Keith Simmons UNC Chapel Hill

Sets, Classes and Extensions: A Singularity Approach to Russell's Paradox

Russell's famous paradox runs as follows. It seems that some sets, like the set of abstract objects or the set of infinite sets, are members of themselves. Other sets, like the set of U.S. Senators or the set of bridges over the Thames, are not. What about the set of all those sets that are not members of themselves? In the familiar way, we reach a contradiction if we suppose that it is a member of itself, and if we suppose that it isn't.

There is a natural assumption underlying Russell's paradox: the assumption that every predicate has an extension. Call this assumption Naive Comprehension, symbolized by

$$\exists y \forall x (x \in y \leftrightarrow \boldsymbol{\varphi} x)$$

where the variables range over sets, and $\boldsymbol{\varphi}$ stands for any predicate. For example, if we put the predicate 'bridge over the Thames' for $\boldsymbol{\varphi}$, Naive Comprehension guarantees a set of bridges over the Thames. But certain substitutions for $\boldsymbol{\varphi}$ lead to trouble. In particular, put '$x \in x$' for $\boldsymbol{\varphi}$, and we are landed in Russell's paradox.

It is easy to slide, as we have just done, between two distinct versions of Russell's paradox, one in terms of sets and one in terms of extensions. In my view, sets and extensions provide very different settings for Russell's paradox. It is far from clear that a resolution of the paradox in one setting carries over to the other. In this paper, I shall explore the differences between sets and extensions. And I shall propose a new approach to Russell's paradox for extensions.

I. ZF set theory

In Zermelo-Fraenkel set theory (ZF) - nowadays the widely received set theory - Naive Comprehension is abandoned in favour of the axiom schema of *Separation*:

$$\forall z \exists y \forall x (x \in y \leftrightarrow x \in z \,\&\, \boldsymbol{\varphi} x)$$

Given a predicate $\boldsymbol{\varphi}$, we are no longer guaranteed a set of elements that are $\boldsymbol{\varphi}$; rather, given a set z, there is a subset y of those members of z that are **φ**.[1] In particular, given any set z, there is a subset of those members of y that are non-self-members. But there is no set of all the non-self-membered sets, on pain of contradiction.

ZF embodies to a degree a certain conception of set. Think of a set as formed this way: we start with some individuals, and collect them together to form a set. Suppose we start with individuals at the lowest level. At the next level, we form sets of all possible combinations of these individuals. And then we *iterate* this procedure: at the next level, we form all possible sets of sets and individuals from the first two levels. And so on. In pure set theory we start with no individuals, just the empty set. Given the set of all sets at a particular level, the next level will contain the members of its power set. Every set appears somewhere in the hierarchy. This conception of set is the *combinatorial* or *iterative* conception.[2]

We can see how the Separation Axiom fits this picture. Given a set at some level of the iterative hierarchy, at the same level there will be the set of all its members that are **φ**.[3] But Naive Comprehension is *false* on this picture. Put the predicate 'set' or the predicate 'x=x' for **φ** in the schema. Either instance is false: there is no set that serves as the extension of these predicates. At no level of the hierarchy do we reach the universal set of *all* sets; it is just "too big". So Cantor's paradox is avoided.[4] And so is Russell's, because there is no set of exactly the non-self-membered sets. According to the combinatorial/iterative conception, no set is a member of itself. (No set can collect or 'lasso' itself.)[5] So the Russell set, if it existed, would be the universal set. But there is no universal set in the iterative hierarchy.

Now ZF provides a set-theoretical basis for mathematics - but we may wonder if we have a satisfactory resolution of Russell's paradox. For one thing, we expect a well-defined predicate to have an extension. In particular, we expect the self-identity predicate to have an extension - but none of the sets in the ZF hierarchy is the extension of 'x=x'. Or again, since ZF provide a clearcut concept of <u>set</u>, we expect the predicate 'set' to have an extension - and in ZF it doesn't. Note further that in ZF we quantify over sets, and so we need a domain of quantification; but again no set in the hierarchy can serve as this domain.[6] And with the combinatorial/iterative conception in mind, why can't we 'collect together' or 'lasso' all the sets in the ZF hierarchy, and form a collection of them all?

There are (at least) two avenues we might explore here. We might investigate the prospects for a set theory with a universal set. In his system New Foundations (NF), Quine takes this first way.[7] Or we might admit collections that are not sets, following Cantor and von Neumann. We will explore these two avenues in turn.

II. Set theory with a universal set

The axioms of Quine's NF are the axiom of Extensionality and an axiom schema of Naive Comprehension restricted to *stratified* instances: that is, any occurrence of '$\in$' must be flanked by variables with consecutive ascending indices. This is very reminiscent of the simple theory of types, but one has to be careful not to push the parallel too far.[8] In the simple theory of types, a formula has to be stratified in order to be well-formed. In Quine's NF, stratification is not a necessary condition of well-formedness - Quine separates the well-formedness of formulas from the stratification condition. We may even be able to prove the existence of sets given by unstratified formulas.[9] But if we are to prove the existence of a set directly from the comprehension schema, the instance must be stratified. Since the formula '$x=x$' is stratified (vacuously so, because '$\in$' does not appear in it), we *can* prove the existence of a universal set directly from the comprehension schema.

How does NF avoid the paradoxes? Russell's paradox is avoided because '$x\in x$' is unstratified. Cantor's paradox is generated by a diagonal argument: given a set X and its power set P(X), we assume a function f from X onto its power set P(X), and consider the set $\{y\in X|y\notin f(y)\}$. But the formula '$y\in f(y)$' is unstratified, and so the diagonal proof breaks down.

So NF embraces a universal set while dodging the paradoxes. But there is a heavy price to pay. One significant cost is the failure of mathematical induction.[10] And there are other problems. Following Rosser, call a set *Cantorian* if it is of the same size as the set of its unit sets. We expect sets to be Cantorian, but in NF there are non-Cantorian sets. The universal set V is one. Cantor's theorem does not hold for non-Cantorian sets; this is clear in the case of V, which is its own power set. And non-Cantorian sets produce other anomalous effects. We expect the relation of *less-to-greater* among cardinal numbers to be a well-ordering - but Specker proved that this relation is not a well-ordering in NF, because of the non-Cantorian sets.[11] Now in NF cardinal numbers are construed in the Frege-Russell way (according to which, for example, the number 3 is the class of all three-membered classes). And the proof that the Frege-Russell cardinals are well-ordered rests on the Axiom of Choice. So Specker's proof also shows that the Axiom of Choice fails in NF.

These difficulties are substantial. Quine himself abandoned NF, and took instead the second avenue - his subsequent system ML distinguised sets and ultimate classes. There is anyway a further matter: although NF supplies a set as the extension of the predicate 'x=x', it does not supply extensions for certain other predicates, like 'x is a non-self-membered set' and 'x is well-founded'.[12] If we concede that these predicates have no extension, then we lose a primary motivation for a set theory with a universal set: to respect the intuition that every predicate has an extension. The intuition is respected for the predicate 'x=x', but not for others. The systems of Church and Mitchell - two other set theories with a universal set - suffer the same limitation.[13]

Instead of abandoning a universal set, we might admit that there are subcollections of the universal set V (for example, the collection of non-self-membered sets, the collection of well-founded sets) that are not sets. Here one might draw on Vopenka and Hayek's notion of a semiset.[14] A semiset is a subclass of a set, and a proper semiset is a subclass of a set that is not itself a set. Semisets are given via properties and predication.[15] So, in the context of a set theory with a universal set, we might admit as semisets the collections given by the predicates '$x \in x$' and 'x is well-founded', where these semisets are subclasses of V that are not themselves sets.[16] But now we're really taking the second avenue: we are drawing a distinction between collections that are sets and those that are not.

III. Sets and classes

Zermelo himself drew a distinction between two kinds of collection, at least implicitly: "Set theory is concerned with a *domain* B of individuals, ... among which are the sets"[17], but "*the domain B is not itself a set*"[18] The distinction can be traced back to Schröder and Cantor.[19] Cantor drew the distinction this way:

> If we start from the notion of a definite multiplicity of things, it is necessary, as I discovered, to distinguish two kinds of multiplicities (by this I always mean *definite* multiplicities).
>
> For a multiplicity can be such that the assumption that <u>all</u> of its elements "are together" leads to a contradiction, so that it is impossible to conceive of the multiplicity as a unity, as "one finished thing". Such multiplicites I call *absolutely infinite* or *inconsistent multiplicities*. ...
>
> If on the other hand the totality of the elements of a multiplicity can be thought of without contradiction as "being together", so that they can be gathered together into "*one* thing", I call it a *consistent multiplicity* or a "set".[20]

Cantor's immediate concern was Burali-Forti's paradox, generated by the ordinal number of the set of all ordinals. What the Burali-Forti contradiction shows, according to Cantor, was that all ordinals form an "inconsistent, absolutely infinite multiplicity" (<u>op</u>. <u>cit</u>., p.115). It is a multiplicity, but it is *not* a set.

The sharper modern distinction is between *classes* and *sets*. There is a class, but not a set, of all ordinals; and there is a class of all sets and a class of all non-self-membered sets, but no such sets. Von Neumann writes of his axiomatization of set theory:

> ... the present set theory deals even with sets (or "functions") that are "too big" ... . Rather than being completely prohibited, they are only declared incapable of being arguments.... This suffices to avoid the antinomies.[21]

And a little later in the paper, von Neumann switches to the terminology of "sets" and "classes":

> ... "sets" are the sets that (in the earlier terminology) are "not too big", and "classes" are all totalities irrespective of their "size". A class is "capable of being an argument" if and only if it is a set.[22]

So for von Neumann all sets are classes, but not all classes are sets. And those classes that are not sets - the so-called *proper classes* - cannot themselves be members. The mere admission of the existence of, for example, a collection V of all sets does not generate paradox. A further assumption is needed, that V can itself be a member (and in particular a member of itself), and this assumption is false for von Neumann's proper classes.

We can formulate an Axiom of Comprehension for classes as follows:

$$\exists A \forall x(x \in A <-> \boldsymbol{\varphi} x),$$

where A is a variable ranging over classes, x is a variable ranging over sets, and there is the following restriction on $\boldsymbol{\varphi}$: $\boldsymbol{\varphi}$ does not contain quantifiers over classes. $\boldsymbol{\varphi}$ can contain class variables (or parameters), but these must not be bound.[23] Comprehension for classes guarantees the existence of a class of all sets (put 'x=x' for **φ**) and the existence of a class of exactly the non-self-membered sets (put 'x∈x' for **φ**). And similarly there is a class of all ordinals and a class of the well-founded sets. If we suppose these classes to be sets, paradoxes ensue. But in von Neumann's system, the arguments that led to paradox instead establish that these classes are proper classes. And we cannot generate a new family of paradoxes for classes. No proper class can be a member, and so no class can have a proper class as a member. In particular, there can be no class of all classes, and no class of all the non-self-membered classes.

In von Neumann's system, then, there are extensions for the predicates 'set', 'non-self-membered set', 'well-founded set', 'ordinal'; there is a well-determined collection of all the ZF sets; and there is a domain for quantification over sets. Further, the Axiom of Choice is provable in von Neumann's system.[24] But these benefits are costly, both technically and intuitively. The price is technically high because it is provable that in von Neumann's system one cannot prove all instances of the induction schema.[25] And the price is intuitively high, for at least two reasons. First, it seems *ad hoc* and counterintutive to say that a proper class cannot be a member. Why, for example, can't we form its unit class? More generally, the ban on finite classes of proper classes does not seem well-motivated - finite classes do not generate paradox.[26] And if the class and all its subclasses exist, then why doesn't its power class exist too? Second, we expect the predicate 'class' to have an extension, just as we expect the predicate 'set' to have an extension. Moreover, since von Neumann quantifies over proper classes, we seem to require such an extension to serve as the domain of quantification.[27] But according to von Neumann's system, there is no class (nor any other kind of collection) of all classes. The problem has just been pushed back: we are left again with a clear-cut concept - now, the concept of class - which has no extension.

Von Neumann's system has been liberalized in various ways. A natural move is to ease the restriction on **φ** in class comprehension by admitting class quantifiers, a move that has been made by Wang, Morse and Kelley. In this stronger system, we can prove the existence of more sets and more classes of sets. But still a class cannot be a member - again we cannot even form the unit class of a proper class. And still there is no extension for the predicate 'class'.

One can liberalize further, and develop systems in which proper classes are members. Following Levy *et al*,[28] we can introduce the notion of a hyperclass, via the schema:

There is a hyperclass whose members are all classes (and sets) that are $\boldsymbol{\varphi}$,

where $\boldsymbol{\varphi}$ is unrestricted.[29] And we can go further still. Levy *et al* present a two-tier theory, with sets in the lower tier and classes in the upper tier. In the upper tier, we find the class V of all sets, and its power class PV, and PPV, and so on.[30] These power classes also exist in Ackermann's system of classes.[31] Since such systems can be constructed without falling foul of the paradoxes, the ban on membership for proper classes seems not just counterintuitive, but also an unnecessarily heavy-handed response to the paradoxes. We need only a ban on *self-membership* (and more generally on *unfounded* classes) - and the ban is preserved even in these liberalized systems. And with this ban we still pay the same high price: there is no class of all classes, no domain for quantification over classes, and no extension for the predicate 'class'.

IV. Extensions

We have discussed a wide variety of set theories and class theories. Setting aside their various technical difficulties, there is one persistent failing that they share - they fail to provide an extension for certain predicates. They fail to do justice to a conception of set quite different in spirit from the iterative/combinatorial conception. The alternative conception goes like this: to any predicate that denotes a well-determined condition or concept (in Frege's sense), such as the predicates 'abstract' or 'set' or 'class', there corresponds the collection of those things to which the predicate applies - the collection of abstract things, the collection of sets, the collection of classes. Call this the *predicative* conception - and call the collection of things to which a given predicate applies the *extension* of the predicate.

It is clear that Frege would be quite opposed to the combinatorial conception of set embodied in ZF. Frege writes:

> I do, in fact, maintain that the concept is logically prior to its extension; and I regard as futile the attempt to take the extension of a concept as a class, and make it rest, not on the concept, but on single things."[32]

And again:

... [T]here cannot be an empty class if we take a class to be a collection or totality of individuals, so that ... the class consists of individuals or individuals make up the class. In the course of this discussion we have once more had it shown to us that this way of talking is logically useless; that the extension of a concept is constituted in being, not by the individuals, but by the concept itself; i.e. by what is asserted of an object when it is brought under a concept."[33]

An early discussion of these two conceptions - the iterative/combinatorial and the predicative - can be found in König.[34] For reasons tied to his mistaken belief that the continuum cannot be well-ordered, König distinguished between two kinds of classes. As an example of one kind, consider the continuum, defined by König as the set of denumerably infinite sequences $(a_1,a_2,... a_k,...)$ of positive integers. König points out that since there are non-denumerably many of these sequences, they cannot all be defined by linguistic expressions, of which there are only denumerably many. Accordingly, König rejects the idea that we must define or describe a set before it can be said to exist; he works instead with the combinatorial idea. According to König, the elements of the continuum are formed by distinct combinations of the positive integers:

"Through the stipulation that a_1, a_2,... are to be replaced by definite positive integers, it becomes a "definite" sequence, an element of the continuum, which cannot become an object of our thought without being conceptually distinct from any other element. The further stipulation that we consider the *totality* of these "well-distinguished" objects then leads to the continuum".[35]

König goes on to consider, by way of contrast, Cantor's second number class, defined by "the totality of all order types of well-ordered sets of cardinality $_0$."[36] According to König, there is a crucial difference between the two classes: the continuum is given combinatorially and its members are "well-distinguished". But the second number class is given by a *property* or *notion*:

"Its 'elements' are determined by the property of being order types of well-ordered sets of cardinality $_0$. To be sure, we know such elements: **ω**, **ω**+1,...; but this property is only an abstraction, at best a means of distinguishing between objects belonging and not belonging to the class; however, it is certainly not a rule according to which every element of [the second number class] can be formed. What is primary, or fundamental, here is the collective notion, which for this very reason, following Cantor's nomenclature, I would not call a "set" but a "class"; it is only *afterward* that elements belonging to the class are constructed that exemplify this notion."[37]

In König's view, the second number class "cannot be considered to be a completed set, that is, a totality of well-distinguished elements that are altogether conceptually distinct"[38]. Rather, it is a "'set in the process of becoming'"[39]. According to König, the same can be said of the classes that lead to paradox, like the class of ordinal numbers.[40] König concludes that "the distinction here drawn between 'set' and 'class' completely resolves the paradoxes."[41]

König's incomplete classes might remind us of Cantor's "inconsistent multiplicities" that cannot be regarded as one finished thing. But König's and Cantor's view of classes seem to be orthogonal. Acording to Cantor, the mark of a class that is not a set is that it is "absolutely infinite" - it is "too big" to be a set, to put it in von Neumann's terms. Cantor points to the size of the class, and not to any failure to distinguish between its members - for Cantor, even inconsistent multiplicities are "definite". For König, on the other hand, a class fails to be a completed set not because of its size but because its members are not properly distinguished. Notice that the second number class is not "too big" to be a set: it is neither a Cantorian inconsistent multiplicity nor a proper class. According to König, it fails to be a completed set because its elements are determined by a property, and a property will only distinguish between objects in the class and those outside it - it will not distinguish between elements of the class.[42]

So König, like Frege, recognizes two distinct conceptions of sets: the combinatorial and the predicative. But there the agreement ends. In König's view, it is only the combinatorial conception that guarantees a completed set of well-distinguished objects. But according to Frege, it is the predicative conception that is fundamental, and the attempt to base set theory on the combinatorial conception is "futile".

V. Sets, Classes, Extensions

We now have in play three types of collections: sets, classes, and extensions. Is any one of these primary? Is there one to which the others can be reduced? Are there distinct intuitions underlying each type?

We can take ZF sets to be collections conceived of iteratively or combinatorially. And we can take extensions as collections conceived of predicatively. There are clearly distinct conceptions here. Does one reduce to the other?[43]

We cannot reduce extensions to ZF sets: as we have seen, there are predicates of ZF (like 'x=x' and 'x∈x') whose extensions cannot be sets. And even if, for the moment, we think of proper classes as glorified sets, still extensions won't be reducible to sets - no system incorporating proper classes provides an extension for 'class' or 'non-self-membered class'. Quine's NF does provide a set as the extension of 'x=x', but as we have seen, there remain other predicates (like 'x∈x') whose extensions cannot be sets of NF.

Moreover, it is natural to suppose that some extensions are members of themselves - for example, the extension of the predicate 'abstract' is itself abstract, and the extension of 'infinite extension' is itself infinite. But, as we saw, no set of ZF is a self-member - Foundation is a part of standard set theory.[44] And I am not aware of any system with classes as well as sets that allows self-membered proper classes.

On the other hand, we cannot reduce sets to extensions. We saw that the iterative hierarchy is generated by the power set operation. And the assumption that, given a set, there exists a set of all its subsets, is not at all underwritten by the notion of predication: we do not require a predicate or a rule for determining each subset. Indeed, on certain natural assumptions, we have to give up the idea that every set is determined by a predicate. For if we assume that in a given language there are denumerably many predicates, then the members of a nondenumerable set will outrun the predicates. More generally, if we suppose that the predicates of a language always form a set of some cardinality, then, since there will always exist sets of greater cardinality, sets will always outrun predicates. This was a moral of König's discussion: the combinatorial conception yields sets that the predicative conception does not.

In this regard, consider also the Axiom of Choice. When in 1904 Zermelo proved the Well-Ordering Theorem (that every set can be well-ordered), he made fully explicit his reliance on the Axiom of Choice, "the principle that even for an infinite totality of [non-empty] sets there always exist mappings by which each set corresponds to one of its elements".[45] We can think of these mappings as functions that 'choose' an element from each set in the totality, yielding a 'choice set' as output. Zermelo's proof met with immediate opposition because it assumed the existence of mappings and choice sets without defining them.[46]

However, if we adopt the iterative conception such worries about the Axiom of Choice seem quite misplaced. Consider a totality of non-empty sets. There will be some level of the iterative hierarchy at which all these non-empty sets appear, and at this same level all the associated choice sets will also appear, since their members all appear at lower levels. Again sets outrun the expressive capacity of our language - there are sets which are not finitely definable, sets for which we cannot supply a rule or a law.

So I think we should regard extensions and sets as independent notions. They embody quite different ways of thinking about collections. We might regard them as primitive notions, or perhaps as alternative, mutually irreducible conceptions of the more general notion of collection.

The same cannot be said of proper classes - it seems that there is no further distinct intuition behind these collections. It may be that proper classes are in part motivated by the predicative conception - for example, we may want an extension for the well-defined predicate 'set' and invoke a proper class to do the job. But there is no new intuition here - and anyway the predicative conception will not be accommodated by proper classes, since proper classes cannot serve as the extension of certain predicates (e.g. 'class'). Extensions cannot be reduced to proper classes.

And proper classes cannot be reduced to extensions. Proper classes are at least in part motivated by the iterative/combinatorial conception. Von Neumann writes:

> "... If we make the sets that are "too big" and incapable of being arguments capable of being arguments in a new system P, we can still circumvent the antinomies if in turn we admit the sets that are formed from all of these and are "still bigger" (that is too big on P) but declare them incapable of being arguments. The idea is partly the same as the one upon which Russell's "hierarchy of types" rests." (p.404)

The idea is that we can expand our system containing proper classes to a system P in which they are members of still bigger classes which, in the system P, cannot themselves be members. And presumably we can keep going in this way. This should remind us of Levy's discussion of hyperclasses, and the existence in Ackermann's system of V, PV, PPV, etc. These systems treat proper classes in the same iterative way that ZF treats sets. Proper classes now look a lot like sets, occupying sufficiently high levels of the iterative hierarchy.[47] We can think of proper classes as an extra layer or series of layers of sets. And if we think of proper classes that way, then they are reducible to sets.

VI. Extension and Paradox

If I am right that the notions of *set* and *extension* are independent and mutually irreducible, then we have two quite different settings for Russell's paradox. Since sets are generated iteratively, there is no universal set. And since no set is a self-member, there is no Russell set of non-self-members because there is no universal set. Russell's paradox does not arise. Parallel remarks hold for classes, including proper classes: there is no class of all classes, and there are no self-membered classes, and consequently no Russell class of non-self-membered classes.

But the parallel breaks down in the case of extensions. Even if there is no universal extension, we must admit self-membered extensions (like the extension of 'abstract') as well as non-self-membered extensions (like the extension of 'teaspoon'). And a paradox is generated when we ask whether the extension of 'is a non-self-membered extension' is self-membered or not. We need a way out of Russell's paradox for extensions. And, since extensions are not reducible to sets or classes, our best strategy is to develop a direct account of extensions.

In what follows, I shall take extensions to be *extensions of predicates*. And I shall take Russell's paradox to present a challenge for our basic concept of extension, just as the Liar paradox challenges our basic concept of truth. In my view, neither the Liar nor Russell's paradox for extensions is a technical problem restricted to formal languages. If a theory is to resolve the Liar, it must be a theory of *our* basic notion of truth; and if a theory is to solve Russell's paradox for extensions, it must be a theory of *our* notion of extension, tied to the predicates of *our* language. It is in English that we express our intuition that the ZF concept of set must have an extension; and it is in English that we formulate Russell's paradox for extensions. English (or Spanish or ...) is the language in which we discuss alternative set theories, and employ notions like set, class, proper class, virtual class, and so on, and in which we express the intuition that a well-defined predicate has an extension.

Consider a simple paradox for extensions, expressed in ordinary English. Imagine that I am confused about my whereabouts, and wish to refer to predicates I believe to be written on the board next door. I write on the board the following two predicates:

(1) moon of the Earth

(2) unit extension of a predicate written on the board in room 213 Caldwell Hall at noon 7/1/98.

I am myself in room 213. The extension of (1) - ext(1) for short - is clear enough. It has one member, so it is a unit extension. But what about the extension of (2)? To determine its extension, we look at each predicate written on the board, and see whether its extension is a unit extension. Ext(1) is a unit extension, so ext(1) is in ext(2). Is ext(2) in ext(2)? Suppose so; then ext(2) has 2 members, and so is not a unit extension, and so is not in ext(2). Suppose not; then ext(2) has 1 member, and so is a unit extension, and so is a member of ext(2). Either way we obtain a contradiction: we are landed in paradox.

Let us take a closer look at the paradox. In general, we give membership conditions for extensions via the schema

x is in the ext($\boldsymbol{\varphi}$) iff x is $\boldsymbol{\varphi}$,

where $\boldsymbol{\varphi}$ stands for a predicate. An instance of this schema is:

> ext(2) is in ext(2) iff ext(2) is a unit extension of a predicate written on the board in room 213 Caldwell Hall at noon 7/1/98.

This instance is utilized in the reasoning that generated paradox. We start by supposing that ext(2) is in ext(2): we suppose that the left hand side of the biconditional holds. So ext(2) has two members, so the right hand side does not hold. So the left hand does not hold - contradiction. Now suppose that ext(2) is not in ext(2); that is, that the left hand side does not hold. Then ext(2) has one member, so the right hand side holds. And so the left hand side holds - and we have a contradiction again.

So we cannot determine an extension for (2), because we cannot consistently give it self-membership conditions. (2) is a pathological predicate, in the sense that (2) cannot consistently be supplied with an extension. That is, we may conclude:

(P) (2) does not have a well-determined extension.

Since (2) does not have a well-determined extension, it does not have a unit extension. In contrast, ext(1) is a unit extension. So the only *unit extension of a predicate on the board* is ext(1). Consider the underlined predicate in the previous sentence - call it (2*). The extension of

(2*) is unproblematic: its only member is ext(1). That is:

(R) (2*) has a well-determined extension.

Now this reasoning calls for explanation.[48] (2) and (2*) are predicate tokens of the same type, and yet one has a determinate extension and the other does not. When we tried to determine an extension for (2), we were landed in contradiction. But an extension for (2*) is readily obtained. How can we explain the difference? We might suppose that we are confronted with a strengthened paradox, and try to block the reasoning in some way. However, the reasoning leading to (R) appears to be valid, and (R) appears to be true. Since the reasoning is natural, we should not block it by artificial, *ad hoc* means. Rather, my strategy is to provide an analysis of the reasoning that preserves its validity and the truth of (R).

The puzzle is this: (2) and (2*) are composed of the very same words with the very same linguistic meaning, yet one is pathological and the other is not. A shift in extension without a shift in meaning suggests some *pragmatic* difference between the two tokens. And there are a number of differences between their respective contexts. There are differences of time and place and perhaps speaker too. Still, the familiar contextual parameters of speaker, time, and place do not tell the whole story.

Another difference is that the two tokens are embedded in different stages of the strengthened discourse. We may split the discourse into two stages. At the first stage I produce the predicates (1) and (2), and we argue to the subconclusion (P); at the second stage, we reason from (P) to (R). In general, the correct interpretation of an expression or a stretch of discourse may depend on the larger discourse in which it is embedded. The reasoning from (P) to (R) is second not merely in temporal order, but in logical order as well. The second stage starts out from a subconclusion, namely (P), established by the first stage of the argument. We can think of the second stage as *reflective* with respect to the first: at the second stage of the reasoning, we reflect on (2)'s lack of a well-determined extension, established by the first stage. This logical order constitutes a difference in the relation that each stage of the discourse bears to the discourse as a whole.

A further difference is found in speaker's intentions. At the two stages of the reasoning, there are different intentions towards (2). At the second stage, our intention is to treat (2) as pathological and see where that leads us. But when I first produced (2) I had no intention of producing a pathological predicate: my intention was to pick out the unit extensions of predicates written next door.

There is another shift between the two stages, a shift of relevant information. The pathological nature of (2) is established only at the completion of the first stage. But this information is available throughout the reflective second stage of the reasoning. The reasoning of the second stage should be interpreted as incorporating this information.

So we distinguish two contexts: one in which we reason to (P), and the second in which we reason from (P) to (R). Call these *the initial context* and *the reflective context* respectively. Between these contexts there is a shift in a number of contextual parameters, shifts in speaker, time, place, discourse position, intention, and relevant information. Recall what we want to explain. Predicates (2) and (2*) are two tokens of the same type, one occurring in the initial context, the other occurring in the reflective context. Though these tokens are composed of the same words with the same meaning, one has a determinate extension and the other does not. A pragmatic explanation is indicated, one that takes account of the shift in the contextual parameters. If we accept the appropriateness of a pragmatic explanation, then we should expect to find a term occurring in the two tokens that is context-sensitive. I propose that, in the absence of any reasonable alternative, we explore the idea that the expression 'extension' is itself the context-sensitive term.

Let 'extension_I' abbreviate 'extension in the initial context', and let 'extension_R' abbreviate 'extension in the reflective context'. We take the occurrence of 'extension' in (2) to be sensitive to the context in which it occurs. Accordingly, (2) is represented as:

(2) unit extension_I of a predicate written on the board in room 213 Caldwell Hall at noon 7/1/98.

Corresponding to the two contexts are two schemas, the I-schema and the R-schema. The I-schema is

$$x \text{ is in } \text{ext}_I(\boldsymbol{\varphi}) \text{ iff } x \text{ is } \boldsymbol{\varphi},$$

where '$\text{ext}_I(\boldsymbol{\varphi})$' abbreviates 'the extension_I of $\boldsymbol{\varphi}$'. At the initial stage of the reasoning, we determine the extension of (2) via the I-schema, and obtain this instance:

$\text{ext}_I(2)$ is in $\text{ext}_I(2)$ iff $\text{ext}_I(2)$ is a unit extension_I of a predicate on the board in room 213

Caldwell Hall at noon 7/1/98.

And we reach a contradiction. We cannot determine an extension for (2) via the I-schema. Our attempt to do so lands us in contradiction. So we conclude:

(P) (2) does not have a well-determined extension_I.

This is the subconclusion we reach at the culmination of the first stage of the reasoning.

We now reflect on the fact that (2) is a pathological predicate. At this reflective stage of the reasoning, we determine which predicates on the board have unit extensions_I. It's clear that (1) has a unit extension_I, and it's clear that (2) does not, because (2) has no determinate extension_I. So the only *unit extension$_I$ of a predicate on the board* is ext(1). In producing (2*) here we have in effect repeated (2). But we have repeated (2) in a new reflective context, in which we no longer determine extensions via the I-schema. Instead we determine an extension for (2*) via the R-schema, according to which

$$x \text{ is in } \text{ext}_R(\boldsymbol{\varphi}) \text{ iff } x \text{ is } \boldsymbol{\varphi}.$$

And via this schema, we can determine an extension for (2*). Consider the instance:

x is in ext_R(2*) iff x is a unit extension_I of a predicate written on the board in room 213 Caldwell Hall at noon 7/1/98.

The right hand side of this biconditional is true for x=ext(1), and false otherwise. And so ext(1) is the sole member of the extension_R of (2*). And we may conclude:

(R) (2*) has a well-determined extension_R.

(2) and (2*) are semantically indistinguishable. The difference between them is purely pragmatic. It is a matter of the schema by which their extensions are determined.[49] At the first stage of the reasoning, the extension of (2) is determined via the I-schema; at the second stage of the reasoning, the extension of (2*) is determined by the R-schema. When I first produce the token (2), I am picking out the extensions_I of predicates on the board in room 213. So when we determine the extensions of these predicates, we use the I-schema. It is the I-schema that is implicated here. We go on find that we cannot determine an extension for (2) via the I-schema, on pain of contradiction. At the second stage of the reasoning, we repeat the words of (2), producing the token (2*) in a new reflective context. To determine the extension of (2*), we use a schema that is sensitive to (2)'s pathology - and here the implicated schema is the reflective R-

schema. That the R-schema is reflective in this way is a product of the reasoning of the first stage, the assessment of (2) as a pathological predicate, and the intention to treat (2) as such. With the shift in context, there is a shift in the implicated schema.

Notice that if we determine the extension of the predicate token (2) by the R-schema, we find that (2), like (2*), has sole member ext(1); and if we determine the extension of (2*) by the I-schema, we find that (2*), like (2), does not have a well-determined extension. The expression 'the extension_R of' denotes an operation that carries predicates to their extensions_R; in particular, it carries (2) and (2*) to a well-determined extension. The expression 'the extension_I of' denotes an operation that is undefined for arguments (2) and (2*). So the expression 'the extension of' is a context-sensitive term that shifts its reference according to context.

VII. A Singularity Proposal

The question naturally arises: What is the relation between different occurrences of the term 'extension'? More specifically, what is the relation between the expressions 'extension_I' and 'extension_R'?

A possible response here is a Tarskian one: when we move from the first stage of the reasoning to the second, we push up a level of language. The expressions 'extension_I' and 'extension_R' belong to distinct languages. On such a hierarchical account, the domain of the extension_R operator properly contains the domain of the extension_I operator.

But there are a number of serious difficulties with the Tarskian approach. Consider how such an account might go. We might start with a fragment of English free of the term 'extension'. This will be the language at the first level, call it L_0. At the next level of language L_1, we can talk about the extensions of predicates of L_0. That is, L_1 contains the expression 'extension of a predicate of L_0', or 'extension_0' for short. In general, $L_{\sigma+1}$ contains the expression 'extension_σ', which applies exactly to the predicates of L_σ.[50]

An immediate worry with this Tarskian account is its artificiality: why suppose that a natural language like English is regimented and stratified in this way? But there are other problems too. Suppose I use the phrase "the extension of 'teaspoon'". According to the present

Tarskian line, this is a phrase of the language L_1, represented by 'the extension$_0$ of teaspoon'. But then I have picked out an extension operator that is massively restricted: only predicates of L_0 are in its scope. Gödel remarks of Russell's type theory that "...each concept is significant only ... for an infinitely small portion of all objects."[51] A similar complaint can be made here. Why suppose that my use of 'extension' is restricted to the predicates of L_0.

Further, it is hard to see how levels can be assigned in a systematic way. How are we to interpret a given expression containing the term 'extension'? To which language does it belong? Except in very simple cases, we will have little basis for an assignment of one level rather than another. And what level should we assign to a global statement like 'Every predicate of English either has a well-determined extension or it doesn't'? Any assignment of a level here will compromise the global nature of the statement.

These difficulties also arise for hierarchical accounts of truth and denotation. But there is also a special difficulty for a hierarchical account of extensions. It is an important mark of extensions that some are self-members. For example, the extension of the predicate 'extension with more than one member' has more than one member, and so is a self-member. But on the hierarchical approach, this predicate belongs to some level of language, and is analyzed as 'extensionσ with more than one member' (where σ is an ordinal) - or in full, 'extension of a predicate of L_σ with more than one member'. And the predicate itself is a predicate of L_{o+1} and not of L_o. So the extension of this predicate is not a self-member - it contains extensions of predicates of L_o only. The hierarchical account cannot accommodate self-membered extensions. A distinctive feature of extensions is regimented away.

I am after an account of extensions which accommodates self-membership, and avoids counterintuitive restrictions. The account I shall offer is in a strong sense anti-hierarchical. The leading idea is that occurrences of 'extension' in pathological predicates are *minimally* restricted. At this point, a pragmatic principle of interpretation comes into play: the principle of *Minimality*. According to Minimality, restrictions on occurrences of 'extension' are kept to a minimum: we are to restrict the application of 'extension' only when there is reason to do so.

Suppose you say: "The extension of 'natural number' is infinite". Here, your use of 'extension' is quite unproblematic. Should (2) be excluded from its scope? Minimality says no.

And this is surely plausible. As we've seen, *within* its context of utterance, (2) is pathological; but outside that context, an extension can be determined for it. Since your utterance is quite unrelated to (2), we have no reason to interpret your utterance as in some way pathologically linked to (2). Assessed from outside its context, (2) does have an extension (whose sole member is ext(1)), and so we have no reason to withhold (2) from the scope of 'extension' in your utterance. It would be a poor interpretation that implicated your utterance in semantic pathology. In general, speakers do not usually aim to produce pathological utterances, or utterances implicated in paradox.[52] By adopting Minimality, we respect this pragmatic fact.

Further, we expect all well-formed predicates to have an extension. In particular, since (2) is a well-formed predicate, we expect it to be within the scope of your use of 'extension'. We do expect any solution to a genuine paradox to require some revision of our intuitions. But the more a solution conflicts with our intuitions, the less plausible that solution will be. Minimality keeps surprise to a minimum. Almost all predicates are within the scope of any use of 'extension'. We are sometimes forced to restrict 'extension' - we must, for example, limit the scope of its occurrence in (2) by excluding (2) from its scope. Still, according to Minimality, we exclude only those predicates that cannot be included.

So my proposal identifies what I shall call *singularities* of the extension operator. In general, given a context C, if an extension of a predicate token F cannot be determined by the C-schema, then F is a *singularity* of 'extension_C'. Further, if the C-schema is the schema which determines the extension of F in its context, then F is *pathological*. So (2) is a singularity of 'extension_I', and (2) is also pathological, since in its context its extension is determined by the I-schema. (2*) is a singularity of 'the extension_I of', but (2*) is *not* pathological, since in its context its extension is determined by the R-schema, and (2*) does have an extension_R.

(2) is a singularity only in a context-relative way - there is an appropriate reflective context in which (2) is in the scope of the extension operator. In the case of our strengthened discourse, we *can*, from our subsequent reflective context, determine a definite extension for (2). (2) is a singularity relative only to the I-schema.

No occurrence of 'extension' is without singularities. For example, suppose you say, innocently enough, "the extension of 'natural number'". But you perversely continue: "unioned with any non-self-membered extension of a nine-word predicate in utterance U", where you stipulate that utterance U is your very utterance. Within your utterance, there is no shift to a reflective context. Consequently, if your first use of 'extension' is represented by 'extension$_C$', your subsequent use is represented by 'extension$_C$' as well. Then the only nine-word predicate in your utterance is represented by: 'non-self-membered extension$_C$ of a nine-word predicate in utterance U'. It is straightforward to check that this predicate token is a singularity of 'extension$_C$': the C-schema cannot provide an extension for it.[53] It may be that there are no *actual* phrases uttered that force restrictions on a given occurrence of 'extension'; there may be no *actual* singularities. But continuations like yours here are always *possible* - and they yield singularities.

We are now in a position to see the anti-Tarskian nature of the singularity proposal. As a consequence of Minimality, the singularity proposal is not hierarchical. Recall the conclusion of our strengthened reasoning, suitably represented:

(R) (2*) has a well-determined extension$_R$.

Consider the occurrence of 'has a well-determined extension$_R$' in (R) - call this predicate token (3). Let us ask: Does (3) have an extension$_I$? (3) is <u>not</u> a singularity of 'extension$_I$'. So by Minimality, (3) is <u>not</u> excluded from the scope of 'extension$_I$'. When we instantiate the I-schema to (3), we obtain:

x is in ext$_I$(3) iff x has a well-determined extension$_R$.

This schema determines the extension$_I$ of (3). Notice in particular that putting (2*) or (2) for x does not lead to trouble. Both (2) and (2*) have the same well-determined extension$_R$, so in each case the right hand side is true, and so both (2) and (2*) are in ext$_I$(3). So a predicate containing an occurrence of 'extension$_R$' has an extension$_I$. For the Tarskian, this would amount to an unacceptable mixing of language levels. According to the singularity proposal, there are no such levels.

Moreover, there are singularities of 'extension$_R$' that are *not* singularities of 'extension$_I$'. Having inferred (R), you might perversely add:

> "Now form the union of the well-determined extension of (2*) and any empty extension of a nine-word predicate in utterance V",

where you stipulate that V is your perverse addition. Let (4) be the predicate token composed of the last nine words of V. Given the context, (4) is represented by "empty extension$_R$ of a nine-word predicate in utterance V". It is easy to check that (4) is a singularity of 'extension$_R$' - the R-schema cannot provide an extension for (4).[54]

But now, just as we reflected on (2), we can reflect on (4). We have determined that (4) is a pathological predicate, and has no extension$_R$, empty or otherwise. Since (4) is the only nine-word predicate in V, it follows that there is no empty extension$_R$ of a nine-word predicate in utterance V. We've just produced a token (4*) of the same type as (4), and in the reflective context we can determine an extension for (4*), and for (4) - the empty extension. And according to Minimality, any neutral schema will determine an extension for (4) as well. In particular, (4) is not a singularity of 'extension$_I$'. (4) has an extension$_I$. On a Tarskian account, the domain of the extension$_I$ operator will be properly included in the domain of the extension$_R$ operator. According to the singularity proposal, neither domain includes the other.

So on my proposal, our simple paradox is to be treated by the identification and exclusion of singularities.[55] We treat everyday English not as a hierarchy of languages, but as a single language. We do not divide up the term 'extension' between infinitely many languages; rather we identify singularities of a single, context-sensitive term.

Gödel noted that Russell's theory brings in a new idea for the solution of the paradoxes:

> It consists in blaming the paradoxes ... on the assumption that every concept gives a meaningful proposition, if asserted for any arbitrary object or objects as arguments.[56]

Gödel goes on to say that the simple theory of types carries through this idea on the basis of a further restrictive principle, by which objects are grouped into mutually exclusive ranges of significance, or types, arranged in a hierarchy.

Gödel suggests that we reject this principle, while retaining the idea that not every concept gives a meaningful proposition for any object as argument:

> It is not impossible that the idea of limited ranges of significance could be carried out without the above restrictive principle. It might even turn out that it is possible to assume every concept to be significant everywhere except for certain 'singular points' or 'limiting points', so that the paradoxes would appear as something analogous to dividing by zero. Such a system would be most satisfying in the following respect: our logical intuitions

> would then remain correct up to certain minor corrections, i.e. they could then be considered to give an essentially correct, only somewhat 'blurred', picture of the real state of affairs.[57]

I take my singularity proposal to be very much in the spirit of Gödel's remarks. And we can claim for it the same satisfying feature: our logical intuitions about extensions are almost correct. It is only in pathological or paradoxical contexts that we may mistakenly suppose that certain predicates have extensions when they do not -- and in such cases our uses of 'extension' require only *minimal* corrections. A second intuition that requires revision is that the term 'extension' is a *constant*. Strengthened discourses indicate that the term shifts its scope according to context. But these shifts are kept to a minimum.

In correcting both these intuitions, we retain a single expression which undergoes minimal changes in its extension according to context. There is no wholesale revision of the notion of extension; no division of 'extension' into infinitely many distinct terms; no splitting of everyday English into an infinite hierarchy of languages.

VIII. Ungroundedness and the identification of singularities

Thus far, I have spoken of pathological predicates and singularities more or less informally, largely by way of examples. I shall now sketch a more general treatment of these notions.

Consider again our simple paradox. Why does predicate (2) lead to trouble? In trying to determine an extension for (2), we are directed to the predicates written on the board - but (2) is one of these predicates. Part of the story, then, appears to be a certain kind of ungroundedness: to determine the extension of (2) we must first determine the extension of (2), among other predicates. We might think in terms of a kind of dependency tree:

```
           (2)
          /   \
      (1)      (2)
              /   \
           (1)     (2)
                  /   .
                        .
                          .
```

where the extension of a predicate higher on a branch depends on the extension of any predicate lower on that branch.

To make these ideas more rigorous, we will need a more precise representation of (2). Suppose we represented the token (2) as an ordered pair <type(2), ext_I>, where the first element is the type of (2), and the second indicates the representation of the occurrence of 'extension' in (2). Then something is missing. This representation does not distinguish (2) from (2*), yet the former predicate token is pathological and the latter is not. There is something more to consider: the schema which determines the token's extension. The implicated schema for (2) is the I-schema (and for (2*) the R-schema). So we can represent (2) more perspicuously as an ordered triple - <type(2), ext_I, ext_I> - where the third element indicates the schema which determines an extension for (2) in its context of utterance. Let us call this the *primary representation* of (2). In general, the *primary representation* of a predicate token containing 'extension' indicates in order the type of the token, the occurrence of 'extension' in the token, and the implicated schema which determines an extension for the token in its context. The primary representation of (2*) is the triple <type(2), ext_I, ext_R>.

We can also assess (2) from contexts other than its own. For example, we can assess (2) from the subsequent reflective context. We may represent such an assessment as the triple <type(2), ext_I, ext_R>. We can think of this as a *secondary representation* of (2), since here an extension for (2) is not determined by its implicated schema. Notice that this secondary representation of (2) is identical to the primary representation of (2*). This is appropriate, since both (2) and (2*) have the same extension_R.

Now, with our dependency tree in mind, let us introduce the notion of a *determinant* of a predicate. As we've seen, to determine an extension for (2), we must first determine the extensions of predicates to which (2) refers - (1) and (2). We will call these predicates the *determinants* of (2). Notice that (2) makes reference to the <u>extensions$_I$</u> of (1) and (2), since the occurrence of 'extension' in (2) is represented by 'extension$_I$'. So to determine what extension (2) has (if any), we need to determine the extensions$_I$ of (2)'s determinants. To determine (2)'s extension, then, the schema implicated for its determinants is the I-schema. In general, let the primary representation of a predicate F be the triple <type(F), $ext_{c\alpha}$, $ext_{c\beta}$>. Let G be a determinant of F. To determine the extension of F, we should determine the extension of G via the c_α-schema.

We are now in a position to introduce the notion of a *primary tree*. The primary tree for (2) looks like this:

<type(2), ext_I, ext_I>
/ \
type(1) <type(2), ext_I, ext_I>
/ \
type(1) <type(2), ext_I, ext_I>
/ .
type(1) .
.

To construct the primary tree for (2), we start with the primary representation of (2), the triple <type(2), ext_I, ext_I>. This is the node at the top of the tree. At the second tier are the determinants of (2), suitably represented. (1) contains no context-sensitive terms, and so it is suitably represented via its type.[58] This is not so for the other determinant of (2), namely (2) itself. Following the remarks of the previous paragraph, (2) is to be assessed via the I-schema. Accordingly, we represent C at the second tier as <type(C), ext_I, ext_I>. This is the primary representation of C again, which in turn generates a third tier of nodes. And so on, indefinitely.

The primary tree for (2) has an infinite branch, *on which the primary representation of (2) repeats*. This indicates that (2) is a pathological predicate. The repetition of its primary representation shows that an extension for (2) cannot be determined by the I-schema - and so we can also say that (2) is a *singularity* of 'the extension$_I$'.

In general, we construct the primary tree for a predicate F as follows. The top node is the primary representation of F - let it be $<\text{type}(F), \text{ext}_{c\alpha}, \text{ext}_{c\beta}>$. At the second tier are the determinants of F, suitably represented. Branches lead from each of these determinants to their own determinants. And so on.

Consider now the primary tree for (2*). To determine an extension for (2*) in its context of utterance, we consider (1) only, since in the reflective context (2) is explicitly excluded as a determinant, because it is pathological. In general, the identification of the determinants of a predicate F is guided by both semantic and pragmatic considerations. On the semantic side, we consider the predicates to which F refers. We take these to be the determinants unless there are overriding contextual considerations, as in the case of (2*).

Accordingly, the primary tree for (2*) is:

$$
\begin{array}{c}
<\text{type}(2), \text{ext}_I, \text{ext}_R> \\
| \\
\text{type}(1)
\end{array}
$$

This well-founded tree indicates the groundedness of (2*).

Suppose now we determine the extension of (2) from some neutral context, quite unrelated to the contexts of (2) and (2*). Let the secondary representation of (2) here be the triple $<\text{type}(2), \text{ext}_I, \text{ext}_N>$. The secondary tree for (2) is this:

$$
\begin{array}{ccccccc}
 & <\text{type}(2), \text{ext}_I, \text{ext}_N> & & & & & \\
/ & & \backslash & & & & \\
\text{type}(1) & & <\text{type}(2), \text{ext}_I, \text{ext}_I> & & & & \\
 & / & & \backslash & & & \\
 & \text{type}(1) & & <\text{type}(2), \text{ext}_I, \text{ext}_I> & & & \\
 & & / & & . & & \\
 & & \text{type}(1) & & & . & \\
 & & & & & & .
\end{array}
$$

(2)'s secondary representation does not repeat on the infinite branch, but its primary representation does.

We can think of (2)'s secondary representation as standing above the circle in which (2)'s primary representation is caught. And we can given expression to this idea via the notion of a *pruned* tree. We *prune* the tree here by terminating the infinite branch at the first occurrence of a non-repeating node.[59] The pruned tree is:

$$\begin{array}{c} \langle \text{type}(2), \text{ext}_I, \text{ext}_N \rangle \\ | \\ \text{type}(1) \end{array}$$

The pruned tree indicates that we can determine an extension for C via the N-schema. We do not restrict the scope of 'extension$_N$' if there is no need. (2) does have an extension$_N$. (2) is *not* identified as a singularity of 'extension$_N$', *in accordance with Minimality*.

We can now give a more general characterization of the notions of *ungroundedness* and *singularities*. Let F be a predicate token, with primary representation <type(F), ext$_{c\alpha}$, ext$_{c\beta}$>. Take F's primary tree, and prune it.[60] If F's primary representation repeats on an infinite branch of F's pruned tree then F is *ungrounded*. Further, F is a *singularity* of 'extension$_{c\beta}$', since the repetition of F's primary representation indicates that the c_β-schema fails to determine an extension for F. It is easy to check that (2)'s pruned tree is just its primary tree - we cannot prune its infinite branch, since it has no non-repeating nodes. So (2) is ungrounded, and it is a singularity of 'extension$_I$'.

(IX). Russell's paradox and extensions

The story is essentially the same for a version of Russell's paradox. Suppose I write on the board:

A. abstract.

B. teaspoon

C. non-self-membered extension of a predicate on the board.

We can reason to C's pathology, and conclude that C does not have an extension_{IN} (where 'extension_{IN}' abbreviates 'extension in the initial context'). So only predicates A and B have extensions_{IN}, and only ext(B) is a *non-self-membered* $\textit{extension}_{IN}$ *of a predicate on the board*. Here, in the present reflective context, we have produced a token of the same type as C, call it C*, which has a well-determined extension_{RE} (where 'extension_{RE}' abbreviates 'extension in the reflective context'). The pruned tree for C is just its primary tree, and this unfounded tree will indicate that C is ungrounded, and a singularity of 'extension_{IN}'. In contrast, the pruned tree for C* (which is just C*'s primary tree) is well-founded, and C* is grounded. Both C and C* have extensions_{RE}; neither have extensions_{IN}.

There is an interesting variant of the present Russell case: retain A and B, but replace C by

C'. self-membered extension of a predicate on the board.

While C' does not generate a contradiction, it is nevertheless pathological. (C' stands to C as the Truth-Teller - "This sentence is true" - stands to the Liar.) According to our account, C' is ungrounded and a singularity of 'extension_{IN}'.

A broader version of Russell's paradox for extensions might be couched in terms of all predicates of, say, English, along the following lines:

> Among the predicates of English there are those, like 'abstract' that have self-membered extensions. And there are others, like 'teaspoon', that have non-self-membered extensions. Now consider the predicate 'non-self-membered extension'. In the usual way, we reach a contradiction if we suppose that its extension is self-membered and if we suppose it is not self-membered. We conclude that the predicate does not have a well-determined extension.

We may now continue:

> Since the predicate has no extension, it is not among the predicates of English with well-determined extensions. Once we set it aside (and any other related pathological predicate), we will be left with just those predicates of English which have well-determined extensions. And among these extensions will be the non-self-membered extensions (like the extension of 'teaspoon'). We have just produced a token of our problematic predicate, and this token has a well-determined extension.

This version of Russell's paradox is presented in terms of predicate-*types*, like 'abstract', 'teaspoon', 'non-self-membered extension'.[61] But according to the singularity account, this last predicate-type contains a context-sensitive term. We cannot determine an extension for the type *simpliciter*; an extension can be determined only for the type paired with a context (or, equivalently, for a token of the type). According to the contextual analysis, at the first stage of the reasoning we try to determine an extension for the type paired with the initial context i. That is, we try to determine an extension for 'non-self-membered extension$_i$' - call this predicate D. Notice that when we suppose that D is a self-membered extension or a non-self-membered extension, we produce tokens of 'extension' that are tied to the initial context. We reach a contradiction because we try to determine that extension of D via the i-schema.

We can capture the second stage of the reasoning as follows:

> Since D has no extension$_i$, it is not among the predicates of English with well-determined extensions$_i$. Once we set D aside (along with any other predicates for which the i-schema cannot provide an extension), we will be left with just those predicates of English that have well-determined extensions$_i$. And among these extensions$_i$ will be the non-self-membered extensions$_i$ (like the extension$_i$ of 'teaspoon'). We have produced a token of the type of D, call it D*, and this token has a well-determined extension$_r$ (where 'extension$_r$' abbreviates 'extension in the reflective context').

We can put all this in terms of dependency trees. The determinants of D are the predicate-types of English and D itself. The primary representation of D is the triple <type(D), ext$_i$, ext$_i$>. The primary tree for D has an infinite branch on which the primary representation of D repeats. D's pruned tree is just its primary tree. This tree indicates that D is ungrounded, and a singularity of 'extension$_i$'.

The primary representation of D* is <type(D), ext$_i$, ext$_r$>. The determinants of D* are the unproblematic predicate-types of English, since D is explicitly excluded as a determinant. So the primary tree for D* is well-founded, and D* is grounded.

As before, the move from 'extension$_i$' to 'extension$_r$' (or from 'extension$_{IN}$' to 'extension$_{RE}$') is not a move to a higher level of language. Neither expression is more comprehensive than the other; each has singularities that the other does not.[62] We resolve the Russell paradox as we resolved our simple paradox: the term 'extension' is a context-sensitive expression that is minimally restricted on any occasion of use.

(X) Universality

Natural languages are remarkably flexible and open-ended. If there is something that can be said, it might seem that a natural language like English has at least the potential to say it. Natural languages evolve; they always admit of expansion, of increased expressive power. Tarski speaks of the "all-comprehensive, universal character" of natural language, and continues:

> The common language is universal and is intended to be so. It is supposed to provide adequate facilities for expressing everything that can be expressed at all, in any language whatsoever... .[63]

I think that the singularity proposal goes a long way to accommodate this intuition. An occurrence of the context-sensitive term 'extension' is as close to universal as it can be without contradiction - it applies to all predicates except its singularities. Moreover, according to the singularity proposal, even predicates that are singularities relative to a given context are in the scope of 'extension' in other contexts (such as an appropriate reflective context or a neutral context). So the application of any occurrence of 'extension' is almost global, and those predicates that prevent its application from being fully global are captured by other uses of 'extension'.

We can take these points a little further. Many have thought that the goal of a universal language is unattainable, because any theory of the semantical predicates of a language must be couched in an essentially richer metalanguage.[64] But if we adopt the singularity proposal, then we are not driven to this conclusion. Let L be that fragment of English that is free of context-sensitive terms. Let $L^{/}$ be the result of adding to L the context-sensitive term 'extension'. We can think of $L^{/}$ as the object language for our singularity account. Now in this paper I have not attempted to give a formal theory of 'extension'. I have only described some notions - e.g. <u>primary tree</u>, <u>pruned tree</u>, <u>groundedness</u>, and <u>singularity</u> - that I take to be central to such a theory. But suppose for a moment that we have a formal singularity theory.[65]

With any theory of context-sensitive terms, there is an inevitable separation of the object language and the language of the theory. We can take the language of the theory to be a classical formal language which quantifies over contexts, and in which context-sensitive terms do not appear. Unlike the object language, the language of the theory is context-independent. Now the language of the theory - call it T - will be in certain ways richer than the object language $L^{/}$. For example, T will contain the predicate

(Ext) 'predicate of $L^{/}$ that has a well-determined extension in some context'.

Among the predicates to which (Ext) applies are all the singularities. Given that any occurrence of the context-sensitive term 'extension' has singularities, the predicate (Ext) will be in this way more inclusive than any occurrence of the context-sensitive predicate 'has a well-determined extension'.

This may tempt us to suppose that (Ext) is more comprehensive than any occurrence of the context-sensitive predicate 'has a well-determined extension' in the object language, and that T is a Tarskian metalanguage for $L^{/}$. But the temptation should be resisted. Let us see why.

By Minimality, (Ext) will not be excluded from the scope of any occurrence of 'extension' in $L^{/}$. Only the singularities of the given occurrence are excluded - and unproblematic predicates expressed in T will not be identified as singularities. (Intuitively, the theory tells us what is <u>excluded from</u> the scope of occurrences of 'extension', <u>not</u> what is included; we take a 'downward' route rather than an 'upward' route.) This shows that T is not a Tarskian metalanguage for $L^{/}$, since ordinary context-sensitive uses of 'extension' apply to predicates like (Ext) that are expressible in T but not in $L^{/}$. Again, according to the singularity proposal, paradox is avoided not by a Tarskian ascent, but by the identification and exclusion of singularities.

Moreover, if we suppose that T is a classical formal language, free of context-sensitive terms, then T cannot contain the expression 'extension of a predicate of T', on pain of contradiction.[66] So we can generate from this formal language a Tarskian hierarchy of formal languages, each containing an extension predicate for the preceding language. But none of the predicates expressible in these languages are identified as singularities, and so none are excluded from the extensions of our ordinary context-sensitive uses of 'extension'. To speak metaphorically, our context-sensitive uses of 'extension' arch over not only the predicates of T, but also all the predicates expressed by the languages of this hierarchy.

So we do not take the formal hierarchy generated from T to explicate our concept of extension. The levels do not correspond to any stratification of the context-sensitive term 'extension'. The singularity proposal abandons this Tarskian route. For the Tarskian, questions about the extent of the hierarchy and quantification over the levels will present special difficulties. Of course, these are substantial questions, quite independently of any particular proposal about 'extension' in English. But they present no special difficulty for the singularity account. According to the singularity account, an ordinary context-sensitive use of 'extension' applies "almost everywhere", failing to apply only to those predicates that are pathological in its context of utterance. By Minimality, when we use the term 'extension' we point to as many predicates as we can point to from our context of utterance. These predicates include those expressible in T, and at any level of the Tarskian hierarchy which can be generated from that theoretical language (whatever the extent of the hierarchy).

To return to universality. With respect to the concept of extension, a universal language would have a term applying to <u>all</u> the predicates of the language. Such a language would be subject to Russell's paradox. According to the singularity account, any use of 'extension' has its singularities. But a use of 'extension' applies everywhere else - even to predicates couched in the language T of the theory, and the hierarchy of languages generated from T. Indeed, a use of 'extension' applies to any predicate of any language, as long as the predicate is not identified as a singularity. If we adopt the singularity proposal, then any use of 'extension' is as close to global as it can be. In this way we respect Tarski's intuition that we can say everything there is to say.

<u>Endnotes</u>

1. Zermelo explains the Separation Axiom this way: "... sets may never be *independently defined* by means of this axiom but must always be *separated* as subsets from sets already given; thus contradictory notions such as "the set of all sets" or "the set of all ordinal numbers" ... are excluded." (Zermelo 1908, p.202).

2. For a detailed discussion of the iterative conception of set, see Boolos 1971.

3. However, the predicative conception still has some force here. Proofs of set existence via Separation depend on predication.

4. To generate Cantor's paradox, suppose that there is a universal set V of all sets. By Cantor's theorem, the power set of V, call it P(V), is strictly larger than V. And we have a contradiction: P(V) contains more sets than the set of all sets. Russell was led to his paradox by reflecting on Cantor's paradox.

5. The 'lasso' figure can be found in Boolos 1971, where it is attributed to Kripke.

6. Zermelo writes: "Set theory is concerned with a *domain* B of individuals ... among which are the sets" (Zermelo 1908, p.201). And "*the domain B is not itself a set*" (<u>op</u>. <u>cit</u>., p.203).

7. See Quine 1937.

8. Forster reports that according to Quine's own account (in Quine 1987), his starting point was the simple theory of types (Forster 1992, p.20).

9. We can, for example, show in an indirect way that $z \cup \{z\}$ is a set, even though the formula '$x \in z \vee x = z$' is unstratified. See Quine 1963, pp.289-90.

10. Define the successor of a set a by $a \cup \{a\}$. Define an inductive set as one to which the empty set $\emptyset$ belongs, and which is closed under the successor operation. Define a natural number as a set which belongs to every inductive set. Then, if $\boldsymbol{\omega}$ is the set of natural numbers, we can take the Induction Principle to be: Any inductive subset of $\boldsymbol{\omega}$ coincides with $\boldsymbol{\omega}$. Suppose now that we want to show that every natural number n has the property $\boldsymbol{\varphi}$. We form the set $\{n \in \boldsymbol{\omega} | \boldsymbol{\varphi} n\}$, and show that this set is inductive. But this way of proceeding requires $\{n \in \boldsymbol{\omega} | \boldsymbol{\varphi} n\}$ to be a set - and in NF we will have in general no guarantee of its sethood when $\boldsymbol{\varphi} n$ is unstratified.

11. Quine also notes that prior to Specker's result, "Rosser and Wang had already shown that no model of NF - no interpretation of '$\in$' compatible with the axioms - could make well-orderings of both the less-to-greater relation among ordinals and that among finite cardinals, unless at the cost of not interpreting the '=' of NF as identity." (Quine 1963, p.294, fn1)

12. Roughly, a set is well-founded if it is not the first link in an endless membership chain. The set of all well-founded sets generates Mirimanoff's paradox - and NF avoids the paradox because in NF the well-founded sets do not form a set. Clearly NF admits sets that are not well-founded: V is one of them, since V is a self-member, and so has a member (namely, V) which has a member (V) which has a member (V), and so on endlessly.

13. In the Church 1974 and the Mitchell 1976 systems, there is, for example, no set of the well-founded sets, no set of the non-self-membered sets, and no set of all ordinals.

14. See Vopenka and Hajek 1972, and Vopenka 1979. Vopenka and Hayek developed their theory of semisets for a set theory without a universal set.

15. Vopenka's leading examples of semisets turn on the phenomenon of vagueness.

> "Professor Charles Darwin teaches us that there is a set D of objects and a linear ordering of this set such that the first element in that set is an ape Charlie, each non-first element is a son of the immediately preceding element, and the last element is Darwin itself. The collection A of all apes belonging to D is not a set; otherwise A would have a last element. But, as everybody knows, sons of apes are apes; thus every member of D, including Mr Darwin, would have to be an ape. Elements of D can be coded in the universe of sets, e.g. by o, {o}, {{o}}, ... etc in such a way that D itself becomes a set from the universe of sets. The class of codes of all apes (element of A) is a proper semiset." (p.33)

Vopenka cites as other examples of semisets the class of all living men (where there is no "crisp boundary" between *not yet born* and *already born*, or between *yet alive* and *already dead*) and the class of all bald men. He writes:

> "Examples of proper semisets have been known for a long time but they were held for anomalies, as e.g. the "bald man paradox". But we meet proper semisets whenever in considering a property of some objects we emphasize its intension rather than its extension." (Vopenka 1979, p.34)

16. There are other examples in NF of subclasses of sets that are not sets, including some that are finite (see Forster 1992, pp.30-1).

17. Zermelo 1908, p. 201.

18. Zermelo 1908, p.203; the emphases are Zermelo's.

19. Schröder 1890 and Cantor 1899, cited in van Heijenoort 1967, p.113.

20. Cantor 1899, p.114; the emphases are Cantor's.

21. Von Neumann 1925, p.401.

22. Op. cit. p.403.

23. It might seem natural to restrict φ to formulas that contain no class variables, only set variables. However this restriction would make the system cumbersome. For a discussion of this point, see Levy et al, in Müller 1976, pp.180-1.

24. The axiom of choice is derivable from Von Neumann's axiom IV 2, which says that A is a proper class iff there is a function with domain A and range V, the class of all sets.

25. See, for example, Levy, in Müller 1976, p.198.

26. As Levy et al write: "The existence of real mathematical objects which cannot be members of even finite classes is a rather peculiar matter... " (Müller 1976, p.201.) And Quine writes: "... the obscure classes are infinite ones, and only the infinite ones give rise to paradox. Our maxim of minimum mutilation then favors admitting all finite classes of whatever things we admit" (Quine 1963, p.51). But von Neumann's system abandons the standard pairing axiom for proper classes, so that we cannot even form the unit class of each thing we admit.

27. In von Neumann's terminology, the proper classes correspond to the II-objects that are not I-objects. And throughout his 1925 paper, von Neumann quantifies over II-objects. His arithmetic construction axioms and logical construction axioms are ways of producing II-objects, and they take this form: "There is a II-object such that ...". See van Heijenoort 1967, pp.399-400.

28. See Müller 1976, p.201.

29. This system is consistent if ZF# is consistent, where ZF# is the system of ZF with the additional axiom "There is at least one inaccessible cardinal".

30. Levy et al, in Müller 1976, p.202. This system can be modelled in ZF# (see previous footnote). Let be a fixed inaccessible cardinal. We interpret our two-tier theory as follows: a 'set' is a member of rank( ), and a "class" is a set.

And we need not stop here. We can develop a theory of classes modelled by ZF plus an axiom that asserts the existence of arbitrarily large inaccessible cardinals - a system of tiers, with the lowest occupied by sets, the next by classes, the next by superclasses, and so on. Such a theory provides a universe for category theory. See Levy et al, in Müller 1976, pp.201ff.

31. For Ackermann's system, see Ackermann 1956. Levy and Vaught have shown that one can prove the existence in Ackermann's system of the unit set of V, PV, PPV, and so on (see Levy 1959, Levy-Vaught 1961).

32. Geach and Black 1952, p.108.

33. Op. cit., p.102. In the same vein, Frege criticizes Grassmann as follows:

> "He forms classes or concepts by logical addition. He would e.g. define 'continent' as 'Europe or Asia [or Africa] or America or Australia'. But it is surely a highly arbitrary procedure to form concepts merely by assembling individuals, and one devoid of

significance for actual thinking unless the objects are held together by having characteristics in common. It is precisely these which constitute the essence of the concept". (Frege 1979, p.34)

See also Frege's introduction to Grundgesetze, in Geach and Black 1952, pp.148-50; and Russell 1919, p.12.

34. König 1905, in van Heijenoort 1967, pp.145-149. Maddy emphasizes the significance of König's discussion in Maddy 1983.

35. Van Heijenoort 1967, p.148.

36. Cantor's first number class I is the class of all finite integers, generated from 1 by the successive addition of units. Cantor's second number class II is generated by addition and by the taking of limits, with the following restriction: given any member of II, its predecessors constitute a set of cardinality $\aleph_0$. The smallest member is **ω**, the limit of the finite integers. After **ω** comes $\omega+1$, $\omega+2$, ... 2ω, $2\omega+1$, ...ω^{ω}, $\omega^{\omega}+1$,... .

37. Van Heijenoort 1967, p.148.

38. *ibid*.

39. *Op. cit*., p.149, fn.1.

40. See *op. cit*., p.148, fn.1.

41. *Op. cit*., p.149.

42. So I would disagree with a claim of Maddy's (in Maddy 1983, p.118), that König's characterization of sets as completed and classes as incomplete is an articulation of the distinction between sets and proper classes.

43. This question is also taken up in Parsons 1974.

44. Aczel's set theory admits nonwellfounded sets, replacing ZF's Axiom of Regularity by an Anti-Foundation Axiom (see Aczel 1987). But still in Aczel's theory, there is no set of all sets, no set of the non-self-membered sets, no set of the well-founded sets, and so on. Aczel's set theory can no more provide extensions for 'set', 'non-self-membered set', 'well-founded set', etc, than can ZF.

45. Zermelo 1904, p. 516.

46. Expressing a view typical of the French constructivists, Lebesgue wrote:

"I believe that we can only build solidly by granting that it is impossible to demonstrate the existence of an object without defining it" (Lebesgue 1905, in Moore 1982, p.314 -

the emphasis is Lebesgue's).
See also *op. cit*., pp.316-7; and Russell 1911, pp.32-33.

47. Recall that Levy's hyperclass system and his two-tier system can be modeled in ZF# (see footnotes 24 and 25). And Ackermann's system is in a strong sense equivalent to ZF - see Levy et al, in Müller 1976, pp.210-2.

48. We can also construct strengthened reasoning about 'true' and 'denotes'. For parallel treatments of the strengthened liar and strengthened paradoxes of denotation, see Simmons 1993 and Simmons 1994. Discussions of strengthened reasoning about 'true' can also be found in Parsons 1974a, Burge 1979, Barwise and Etchemendy 1987, and Gaifman 1988 and 1992.

49. Compare Burge's discussion of strengthened liar reasoning in Burge 1979, section II.

50. Russell's paradox is now avoided. Consider the predicate 'non-self-membered extension$_o$' where the subscript indicates level. This applies only to the predicates of L_o, and the predicate itself is a predicate of L_{o+1}, not of L_o. And so the question of its own extension - whether or not it is a self-member - does not arise.

51. Kurt Gödel 1944, in Schilpp 1944, p.149.

52. Of course, philosophical discussions of paradoxes provide exceptions to this general rule.

53. The relevant instance of the C-schema, viz.

ext_CF is in ext_CF iff ext_CF is a non-self-membered extension of a nine-word predicate in utterance U,

readily generates a contradiction.

54. Consider the schema:

x is in ext_R(4) iff x is an empty extension$_R$ of a nine-word predicate in this utterance.

If we put ext_R(4) for x, we readily obtain a contradiction: ext_R(4) has itself as a member iff it is empty.

55. For a similar treatment of the Liar paradox and paradoxes of denotation, see Simmons 1993 and 1994.

56. Kurt Gödel 1944, in Schilpp 1944, p.228.

57. *Op. cit*., p.229.

58. For simplicity, we will in general ignore all context-sensitive expressions other than `extension'. So we will always represent a predicate via its type unless it contains an occurrence of `extension'.

59. In general, suppose we have a primary tree all of whose infinite branches contain repeating nodes. Then the *pruned tree* is obtained by terminating each infinite branch at the first non-repeating node.

60. If the primary tree has any infinite branch on which a node repeats, then we prune that branch, terminating it at the first non-repeating node. We need the qualification "on which a node repeats" because we do not prune infinite branches on which no node repeats. Consider the following example. Suppose that at the Great Rock on Monday, someone says:
(F) extension of a predicate uttered at the Great Rock tomorrow.
And that is all that is said at the Great Rock on Monday. On Tuesday, just one thing is said at the Great Rock, viz., a token of the same type as (F). And so on, ad infinitum. The primary tree for (F) (and for all subsequent predicates) is composed of a single infinite branch, on which no node repeats. Intuitively, F is pathological, and its infinite tree indicates that. And so we do not wish to prune the tree, even though no node on its infinite branch repeats.

There is another way in which the notion of a pruned tree needs refinement. Suppose I write the following three predicates on the board:

(G) moon of the Earth.

(H) natural number between 5 and 7.

(I) unit extension of a predicate on the board.

It is easy to see that the primary tree for (I) has an infinite branch. But it is arguable that (I) is not pathological. For ext(G) and ext(H) are both unit extensions, and so ext(I) has at least two members. But then ext(I) is not a unit extension, and so is not a self-member. So we may conclude that (I) has a well-determined extension, with members exactly ext(G) and ext(H). (An analogous situation arises with the Liarlike sentence "Snow is white or this disjunct is false", which is arguably true in virtue of the first true disjunct.)

Let me sketch a way of accommodating the case of (H). A general treatment is beyond the scope of this paper, but I hope that what follows will indicate its direction. The determinants of (I) are (G), (H) and (I). The primary tree for (I) shows that (G) and (H) have well-determined extensions, since they do not appear on infinite branches. We now form a conjunction composed of conjuncts as follows. Take those determinants that have well-determined extensions. For each determinant x, write 'I(ext(x))' if ext(x) has the property denoted by I, and write -I(ext(x)) if ext(x) does not have the property denoted by I. In the present case, we obtain the conjunction:

I(ext(i))&I(ext(ii)).

Now determine whether this conjunction implies that ext(iii) is a self-member, or that ext(iii) is not a self-member. If the conjunction implies one or the other, then we prune the primary tree for (I) by removing all nodes of the infinite branch except for the top node. Since the conjunction implies that ext(iii) is not a self-member, we remove the infinite branch. The idea is that we eliminate (iii) as a determinant, since (i) and (ii) are all we need to establish an extension for (I). Now the pruned tree for (iii) has no infinite branches, and so (iii) is not ungrounded.

61. We might also treat a version of Russell's paradox couched in terms of tokens. We might begin: "Consider all predicate tokens of English. Some of these tokens have a non-self-membered extension. Now consider the token predicate composed of the last two words of the

previous sentence."

62. The predicate 'non-self-membered $extension_i$' is a singularity of '$extension_i$' but not of '$extension_r$'. But there are also predicates that are singularities of '$extension_r$' but not of '$extension_i$'. For example, suppose we produce a token in the reflective context representable as 'non-self-membered $extension_r$'. This will be a singularity of '$extension_r$', and it will not have an $extension_r$. We can in turn reflect on this token, identify it as pathological, and provide an extension for it via an appropriately reflective schema. And, in accordance with Minimality, since the token is not identified as a singularity of '$extension_i$', it will also have the same $extension_i$.

63. Tarski 1969, p.89. See also Tarski 1983, p.164.

64. See, for example, Kripke 1975, in Martin 1984, p.79 and footnote 34.

65. In Simmons 1993, such a theory is given for `true'.

66. Suppose there is a term of T - `extension of a predicate of T' - that applies to all predicates of T. Then the predicate `non-self-membered extension of a predicate of T' is itself a predicate of T. And the contradiction associated with Russell's paradox is easily generated.