppose that we are considering the existence of $\{x : \phi^s\}$, where ϕ is a formula of the language of TST with $\in$ translated as E , and s is a strictly increasing sequence of types¹⁴. Note that the TTT type of x is some $s(j)$ and the TTT type of $\{x : \phi^s\}$ is $s(j+1)$. For any $s(j+1)$ -allowable permutation π , the truth value of each subformula of ϕ will be preserved if we replace each u of type $s(i)$ with $\pi_{A_{s,i}}(u)$, where $A_{s,i}$ is the set of all $s(k)$ for $i \leq k \leq j+1$ [there being no variables of type higher than $s(j+1)$]: $\pi_{A_{s,i}}(x)E\pi_{A_{s,i+1}}(y)$ is equivalent to $(\pi_{A_{s,i+1}})_{s(i)}(x)E\pi_{A_{s,i+1}}(y)$, which is equivalent to xEy by the observation above. The formula ϕ will contain various parameters a_i of types $s(n_i)$ and it is then evident that the set $\{x : \phi^s\}$ will be fixed by any $s(j+1)$ -allowable permutation π such that $\pi_{A_{s,n_i}}$ fixes a_i for each i . But this means that $\{x : \phi^s\}_{s(j+1)}$ is symmetric and belongs to type $s(j+1)$: we can merge the supports of the a_i 's (with suitable raising of indices) into a single $s(j+1)$ -support. Notice that we assumed the predicativity condition that no variable more than one type higher than x appears (in the sense of TST).

This procedure will certainly work if the set definition is predicative (all bound variables are of type no higher than that of x , parameters at the type of the set being defined are allowed), but it also works for some impredicative set definitions. $\square$

There are easier proofs of the consistency of predicative tangled type theory; there is a reason of course that we have pursued this one.

It should be noted that the construction given here is in a sense a Fraenkel–Mostowski construction, though we have no real need to reference the usual FM constructions in ZFA here. Constructions analogous to Fraenkel–Mostowski constructions can be carried out in TST using permutations of type 0; here we are doing something much more complicated involving many permutations of type -1 which intermesh in precisely the right way. Our explanation of our technique is self-contained, but we do acknowledge this intellectual debt.

¹⁴We use notation $s(i)$ instead of s_i to avoid undue nesting of subscripts, in what follows.

4.6 Impredicativity: verifying the axiom of union

What remains to complete the proof is that typed versions of the axiom of set union hold. That this is sufficient is a fact about predicative type theory. If we have predicative comprehension and union, we note that for any formula ϕ , $\{\iota^k(x) : \phi(x)\}$ will be predicative if k is taken to be large enough, then application of union k times to this set will give $\{x : \phi(x)\}$. $\iota(x)$ here denotes $\{x\}$. It is evidently sufficient to prove that unions of sets of singletons exist. So what we need to show is the following result.

Proposition 4.50. If $\alpha > \beta > \gamma$ and $G \subseteq \tau_\gamma$, and

$$\{\{g\}_\beta : g \in G\}_\alpha$$

is symmetric (has an α -support, so belongs to τ_α), then G_β is symmetric (has a β -support, so belongs to τ_β).

Proof. Suppose that $\{\{g\}_\beta : g \in G\}_\alpha$ is symmetric. It then has a strong support S . We claim that $S_{(\beta)}$, defined as $\{(z, C) : \max(C) = \beta \wedge (z, C \cup \{\alpha\}) \in S\}$, is a β -support for G_β , where without loss of generality we assume that S has been reordered such that the domain of $S_{(\beta)}$ (and $(S_{(\beta)})_{(\gamma)}$) is an ordinal.

Fix some $g \in G$. There is a strong γ -support T for g which extends $(S_{(\beta)})_{(\gamma)}$.

Suppose that the action of a β -allowable permutation π fixes $S_{(\beta)}$. Our plan is to use Freedom of Action technology to construct an α -allowable permutation π^* whose action on S is the identity and whose action on $T^{\uparrow\{\alpha, \beta\}}$ precisely parallels the action of π on $T^{\uparrow\beta}$.

If this is accomplished, then the action of π^* fixes S and so fixes

$$\{\{g\}_\beta : g \in G\}_\alpha,$$

while at the same time $(\pi_\beta^*)_\gamma$ agrees with π_γ at g . This implies that

$$\pi_\gamma(g) = (\pi_\beta^*)_\gamma(g) \in (\pi_\beta^*)_\gamma "G = G.$$

The same construction can be done for each g , and for π^{-1} in place of π , so π fixes G_β .

Now, we show how to construct the allowable permutation π^* by the Freedom of Action Theorem (4.25), by defining the following α -action ξ .

For every $(\{x\}, A) \in \mathbf{rng}(S)$, we define $\xi_A(x) = x$, and for every $(N, A) \in \mathbf{rng}(S)$ for N a near-litter, we define $\xi_A(N) = N$. Extend $T^{\uparrow\beta}$ to a strong β -support T^* (as in proposition 4.28), then if $(\{x\}, B) \in \mathbf{rng}(T^*)$, we define $(\xi_\beta)_B(x) = \pi_B(x)$, and if $(N, B) \in \mathbf{rng}(T^*)$ for N a near-litter, we similarly define $(\xi_\beta)_B(N) = \pi_B(N)$. These two definitions for ξ overlap for support elements that appear in both S and T^* , but whenever this occurs, the two defined images coincide because π fixes $S_{(\beta)}$.

It now suffices to check that this action is coherent; then, by the Freedom of Action Theorem, we will obtain π^* with the property that $\pi^*[S] = S$ and $\pi_\beta^*[T^*] = \pi[T^*]$, so in particular, $(\pi_\beta^*)_\gamma[T] = \pi_\gamma[T]$.

If $\xi_A(N)$ is defined, we have two cases: either (N, A) appears in S , in which case $\xi_A(N) = N$, or $B = A \setminus \{\alpha\}$ has maximum element β and (N, B) occurs in T^* with $\xi_A(N) = \pi_B(N)$.

In the first case, the case for flexibility is trivial, and if $N \sim f_{\delta, \min(A_1)}(x, U)$ with $\delta < \min(A_2)$, then as S is a strong support, the range of $U^{\uparrow A_2}$ is a subset of the range of S , as required.

We now consider the second case. If N is A -flexible, we claim that $\pi_B(N)$ is also A -flexible. Suppose that $\pi_B(N) \sim f_{\delta, \gamma}(x, U)$ for $\gamma = \min(A_1)$ and $\delta < \min(A_2)$. Then by the coherence condition, assuming B has length at least 3,

$$N \sim \pi_B^{-1} \circ f_{\delta, \gamma}(x, U) \sim f_{\delta, \gamma}((\pi_{B_2}^{-1})_\delta(x), (\pi_{B_2}^{-1})_\delta[U]),$$

so N is A -inflexible, giving a contradiction. In the case that B does not have length at least 3, it must have length exactly 2 because A has length at least 3, so we have $A = \{\alpha, \beta, -1\}$ and $B = \{\beta, -1\}$. Note that (N, B) occurs in T^* , so it either came from a support condition in T , which is impossible as $B = \{\beta, -1\}$ does not contain γ , or it came from a type-raised copy of an η -support we needed to include in T^* to make it a strong support. But $\eta \neq -1$ as -1 -supports cannot contain near-litter support conditions, and if η were a proper type index, we would have $-1 < \eta < \beta$ and $\eta \in B$, also giving a contradiction.

Finally, suppose that N is A -inflexible, so $N \sim f_{\delta, \gamma}(x, U)$ for $\gamma = \min(A_1)$ and $\delta < \min(A_2)$. By the argument of the previous paragraph, we must conclude that A is of length at least 4, and so $\gamma = \min(B_1)$ and $\delta < \min(B_2)$. As T^* is strong, it contains a suitably type-raised copy of U , and so we are done by definition of the action of ξ on T^* . $\square$

This completes the proof. In the formal proof in Lean, what is actually done is a proof that each of the assertions in the finite axiomatization of Hailperin in the version discussed in subsection 2.4 holds in all typed versions in our structure for the language of TTT, so it is in fact a model of TTT. The axioms in Hailperin other than the axiom of type lowering are predicative comprehension axioms and admit demonstration by the methods of section 4.5, done explicitly without metamathematics. In the formalization, the axiom of type lowering, which contains rather more content than the axiom of set union restricted to sets of singletons which is proved here, is proved by first proving the existence of an iterated image under elementwise application of the singleton operation of the desired set, whose definition is predicative, then repeatedly applying the result of this section that sets of singletons have unions.

5 Conclusions, extended results, and questions

In a series of subsections, we discuss further results supported by our model construction and further questions.

5.1 Ways in which this resolution of the NF consistency problem is unexciting

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

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

5.2 Questions and a result about global choice-like statements

Questions of consistency with NF of global choice-like statements such as “the universe is linearly ordered” or the prime ideal theorem cannot be resolved by the method used here (at least, not without major changes). Our models falsify the existence of a global linear order at the outset: a litter cannot be linearly ordered. We see no reason to believe that the universe cannot be linearly ordered in NF or that the prime ideal theorem cannot hold, but we do not know how to model these situations.

We do have one quite interesting theorem which appears to show that most practical applications of the axiom of choice are supported in our models of TTT and so with care in the resulting models of NF.

Note that in our models of TTT, any relation which the internal theory thinks is a well-ordering actually is a well-ordering. The reason is that in the metatheory any linear order which is not a well-ordering admits a countable set with no minimal element (the range of a strictly decreasing sequence). This requires choice, but choice holds in the metatheory. So, if a relation represented in the model of TTT is not a well-ordering from the standpoint of the metatheory, there is a countable subcollection of its domain which has no minimal element, and this countable subset is actually realized in the model of TTT, because our models of TTT are countably complete, so the model of TTT sees that the relation is not a well-ordering.

Essentially the same argument shows that any relation which a model of TTT constructed by our method thinks is well-founded is actually well-founded

in the sense of the metatheory.

Theorem: In our models of TTT, the power set of any well-orderable set can be well-ordered.

Proof: Suppose that X is a set in a model of TTT constructed by our method which admits a well-ordering (defined along a particular type sequence). The well-ordering $\leq_X$ of X which we are given must have a support. A permutation which fixes the support of $\leq_X$ will fix $\leq_X$ and, because $\leq_X$ is a well-ordering in the sense of the metatheory, its action on the type of the elements of X (in the type sequence with respect to which $\leq_X$ is a well-ordering) must fix each element of X , so in fact we have a common support for all elements of X , which in turn gives a common support for all subcollections of X in the sense of the metatheory as implemented in the appropriate type in the particular type sequence, so all of these are sets in the model of TTT (in the appropriate type in the particular type sequence) and further gives a support for the well-ordering of the power set of X induced by any well-ordering of the power set of X in the sense of the metatheory, so the power set of X is well-ordered in the model of TTT.

It is worth remarking that this theorem has been proved in the Lean formalization.

It is an interesting incidental remark that the statement that power sets of well-orderable sets are well-orderable is equivalent in NF to the assertion that a single big power set, the power set of the set of ordinals, can be well-ordered, and this must hold in models of NF constructed from our models of TTT.

This form of choice seems to support all practical applications of choice in mathematics outside of set theory, and a lot of the applications inside set theory.

Applying the same technique to sets with well-founded extensional relations on them can be used to show that all $\beth$ numbers are alephs, and that the internal representation of an initial segment of the cumulative hierarchy using equivalence classes of well-founded extensional relations with top which is possible in TTT will agree precisely with the cumulative hierarchy in the metatheory as far as it goes (the internal representation will always give a proper initial segment of the hierarchy in the metatheory, of course).

5.3 Strong unstratified axioms for NF derived from strong partition properties

NF with strong axioms such as the Axiom of Counting (introduced by Rosser in [15], an admirable textbook based on NF), the Axiom of Cantorian Sets

(introduced in [5])¹⁵ or my¹⁶ axioms of Small Ordinals and Large Ordinals (introduced in my [7], which pretends to be a set theory textbook based on NFU) should be obtainable by choosing λ large enough to have strong partition properties, more or less exactly as I report in my paper [8] on strong axioms of infinity in NFU: the results in that paper are not all mine, and I owe a good deal to Robert Solovay in that connection (unpublished conversations and [20]). I am essentially certain that models of NF satisfying these axioms can be obtained in this way, but spelling out the details in full would take more space than is appropriate in this already long paper. In any case the Axiom of Counting holds in ω -models, whose existence we have already established.

That NF has α -models for each standard ordinal α follows by the same methods Jensen used for NFU in his original paper [11]: the proof we give in subsection 2.3.1 for the existence of α -models is essentially the same as Jensen's, and a model of TTT constructed by our method will implement each ordinal α of the metatheory which is less than κ (and we can choose κ larger than any fixed α and λ with cofinality greater than $|\alpha|^{|\alpha|}$).

No model of NF can contain all countable subsets of its domain. The index reference for this fact is [16]. The easiest way to give an example involves sketching the resolution of the Burali-Forti paradox under stratified comprehension. Ordinals are defined as equivalence classes of well-orderings under isomorphism. The natural order on the ordinals is a well-ordering and has an order type Ω . We define an operation T on ordinals: $T(\alpha)$ is the order type of $R' = \{(\{x\}, \{y\}) : (x, y) \in R\}$ for $R \in \alpha$. Now the Burali-Forti paradox seems to be a danger here: consider the order type of the natural order on the ordinals restricted to ordinals less than Ω : if this is Ω we have a contradiction. What actually happens is that our belief that the order type of the natural order on ordinals $< \alpha$ is α is unfounded here. What can actually be proved is that the order type on the natural order on the ordinals less than α is $T^4(\alpha)$ [the exponent can be lowered to 2 if Quine's type level pair is used]. Then we have our example of a countable subset of the domain of our model which cannot be a set: $\Omega > T^4\Omega > T^8\Omega \dots$, so $\{T^{4n}(\Omega) : n \in \mathbb{N}\}$, a countable collection of ordinals of the model, has no smallest element and so cannot be a set. All well-typed combinatorial consequences of closure of a model of TST under taking subsets of size $< \kappa$ will hold in our models, but the application of compactness which gets us from TST + Ambiguity to NF forces the existence of externally countable proper classes, a result which has long been known and which also holds in NFU.

¹⁵Getting Cantorian Sets or Large Ordinals to hold is sensitive to the relationship between κ and λ . It requires the hypothesis that $\kappa < \lambda = \mu$ to avoid outright refutability of this axiom in resulting models of NF, and then a large cardinal hypothesis. We note that for the moment we need to use Henson's original formulation of the Axiom of Cantorian Sets, that any *well-orderable* cantorian set is strongly cantorian; the situation for non-well-orderable sets is not yet clear to us.

¹⁶The first author is speaking.

5.4 Some esoteric questions answered

We mention some esoteric problems which our approach solves. The Theory of Negative Types of Hao Wang (TST with all integers as types, proposed in [24]) has ω -models; an ω -model of NF gives an ω -model of the theory of negative types immediately. The question of existence of ω -models of the theory of negative types was open.

In ordinary set theory, the Specker tree of a cardinal is the tree in which the top is the given cardinal, the children of the top node are the preimages of the top under the map $(\kappa \mapsto 2^\kappa)$, and the part of the tree below each child is the Specker tree of the child. Forster proved (in [9]) using a result of Sierpiński that the Specker tree of a cardinal must be well-founded (a result which applies in ordinary set theory or in NF(U), with some finesse in the definition of the exponential map in NF(U)). Given Choice, there is a finite bound on the lengths of the branches in any given Specker tree. Of course by the Sierpiński result a Specker tree can be assigned an ordinal rank. The question which was open was whether existence of a Specker tree of infinite rank is consistent. It is known that in NF with the Axiom of Counting the Specker tree of the cardinality of the universe is of infinite rank (this is also in [9]). Our methods show the existence of a model of NF with the Axiom of Counting. This gives us a model of TST in which the Specker tree of some cardinal is of infinite rank. It is straightforward to produce a model of bounded Zermelo set theory either with atoms or with extensionality and failure of foundation from a model of TST, which will satisfy the same stratified sentences: identify the elements of type 0 either as atoms or with their own singletons, then implement elements of each type $i + 1$ as subsets of type i as guided by the original model of TST, identifying type $i + 1$ objects with elements of lower types which have been assigned the same extension. This shows that we can get a model of bounded Zermelo set theory with atoms or without foundation in which there is a cardinal with Specker tree of infinite rank. It is straightforward with our methods to get a model of TTT with the Axiom of Counting with inaccessibles cofinal in λ : take $\kappa > \lambda$ with inaccessibles cofinal in λ and build an ω -model as described above. In a model of NF derived from an ω -model the Axiom of Counting will hold. In this model we will have an inaccessible greater than a cardinal with Specker tree of infinite rank, and so a model of ZFA in which there is a cardinal with Specker tree of infinite rank.

The further question remains as to whether it is consistent with ZF that there is a cardinal with Specker tree of infinite rank. We would guess that this could be shown to be true using symmetric forcing extensions, but we are not expert at this and leave the question to others.

5.5 The question of consistency strength

We believe that NF is no stronger than TST + Infinity, which is of the same strength as Zermelo set theory with separation restricted to bounded formulas ([10]). Our work here does not show this, as we need enough Replacement for existence of $\beth_{\omega_1}$ at least. We leave it as an interesting further task, possibly

for others, to tighten things up and show the minimal strength that we expect holds.

5.6 The wide open question: what do models of NF look like in general?

Another question of a very general and amorphous nature which remains is: what do models of NF (or TTT) look like in general? Are all models of NF in some way like the ones we describe, or are there models of quite a different character? There are very special assumptions which we made by fiat in building our model of TTT which do not seem at all inevitable in general models of this theory.

5.7 Postscript

Inevitably, philosophical issues come up in connection with a system of set theory proposed by a philosopher. I get a lot of congratulations for vindicating Quine's foundational agenda, but in fact this is not part of my purpose here. It is not even clear to me that Quine *had* a foundational agenda in which his technical proposal of this set theory had a special place.

The results of this paper show that if one really wanted to, one could use NF as a foundation for mathematics. It is odd that it disproves Choice, but the principle "power sets of well-orderable sets are well-orderable", which holds in models constructed in this way, supports most applications of Choice.

Our opinion is that Quine's proposal of NF was based on a mistake. He discusses whether to assume strong extensionality in the original paper, and his explanation of reasons for choosing strong extensionality contains an actual mathematical error: he maintains that if urelements were present he could re-define membership so as to treat each urelement as its own sole element; this does work in ordinary set theory (though not quite in the way he describes) but in stratified set theory it is not a legitimate move, as the proposed redefinition of membership is unstratified and one cannot rely on comprehension holding for the reinterpreted membership relation. Also, a collection of self-singletons has cardinality less than that of $\iota^n V$ (the collection of n -fold singletons) for any n and so is quite small. None of this is conjectural: the proposal that Quine makes for eliminating urelements fails in all models of NFU built using Jensen's construction. We believe that the correct system to propose was NFU, and if he had proposed NFU, the history of this kind of set theory might have been different.

NFU is serviceable as a foundation for mathematics, and consistent with Choice and various strong axioms of infinity. It is odd that NFU + Choice proves that there are urelements (the same odd fact that NF disproves Choice), but no more than odd. The type discipline of stratification is something that one must work to get used to, and it has been remarked that working with indexed families of sets is extremely awkward in NFU (and would be similarly awkward in NF). The view of the world which NFU supports is basically the same as that of

ZFC: the natural models of NFU are obtained by considering an initial segment of the cumulative hierarchy with an external automorphism which moves a level, and NFU can interpret discussion of exactly such a structure internally.

Unlike NF, NFU on introspection can tell one quite a lot about what a model of NFU might look like (as ZFC can with its own awareness of initial segments of the cumulative hierarchy; in NFU, analysis of the isomorphism classes of well-founded extensional relations gives both an interpretation of an initial segment of the Zermelo style universe [with an external endomorphism in the form of the operation passing from the relation type of R to the relation type of R^u] and an interpretation of NFU itself if one has strong enough assumptions). NF does not appear to give similar internal hints as to how a model of NF might be constructed (it can, being an extension of NFU, internally construct a lot of information about an interpretation of NFU with lots of urelements, but there does not seem to be an obvious way to find out how an extensional world is constructed from internal evidence in NF itself).

So, we do not believe this paper is a philosophical milestone. If there was one, it happened in 1969 when Jensen showed that NFU is consistent, and nobody noticed.

We do believe that there are interesting questions to investigate about NF. Our paper does not settle all the questions about this system which have developed in the minds of people who have worked with it since 1937. In fact, the construction is based on special assumptions and does not seem to give much of an idea of what general models of this theory might look like.

It might be viewed as philosophically interesting that this proof was formally verified. That is really an advertisement for a quite different foundational system, the logic of the Lean proof verification system. That it needed to be formally verified I believe reflects an interesting complexity in the mathematics here, not only the deficiencies of the first author as an expositor.

References

- [1] Crabbé, M. [1992a] On NFU. *Notre Dame Journal of Formal Logic* 33, pp 112-119.
- [2] Scott Fenton. New Foundations set theory developed in metamath. 2015. <https://us.metamath.org/nfeuni/mmnf.html>
- [3] Forster, T.E. [1995] Set Theory with a Universal Set, exploring an untyped Universe Second edition. Oxford Logic Guides, Oxford University Press, Clarendon Press, Oxford.
- [4] Hailperin, T. [1944] A set of axioms for logic. *Journal of Symbolic Logic* 9, pp. 1-19.
- [5] Henson, C.W. [1973a] Type-raising operations in NF. *Journal of Symbolic Logic* 38 , pp. 59-68.
- [6] Holmes, M.R. “The equivalence of NF-style set theories with ”tangled” type theories; the construction of omega-models of predicative NF (and more)”. *Journal of Symbolic Logic* 60 (1995), pp. 178-189.
- [7] Holmes, M. R. [1998] Elementary set theory with a universal set. volume 10 of the Cahiers du Centre de logique, Academia, Louvain-la-Neuve (Belgium), 241 pages, ISBN 2-87209-488-1. See <https://randall-holmes.github.io/errata.txt> for an on-line errata slip. By permission of the publishers, a corrected text is published on-line at <https://randall-holmes.github.io/head.pdf>; an official second edition will appear online eventually.
- [8] Holmes, M. R. [2001] Strong Axioms of infinity in NFU. *Journal of Symbolic Logic*, 66, no. 1, pp. 87-116.
(“Errata in ‘Strong Axioms of Infinity in NFU’ ”, *JSL*, vol. 66, no. 4 (December 2001), p. 1974, reports some errata and provides corrections).
- [9] Forster, Thomas E. [1976] ”N.F.” Ph.D. thesis, University of Cambridge.
- [10] Kemeny, J.G. [1950] Type theory vs. set theory (abstract of Ph.D. thesis). *Journal of Symbolic Logic* 15, p. 78.
- [11] Jensen, R.B. “On the consistency of a slight(?) modification of Quine’s NF”. *Synthese* 19 (1969), pp. 250-263.
- [12] Quine, W.V. [1945] On ordered pairs. *Journal of Symbolic Logic* 10, pp. 95-96.
- [13] Quine, W.V., “New Foundations for Mathematical Logic”. *American Mathematical Monthly* 44 (1937), pp. 70-80.
- [14] Quine, W.V. *Mathematical Logic*. Norton, 1940. revised ed. Harvard, 1951.

- [15] Rosser, J. B. [1978] Logic for mathematicians, second edition. Chelsea Publishing.
 - [16] Rosser, J.B. and Wang, H. [1950] Non-standard models for formal logic. *Journal of Symbolic Logic* 15, pp. 113-129.
 - [17] Russell, Bertrand. *Principles of Mathematics*. Routledge Classics 2010 (originally published 1903).
 - [18] Russell, B.A.W. and Whitehead, A. N.[1910] *Principia Mathematica*. Cambridge University Press.
 - [19] Scott, Dana, “Definitions by abstraction in axiomatic set theory”, *Bull. Amer. Math. Soc.*, vol. 61, p. 442, 1955.
 - [20] Solovay, R, “The consistency strength of NFUB”, preprint on [arXiv.org](https://arxiv.org/abs/math/9707207), [arXiv:math/9707207](https://arxiv.org/abs/math/9707207) [math.LO]
 - [21] Specker, E.P. “The axiom of choice in Quine’s new foundations for mathematical logic”. *Proceedings of the National Academy of Sciences of the USA* 39 (1953), pp. 972-975.
 - [22] Specker, E.P. [1962] “Typical ambiguity”. *Logic, methodology and philosophy of science*, ed. E. Nagel, Stanford University Press, pp. 116-123.
 - [23] Tarski, Alfred. Einige Betrachtungen uber die Begriffe der co-Widerspruchsfreiheit und der co-Vollständigkeit, *Monatsh. Math. Phys.* 40 (1933), 97-112.
- I believe this is the right paper. There is a treatment due to Gödel at about the same time. The reason that I think Tarski is the originator of TST is that Gödel’s theory is ω -order arithmetic (the lowest type is further equipped with the Peano axioms) which is not the same theory, and crucially does not have the ambiguity properties which motivate NF.
- [24] Wang, H. [1952] Negative types. *MIND* 61, pp. 366-368.
 - [25] Wiener, N. ‘A Simplification of the Logic of Relations,’ *Proceedings of the Cambridge Philosophical Society*, 17 (1914) 387-90
 - [26] Sky Wilshaw, Yaël Dillies, Peter LeFanu Lumsdaine, et al. New Foundations is consistent. GitHub repository. <https://leanprover-community.github.io/con-nf/>.

Initial work was done by a group, as indicated by the authorship; since the beginning of 2023, the work, which now amounts to the majority of what has been accomplished, has been done by Sky Wilshaw alone, and the previous work has been very largely reorganized by her.