Axiomatising set theory with a universal set

Thomas Forster

June 27, 1997

I have converted this document from the VUWRITER format in which it originally appeared: partly by hand, party by use of a program written by Paul Taylor It might be of some historical interest.

This paper has grown out of a talk I took round various hospitable logic seminars in England in the winter term of 1981. I would like to thank members of the Bedford, Cambridge and Leeds seminars for illuminating discussions which have helped this document into its present form. It appeared in [3].

There are various problems in Set theory with a universal set: chief of them is persuading people that there is anything worth studying. I shall say less on this subject than I would like, for you who have already read this far are presumably at least willing to listen. However I will allow myself one gibe: ZF is obviously the core of any sensible axiomatic set theory of well founded sets, but that is not to say that wellfounded sets are all there are. Much of the plausibility of ZF as an axiomatisation for *all* of set theory arises from mistaking arguments for the first for arguments for the second.

Although it is now over 40 years since the first axiomatic set theory with a universal set was published, there is still no agreement on even a core for an axiomatisation of set theory with $V \in V$. In this paper I present some (i hope) persuasive motivation for some axioms. The programme is best begun by looking at the most basic problem of all, namely

The problem of identity in set theory with a universal set

The problem of identity in set theory with a universal set is the same as in the more general case of illfounded set theory. Indeed I shall not make much of the difference since there seems little motivation for illfounded set theory unless one is interested in set theory with a universal set.

The axiom of extensionality summarises all that conventional wisdom has had to say about "=" in set theory. It is the closest we come to saying in any formal sense that sets are that-which-is-extensional. A set is just the collection of its members, that and no more. (The most tough-minded expression of this point of view, which I shall adopt below, is that *the only thing a set theorist can know about = is that it is a congruence relation w.r.t.* $\in$.) Thus extensionality, in conjunction with the axiom of foundation, enables us to decide when $x = y$ by seeing if their members are identical. The regress we are launched on here must terminate because the ranks of the things we are looking at is reduced by the induction step. I am becoming more and more convinced that the appeal of the axiom of foundation is simply that it provides us with this elegant recursive characterisation of identity and thereby spares us the need to give the matter any further thought. Historically this restriction has been fruitful as it has enabled us to concentrate our efforts on those parts of set theory where results could be obtained quickly and applied widely. Of late the profuse growth of parts of wellfounded set theory of no interest to outsiders has begun to suggest that it has had all the help it needs and that the time is now ripe to reopen the fundamental questions we have ignored since the turn of the century.

The problem then is that when $V \in V$, the regress I spoke of in the last paragraph ("$x = y$? Are all their members identical? Are all *their* members identical? ...") cannot be relied upon to terminate. Here we can profitably introduce some ideas from game theory. Notation and terminology here will be standard except for the use of the word 'Wins' with upper-case 'W' to mean "...has a winning strategy for..." and that a strategy is not a thing that says "when here, do this" but only "when here, do one of these", namely a thing variously known as a nondeterministic strategy or a restraint etc. This is because, as have argued elsewhere ([1]), AC is probably false in any sensible set theory with a universal set, so if we use strategies in the standard sense (which can often be little more then thinly disguised choice functions) we are liable to find that a game has no Winning strategy in the standard sense—for reasons that have nothing to do with the game itself but derive from our uneccessarily strong notion of strategy and the large and uncertain universe in which the game dwells.

The first game here will be notated $G_{x=y}$ ("The identity game") to commemorate the fact that it is being played to decide whether or not x=y.

Player `II` moves first, choosing a subset $R_1 \subseteq (x \times y)$ s.t. R_1"$V = x$ and R_1^{-1}"$V = y$. At each stage Player `I` picks an ordered pair $\langle x_n, y_n \rangle$ from `II`'s previous choice. Subsequently Player `II` chooses R_{n+1} a subset of $x_n \times y_n$ s.t. R_{n+1}"$V = x_n$ and R_{n+1}^{-1}"$V = y$. Player `II` loses if she is confronted with a pair $\langle x_n, y_n \rangle$ one of which is empty and the other not. `I` loses if he picks $\langle x_n, y_n \rangle$ both of which are empty (notice that this allows for the existence of urelemente). If the game goes on for ever `II` wins. The idea is that `II` is trying to prove $x = y$ and that `I` is trying to prove $x \neq y$. In earlier versions of this paper I had `II` pick bijections rather than relations, the rationale being that a set cannot have two identical members. The present version is better all the same though, because it does not compel us to decide, in order to know what the rules are to be, whether or not some things u, v in the transitive closure of x are identical (which we could discover only by playing $G_{u=v}$). Another way of putting this is to say that to specify formally the rules governing II's moves in the version of the game where she has to play bijections would use the "=" symbol whose meaning is explained only by the game which is yet to be played.

Player II's choice of R_{n+1} when faced with x_n and y_n is obtained by partitioning x_n, y_n into equivalence classes under identity and then pairing the equivalence classes for x_n with those for y_n in an appropriate way. On the face of it, this suggests that we should always require II's choice R to satisfy a condition

$$(uRv \wedge u'Rv \wedge u'Rv') \rightarrow uRv'$$

but it is not hard to persuade oneself that the resulting games are equivalent, and the proof is omitted.

$G_{x=y}$ is a open game. That is to say, if player I wins at all, he has done so after finitely many moves. So I or II must have a winning strategy. It is not hard to see that the relation

II Wins $G_{x=y}$

is a congruence relation w.r.t $\in$. Indeed it looks a very good candidate for a definiens of '$x = y$'. However there are good reasons for looking for something even stronger. Let us define j an operator on maps so that $(j`f)`x = f``x$, and let us define, for each $n \in \mathbb{N}$, an equivalence relation $\sim_n$ by

$x \sim_n y$ iff $(\exists\tau)(\tau$ is a permutation of $V \wedge (j^n`\tau)`x = y)$

$x \sim_\infty y$ iff $x \sim_n y$ for all n.

...and we invoke the notations $[x]_n$, $[x]_\infty$ for equivalence relations as usual. The importance of n-congruence derives from the fact that if $x \sim_n y$ then x and y satisfy the same stratified formulæ in which they both appear at type n. This derives from a theorem, important in the folklore of *NF*, that

LEMMA **1**

$$\phi(x, y, z \ldots) \longleftrightarrow \phi((j^m`\tau)`x, (j^n`\tau)`y, (j^k`\tau)`z, \ldots)$$

where m, n, k ... *are the types of* x, y, z ... *in* ϕ.

This fact, which will not be proved here, will be used later on. $x \sim_n y$ says that the top n "layers" of x and y look the same. This being the case, extensionality would lead us to be very suspicious of having x and y with $x \neq y$ but $x \sim_\infty y$. Since there is no obvious way of constructing a Winning strategy for II in $G_{x=y}$ given merely that $x \sim_\infty y$ a tougher definition of identity will be required.

Consider again the game $G_{x=y}$. Let us suppose I has a winning strategy. Let us consider the tree of all plays obtained by I using his Winning strategy and II doing anything legal. This tree is wellfounded since all plays (branches) terminate after a finite number of steps (the use of DC here may or may not be significant: see the next game below where a similar problem occurs) and occordingly has a rank. Let us nbotate this ordinal "$\in_{x,y}$". $\in_{x,y}$ looks rather like a truth-value of $x = y$ but the ide of ordinals as truth-values of anything is profoundly repugnant and suggests that we have too much structure, some of it spurious. Fortunately we have the following crucial fact:

If $\in_{x,y}$ and $\in_{y,z}$ are both infinite, so is $\in_{x,z}$

Proof:

"$\in_{x,y}$ is infinite" simply says that for each natural number n, player `II` has a strategy that enables her to postpone defeat until n moves have been played. If `II` has such strategies for $G_{x=y}$ and for $G_{y=z}$ then she has strategies for $G_{x=z}$ by composition. ■

What this means is that the relation "$\in_{x,y}$ is infinite" is an equivalence relation and hence must be the identity (!) That is to say

Axiom of strong Extensionality

$(\forall x)(\forall y)(x = y \longleftrightarrow (\forall n \in \mathbb{N})($`II` has a strategy to postpone defeat in $G_{x=y}$ for n moves))

This axiom deliberately expunges a lot of structure. If we had instead defined $x = y$ to be "`II` has a Winning strategy in $G_{x=y}$" then we would have lots of exciting equivalence relations to play with, since ω cannot be the only ordinal α such that $\in_{x,y} \geq \alpha$ is an equivalence relation, but we would nothave ensured that $\sim_\infty$ is equality.

A Quine atom is an object $x = \{x\}$. Strong extensionality prevents there being more than one Quine atom. Indeed it prevents there being more than one object whose transitive closure does not contain the empty set. It also excludes the possibility of $\in$-automorphisms of V.

$G_{x=y}$ has generalisations which can be useful when defining identity in a non-recursive way in models that we obtain by deleting objects e.g., urelemente, from some given model. First we identify objects whose symmetric difference consists entirely of things to be deleted. Then we delete all but one from each equivalence class, and iterate. The same effect can be achieved by playing a version of $G_{x=y}$ where the domains and ranges of the relations played by `II` avoid objects which are to be deleted.

The next game we consider has a simpler structure. This game, played with an initial set x, is notated G_x and in it each player (`I` starting) picks a member of the other player's last choice until the game is ended by one trying to pick a member of an empty set (the game can be played in universes with urelemente) and thereby losing. If the game goes on for ever it is a draw. When $x \in x$ `I` and `II` can go on picking x for ever and thus draw, but one has the feeling that empty sets ought to be sufficiently dense in the transitive closure of any x for one player or the other to be able to force a win. Let us adopt the definitions

$\texttt{I} =_{\rm df} \{x : \texttt{I} \text{ Wins } G_x\}$; $\texttt{II} =_{\rm df} \{x : \texttt{II} \text{ Wins } G_x\}$

(we will leave open for the moment whether I and II are to be sets or proper classes) Obviously $x \in \texttt{I}$ iff $(\exists y \in x)(y \in \texttt{II})$ and dually $x \in \texttt{II}$ iff $(\forall y \in x)(y \in \texttt{I})$. We can rewrite this as $\texttt{I} = B``\texttt{II}$ ($B`x =_{\rm df} \{y : x \in y\}$) and $\texttt{II} = \mathcal{P}`\texttt{I}$. If, with a view to readability, we invent a new function letter $b$ so that $b`x = \bigcup B``x$ (the 'b' is an upside-down 'p' to remind us that b corresponds to $\exists x \in \ldots$ and p to $\forall x \in \ldots$) we can rewrite this as $\texttt{I} = b`\texttt{II}$ and $\texttt{II} = \mathcal{P}`\texttt{I}$. This is rather reminiscent of the fact that $x \in WF$ iff $(\forall y \in x)(y \in WF)$ and $x \in -WF$ iff $(\exists y \in x)(y \in -WF)$. Apart from the elegant characterisation this enables us to give in a language where formulæ can have themselves as proper subformulæ, it invites us to consider what happens if we stick in yet more quantifiers, for example

$x \in X$ iff $(\exists y \in x)(\forall w \in y)(\exists u \in w)(u \in Y)$

$x \in Y$ iff $(\forall y \in x)(\exists w \in y)(\forall u \in w)(u \in Y)$

or, for short, $X = bpbY$ and $Y = pbpX$. In this case I and II are no longer unique solutions since $b`Y$ for $X$ and $\mathcal{P}`X$ for Y will satisfy the same identity. This is rather reminiscent of the way e^x splits into $sinh(x)$ and $cosh(x)$ when we require not $f = Df$ but merely $f = D^2f$. Both in that case and here we find that by increasing the number of iterations new roots will appear. This parallel will not be taken further here. Once we notice that $\Lambda \in \texttt{II}$ and $V \in \texttt{I}$ the discussion above suggests the following recursive construction:

$$\texttt{I}_0 = \{V\}; \quad \texttt{I}_{\alpha+1} = \bigcup B``\texttt{II}_\alpha$$

$$\texttt{II}_0 = \{\Lambda\}; \quad \texttt{II}_{\alpha+1} = \mathcal{P}`I_\alpha$$

…taking sumsets at limit ordinals. It is not hard to show by induction that $\texttt{I}_\alpha$ and $\texttt{II}_\alpha$ are increasing sequences under inclusion. Let us associate with each object in I or II its *rank*, the least α such that it belongs to $\texttt{I}_\alpha$ or $\texttt{II}_\alpha$. We need to show that everything in I or II does indeed have a rank. The proof is analogous to that in ZF that every wellfounded set has a rank.

Suppose $x \in \texttt{II}$ is unranked. Then every $y \in x$ is in I but then some $y \in x$ is unranked, otherwise the rank of x is just $\sup_{y \in x} \text{rank}`y + 1$. Similarly if $x \in \texttt{I}$ is unranked. But this illfoundedness enables (in either case) the "losing" player to construct a strategy ("play unranked sets") which results in an infinite play and a draw, contradicting the existence of a winning strategy. This justifies the definition if I and II as the union of their partial sums over the ordinals.

REMARK 2 *$\texttt{I}_\alpha$ and $\texttt{II}_\beta$ are disjoint for all α, β*

Proof: Suppose α and β are minimal counterexamples, then we have $x \in \texttt{II}_\alpha$ and $x \in \texttt{I}_\beta$. So there is $y \in x$ such that $y \in \texttt{II}_\delta$ for some $\delta < \beta$. But any such y (since $y \in x \in \texttt{II}_\alpha$) must also be in $\texttt{I}_\gamma$ for some $\gamma < \alpha$ contradicting minimality of α and β. This enables us to construct a canonical strategy for the winning player.

The minimal strategy

"When confronted with x, play anything in $x \cap$ II of minimal rank"

It is well-known that the rank of a well-founded set can be defined either as the rank of $\in |TC`x$ considered as a wellfounded relation or as the least ordinal α such that $x \in V_{\alpha+1}$. There is a corresponding result here: let the pseudorank of x be the least ordinal α such that $x \in \mathtt{I}_{\alpha+1} \cup \mathtt{II}_{\alpha+1}$. The pseudorank of x is the same as the rank of the tree of plays obtainable in G_x by the winning player using her minimal strategy and the other player doing anything at all. The proof is an easy induction on rank and is left to the reader. The reader may also wish to verify that any wellfounded set of rank α will have pseudorank α too. The proofs all have such an engaging familiarity to them that it suggests one should adopt, as an analogue of the axiom of foundation the following

Axiom of $\in$-determinacy

$V = \mathtt{I} \cup \mathtt{II}$

There is a slight blemish to the parallel between the axioms of $\in$-determinacy and foundation, namely that $\in$-determinacy tells us that we can associate with each set x a canonical tree which is wellfounded in the weak sense that every path through it is finite. This involves DC in subsequent proofs. We could frame $\in$-determinacy in a way that gets round this by defining recursively, on the tree of possible plays in G_x, a two-valued function f such that $f`y = 0$ says "II has a Win from stage $y$" and $f`y = 1$ says "I has a Win from stage y". The new version of $\in$-determinacy would then say that for all x, this function is defined on the whole of the tree of plays of G_x.

To lend plausibility to this axiom, we can prove it for a large natural class of sets. (For the definition of n-symmetric, see below)

PROPOSITION **3** *Let X be n-symmetric, with n even (odd). Then either* I (II) *Wins G_X in $n+2$ moves or* II (I) *Wins G_X in $n+3$ moves.*

Let us take the case $n = 6$ as a typical illustration. Let '$\Phi(y, x)$' be short for

$$(\exists x_5 \in x)(\forall x_4 \in x_5)(\exists x_3 \in x_4)(\forall x_2 \in x_3)(\exists x_1 \in x_2)(x_1 \subseteq y).$$

Since II Wins G_x for any $x \subseteq B`\Lambda$, $\Phi(B`\Lambda, X)$ will certainly imply that I Wins G_X (in eight moves in fact). '$\Phi(y, x)$' is a stratified wff in which 'x' is of type 6 and 'y' of type 1. By an application of lemma 1 (the automorphism lemma for definable sets) we have

$$\Phi(B`\Lambda, X) \longleftrightarrow \Phi((j`\pi)`(B`\Lambda), (j^6`\pi)`X)$$

for any permutation π. But X is by hypothesis 6-symmetric, which is to say $X = (j^6`\pi)`X$ for any π, so this becomes

$$\Phi(B`\Lambda, X) \longleftrightarrow \Phi((j`\pi)`(B`\Lambda), X).$$

Now suppose I does not have a strategy to Win in eight moves. Then $\neg\Phi(B`\Lambda, X)$ and indeed $\neg\Phi((j`\pi)`B`\Lambda, X)$ for any permutation π.

We now seek a permutation π so that $(j`\pi)`(B`\Lambda) = -\mathcal{P}`B`\Lambda$. This is easy because $(B`\Lambda)$ and $-\mathcal{P}`B`\Lambda$ are the same size, as are their complements. So $\neg\Phi(-\mathcal{P}`B`\Lambda, X)$ which is

$$(\forall x_5 \in x)(\exists x_4 \in x_5)(\forall x_3 \in x_4)(\exists x_2 \in x_3)(\forall x_1 \in x_2)\neg(x_1 \subseteq -\mathcal{P}`B`\Lambda)$$

and the matrix simplifies to $(\exists x_0 \in x_1)(\forall x_{-1} \in x_0)(\Lambda \in x_{-1})$ which is to say II Wins in nine moves. The proofs for other finite n are similar ■

$\in$-determinacy can thus have no counterexamples which are sets definable by straftified expressions.

$\in$-determinacy gets rid of Quine atoms for us (only one play possible in G_x if $x = \{x\}$ and that is a draw!), but there are equally pathological objects that it does not get rid of, such as $x = \{x, \Lambda\}$. Such an object clearly belongs to I so it does not contradict $\in$-determinacy. Strong extensionality limits the number of such objects to 1 but does not get rid of them altogether. We shall find such an axiom in the next section where the discussion has been broadened a bit.

I am going to introduce some canonical objects, canonical in the sense that they are distinguished representatives of their kind generated in a very natural way by the theory. In what sense any of them are to be sets will be left open.

1 The canonical topology

The pseudorank function given by $\in$-determinacy may eventually give us again some constructive control over V but until we have that sort of wellfoundedness available again it is more natural to look instead from the top downwards and classify sets according to what their top few layers look like. For this we naturally turn to the n-equivalence classes of [1]. We topologise V by taking as basis all sets of the form $[x]_n$. All neighbourhoods will in fact be clopen. If we use Quine ordered pairs (so that $V = V \times V$) we find that the product topology on V^2 is identical to the topology on V. The fact mentioned earlier, namely that

$\phi(x, y, z \ldots) \longleftrightarrow \phi((j^n`\pi)`x, (j^m`\pi)`y, (j^k`\pi)`z \ldots)$ where $n, m, k \ldots$ are the types of $x, y, z \ldots$ in ϕ can accordingly be minuted as

REMARK 4 *Functions defined by stratified formulae are continuous in the canonical topology.*

The bad news this brings is the following:

REMARK 5 *The canonical topology is not compact*

In the presence of AC_2 we can find a permutation π of V so that π and $j`\pi$ are conjugate, with $\pi \neq$ the identity (see [5]). This amounts to saying $[\pi]_k = [j`\pi]_k$ for some fixed small k. Also we can show, for any n, that

$[\delta]_n = [j`\delta]_n \rightarrow [j`\delta]_{n+1} = [j^2`\delta]_{n+1}$ for any permutation δ.

From this it follows that $[\pi]_k$, $[j`\pi]_{k+1} \ldots [j^n`\pi]_{k+n} \ldots$ is a nested sequence of closed sets, whose intersection must be nonempty (by compactness) and a singleton $\{a\}$ (by the axiom of strong extensionality) It then follows that $a = j`a$, which is to say that a is an automorphism. Any two objects that are interchanged by an automorphism must be ∞-equivalent, so strong extensionality will imply that there are no automorphisms, contradicting compactness.

Definition 6 *A set is symmetric if it is isolated in the canonical topology.* *x is n-symmetric if $[x]_n = \{x\}$.*

That is to say, x is symmetric if it is n-symmetric for some n. The terminology "symmetric" is motivated by the fact that an n-symmetric set is fixed by lots of permutations of V, namely all those that are j^n of something. All sets definable by stratified expressions will be symmetric. This suggests that the family of symmetric sets might be an appropriate model for some set of axioms we might wish to develop. This possibility is discussed in [1] where it is shown in *NF* that if SYMM (the family of symmetric sets) is extensional (i.e., if x, y are distinct members of SYMM then $x \Delta y$ meets SYMM) then it is a submodel of V elementary for stratified wffs and that AC_2 must fail. One could motivate an axiom V=SYMM along the following lines: the second axiom of strong extensionality implies that $[x]_\infty = \{x\}$ for all x and $V = SYMM$ says that for each x there is some n such that $[x]_n = \{x\}$ already. So "$V = SYMM$" is a natural strengthening of strong extensionality. However its consequences are too bizarre for that to be a sufficient reason for adopting it.

A *permutation model* obtained from V and a permutation π in it (notated V^π) is the structure obtained by keeping the same elements but rewriting $\in$ so that $x \in y$ (in the new sense) iff $x \in \pi`y$ (in the old sense). Such models have been of great help in the devising of relative independence and consistency results in *NF* since the transition to a permutation model preserves all stratified sentences true in the original model, and all the axioms of *NF* are stratified. To procede further we shall need some notation. Let γ be an arbitrary permutation.

Definition 7 $\gamma_0 = identity$; $\gamma_{n+1} = (j`\gamma_n) \circ \gamma$

Now we can express the following piece of folklore

$$V^\gamma \models \phi(x, y, z, \ldots) \longleftrightarrow V \models \phi(\gamma_n`x, \gamma_k`y, \gamma_m`z \ldots)$$

where n, k, $m \ldots$ are the types of 'x', 'y', 'z' ....in ϕ. In particular, $V^\gamma \models x \sim_n y$ iff $\gamma_n`x \sim_n \gamma_n`y$. We shall now try to identify n-equivalence classes across permutation models. We shall need a analogue of the j operation for maps $\pi : V^\sigma \longleftrightarrow V^\gamma$. Call it $m`\pi$ (a nonce notation). We have

$(m`\pi)`x = \pi``x$ (in the sense of V^γ)

so

$m`\pi = \gamma^{-1} \circ (j`\pi) \circ \sigma$

so that to say that x (in V^σ) is n-equivalent to y (in V^γ) becomes

$$(\exists\pi)(y = ((\gamma_n)^{-1} \circ (j^n`\pi) \circ \sigma_n)`x)$$

If we set π = identity we see that $(\gamma_n^{-1} \circ \sigma_n)`x$ is an object in V^γ which has the same n-equivalence class in V^γ as x does in V^σ. In other words, V^σ and V^γ have the same n-equivalence classes. Another way of putting this is to say that the canonical topologies in V^σ and V^γ have the same lattice of open sets and that the only difference is in which sequence of closed sets have empty intersection. By judicious choice of τ we can arrange to V^τ to have, or not to have, a Quine atom. Assuming the second axiom of strong extensionality the (non)-existence of Quine atoms is equivalent to the following nested sequence of closed sets:

$$\iota``V, \iota^2``V, \ldots \iota^n``V \ldots$$

having empty (nonempty) intersection. (ι is the singleton function.) This motivates a partial order of permutations where $\sigma \leq \tau$ if more intersections of closed sets are empty in V^σ than in V^τ. Define

$$\sigma \leq \tau \longleftrightarrow (\exists f)(\forall x)(\forall n)(\exists\pi)(f`x = (\tau_n)^{-1} \circ (j^n`\pi) \circ \sigma_n`x))$$

Thus σ precedes τ iff we can find a function f which sends each $x \in V^\sigma$ to something $f`x \in V^\tau$ which is $n$-equivalent to it for each $n$. It is mechanical to verify that $\leq$ is transitive (take compositions). It is not actually antisymmetrical because $\sigma \leq j`\sigma \leq \sigma$. (σ itself is an isomorphism between V^σ and $V^{j`\sigma}$) $\leq$ has an automorphism generated by -, the complementation function. - commutes with everything in $J_1$ so $j^n`-$ commutes with everything in J_{n+1}. We can use this fact to verify that $\sigma \leq \tau$ iff $- \circ \sigma \leq - \circ \tau$.

Permutations can be used to give us models free of rubbish like Quine atoms. One might feel that any creature that can be thus eradicated is probably something we are better off without. This motivates the **Axiom of minimality** $= \leq \tau$ for all permutations τ.

Minimality is the promised axiom for getting rid of things like $x = \{x, \Lambda\}$. The sweep made by minimality may be a lot cleaner even than that, since no-one has yet proved that if V contains an infinite von Neumann ordinal then so must all its permutation models. If this is not true, and infinite Von Neumann ordinals can indeed be got rid of, then Minimality will have the consequence that there are none, and that all von Neumann ordinals are finite. The reader may feel that the absence of infinite Von Neumann ordinals is unfortunate, but to do arithmetic one does not need Von Neumann ordinals any more than one needs fingers. There are various ways out of this: one could adopt the hard-nosed point of view expressed above, that von Neumann ordinals are essentially irrelevant to ordinal arithmetic even, and are no more mathematically necessary than fingers. There is also the possibility that further research will reveal that infinite von Neumann ordinals can not, in fact, be got rid of, or that more research will tell us so much more about $\leq$ that we realise that minimality is quite inappropriately strong for quite other reasons.

If we consider the special case of minimality that asserts that identity $\leq$ complementation and bear in mind that $\sigma \leq \tau$ iff $- \circ \sigma \leq - \circ \tau$ we infer that V and V^- must have the same canonical topology. This particular case has other motivations. Let $\hat{\phi}$ (read "ϕ-dual") be the result of replacing all occurrences of $\in$ in ϕ by $\notin$ and vice versa. It is evident that $\hat{\phi}$ is logically valid iff ϕ is too. Also that $\hat{}$ is an involution which respects interdeducibility. There will be a corresponding notion of the dual $\hat{\mathcal{M}}$ of a structure $\mathcal{M}$. A structure isomorphic to its dual will be said to be self-dual. In [4] Specker considers the behaviour of such dualities and lists three possibilities for theories T whose languages admit such a duality

1. for all ϕ, $T \vdash \phi$ iff $T \vdash \hat{\phi}$
2. for all ϕ, $T \vdash \phi \longleftrightarrow \hat{\phi}$
3. All models of T are self-dual (to which we may as well add, since large objects can be sets here...)
4. All models $\mathcal{M}$ of T contain an isomorphism $\mathcal{M} \simeq \hat{\mathcal{M}}$

Maps as in (iii) or (iv) are *antimorphisms*. An antimorphism as in (iii) is an *external* antimorphism, one as in (iv) is an *internal* antimorphism. Such antimorphisms are discussed in [1] where it is proved that the existence of an internal antimorphism is inconsistent with AC_2. That proof does not go through if the antimorphism is external (not a set of the model).

REMARK **8** *Any antimorphism of V is unique*

Proof: The composition of any two antimorphism is an automorphism and therefore the identity. Pending further progress on the minimality front we can at least adopt the following special case of it: **Axiom of duality**: There is a unique antimorphism.

Duality and strong extensionality together get rid of another class of pathological object, the Boffa atom. x is a Boffa atom if $x = \{y : x \in y\}$. The reader may verify that if x is a Boffa atom and π is an antimorphism then $\pi`x$ is also a Boffa atom. Also that $x \in x$ iff $\pi`x \notin \pi`x$. Now let $\pi$ be an antimorphism and $x$, *per impossibile*, a Boffa atom. $\pi`x$ is also a Boffa atom, and one of them is self-membered and the other not. Notice $x \in \pi`x$ iff $\pi`x \in x$, since they are Boffa atoms. Now, by duality, if there is a pair of Boffa atoms that are members of each other, there must also be a pair that are not members of each other. So there are two distinct, self-membered Boffa atoms. But this is impossible, as `II` Wins $G_{x=y}$ when x, y are self-membered Boffa atoms, by playing a map σ where $\sigma`z = z$ when $z$ contains both $x$ and $y$ or neither, and $\sigma`z = (z \setminus \{x\}) \cup \{y\}$ if $x \in z$ and conversely if $y \in z$.

If there is to be a unique antimorphism we had better set about finding it. If π is an antimorphism it must satisfy the identity: $\Phi: \ \pi = (j`\pi \circ -)$. This suggests that we devise π by approximation thus

$$\pi = \ldots j^n` - \circ \ldots j` - \circ - .$$

The infinitary expression on the right hand side is easily seen to satisfy the identity Φ. We now note that - is of order 2 and so is $j^n`-$ for any $n$. Also that $j^n`-$, $j^k`-$ commute with one another for all n, k, so we can rewrite the n^{th} finite approximation to the right hand side as

$$\pi_n : - \circ j` - \circ j^2` - \circ \ldots . j^n` -$$

If we apply this permutation to a k-symmetric set, with $k < n$, we can ignore the last $k - n$ terms on the right, since they will not move anything that is k-symmetric. So if x is k-symmetric $\pi_n`x = \pi_m`x$ for any $n, m > k$ and it is this eventually constant value of the π_n that we take to be the value of the canonical antimorphism for argument x. It is now easy to verify that the canonical antimorphism is, indeed, an antimorphism *on the symmetric sets*. Any attempt to extend it to all sets meets only partial success.

The set-theoretical treatment above has been far from rigorous, and no consistency proofs are on offer. This second point should be seen as good news rather than bad, since rather than saying to us that there are no sensible set theories with a universal set, it tells us that they offer us a glimpse of a world so different that interpretation of it in terms of the old are not easy to come by. Besides, history shows that where the available mathematics is sufficiently absorbing, mathematicians are much more likely to get on with developing it than worry about whether it is consistent or not. The philosophical ramifications are simply too tempting to be ignored indefinitely.

[1] Forster, T.E. The status of the Axiom of Choice in Set Theory with a universal set. Submitted to JSL 1979.

[2] Forster. T.E. Quine's NF, an introduction. Cahiers du Centre de Logique Nr 5 CABAY, Louvain-la-Neuve 1982.

[3] La Theorie des Ensembles de Quine, Proceedings Louvain-la Neuve 1981, Cahiers du Centre de Logique Nr 4, CABAY, Louvain-La-Neuve 1982

[4] Specker, E.P. Dualitat. Dialectika 17 (1958) pp 451-465

[5] Forster, T.E. Further consistency and independence results in NF obtained by the permutation method. JSL 1983 pp 231-4