

2018

Quantification and Paradox

Edward Ferrier

Follow this and additional works at: https://scholarworks.umass.edu/dissertations_2

 Part of the [Logic and Foundations of Mathematics Commons](#)

Recommended Citation

Ferrier, Edward, "Quantification and Paradox" (2018). *Doctoral Dissertations*. 1165.
https://scholarworks.umass.edu/dissertations_2/1165

This Open Access Dissertation is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Doctoral Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

QUANTIFICATION AND PARADOX

A Dissertation Presented

by

EDWARD FERRIER

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

February 2018

Philosophy

© Copyright by Edward Ferrier 2018

All Rights Reserved

QUANTIFICATION AND PARADOX

A Dissertation Presented

by

EDWARD FERRIER

Approved as to style and content by:

Kevin Klement, Chair

Phillip Bricker, Member

Christopher Meacham, Member

Robert Kusner, Member

Joseph Levine, Department Head
Philosophy

DEDICATION

To all my wonderful colleagues and teachers who were so patient with me.

ABSTRACT

QUANTIFICATION AND PARADOX

FEBRUARY 2018

EDWARD FERRIER

B.A., THOMAS AQUINAS COLLEGE

B.A., UNIVERSITY OF NORTH CAROLINA GREENSBORO

M.A., UNIVERSITY OF MASSACHUSETTS AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Kevin Klement

I argue that *absolutism*, the view that absolutely unrestricted quantification is possible, is to blame for both the paradoxes that arise in naive set theory and variants of these paradoxes that arise in plural logic and in semantics. The solution is *restrictivism*, the view that absolutely unrestricted quantification is not possible.

It is generally thought that absolutism is true and that restrictivism is not only false, but inexpressible. As a result, the paradoxes are blamed, not on illicit quantification, but on the “logical” conception of *set* which motivates naive set theory. The accepted solution is to replace this with the “iterative” conception of *set*.

I show that this picture is doubly mistaken. After a close examination of the paradoxes in chapters 2–3, I argue in chapters 4 and 5 that it is possible to rescue naive set theory by restricting quantification over sets and that the resulting restrictivist set

theory is expressible. In chapters 6 and 7, I argue that it is the iterative conception of *set* and the thesis of absolutism that should be rejected.

TABLE OF CONTENTS

	Page
ABSTRACT	v
 CHAPTER	
1. ABSOLUTISM AND RESTRICTIVISM	1
1.1 Introduction	1
1.2 Absolutism	7
1.2.1 Defining absolutism	7
1.2.2 Motivating absolutism	9
1.3 Non-restrictivist arguments against absolutism	11
1.4 Restrictivism	14
1.4.1 Defining restrictivism	14
1.4.2 A quick argument for restrictivism?	15
1.4.3 Indefinite extensibility	22
1.4.4 Generality relativism and limitivism	27
2. THREE SET-THEORETIC PARADOXES	29
2.1 Basic terminology and definitions	30
2.2 Cantor's paradox	33
2.2.1 Proof of Cantor's theorem	33
2.2.2 Derivations of Cantor's paradox	34
2.2.2.1 Cantor I	36
2.2.2.2 Cantor II	39
2.2.2.3 Cantor III	41
2.3 Burali-Forti's paradox	42
2.3.1 Derivations of Burali-Forti's paradox	46

2.3.1.1	Burali-Forti I	47
2.3.1.2	Burali-Forti II	49
2.3.1.3	Burali-Forti III	51
2.4	Russell's Paradox	51
2.4.1	Derivation of Russell's Paradox	52
2.5	Genuine paradoxes or mere reductios?	52
2.5.1	Cantor, Burali-Forti and Russell	56
2.5.2	Paradoxes of Logic and Paradoxes of Set Theory	60
2.6	Appendix	62
3.	RUSSELL'S SCHEMA	69
3.1	Russell's schema	69
3.1.1	Self-reproductive properties	71
3.2	Putting the paradoxes into the form of Russell's schema	73
3.2.1	Russell's schema: Russell's paradox	73
3.2.2	Russell's schema: Cantor's paradox	74
3.2.2.1	First translation	75
3.2.2.2	Second translation	76
3.2.3	Russell's schema: Burali-Forti's paradox	77
3.3	The vicious circle principle	78
3.4	Logical relations between No Function and No Set	81
3.4.1	f_1 and f_2 : No Function is equivalent to No Set	81
3.4.2	f_3 and f_4 : No Function is independent from No Set	82
3.4.3	A uniform solution	85
4.	THE LOGICAL CONCEPTION OF <i>SET</i>	86
4.1	Concepts	87
4.2	Extensions	91
4.2.1	Involvement and membership	94
4.2.2	Identity	94
4.3	The priority of concepts to extensions	96

4.3.1	Extensions as pluralities	99
4.4	Inconsistency of naive set theory	104
4.5	Limitation of Size	107
4.5.1	Failure to respect the intuitive notion of an extension	109
4.5.2	Circular explanations	113
4.6	First-order set restrictivism	117
4.6.1	A modalized variant of FC	122
4.6.2	The All-in-One principle	125
4.6.3	Explanation and Expressibility	127
5.	EXPRESSING RESTRICTIVISM	130
5.1	Semantic ascent	131
5.1.1	Williamson	132
5.1.2	Fine	134
5.2	Ambiguous assertion: any and every	141
5.3	The generality of any-statements	143
5.3.1	An alternative account	147
5.4	Expressing restrictivism	149
5.5	Concluding Remark	153
6.	THE ITERATIVE CONCEPTION OF <i>SET</i>	156
6.1	The process of set formation	160
6.1.1	Priority and Explanation	167
6.2	Iterative set theories	171
6.2.1	How Class Comprehension blocks the paradoxes	174
6.2.2	Is NBG set theory restrictivist?	176
6.2.3	How Separation blocks the paradoxes	178
6.2.4	Expressing the Iterative Conception: ZF vs. NBG	180
6.3	Motivating the set existence axioms	181
6.4	Accounts of Priority	200
6.4.1	The constructivist account of priority	201
6.4.2	The modal account of priority	203

6.4.2.1	Modal set theory	209
6.4.2.2	Understanding the modal operators	211
6.4.3	The dependence account of priority	218
6.5	Conclusion	219
7.	SETLESS VARIANTS OF THE PARADOXES	221
7.1	Russell's Paradox: Semantic Version	222
7.1.1	Derivation of the contradiction	224
7.1.2	Evaluation	224
7.2	Cantor's Paradox: Plural Version	227
7.2.1	Cantor's theorem: plural version	228
7.2.2	Derivation of the contradiction	231
7.2.3	Evaluation	232
7.3	Burali-Forti Paradox: Plural Version	235
7.3.1	Derivation of the contradiction	236
7.3.2	Evaluation	237
	BIBLIOGRAPHY	239

CHAPTER 1

ABSOLUTISM AND RESTRICTIVISM

1.1 Introduction

Quantification in natural languages is almost always restricted, either explicitly or implicitly. Suppose I am dropping off a friend (Smith) at the airport. Before driving away, I want to be sure he has all his things. So I ask, “Do you have all your things?” In doing so, I *explicitly* restrict the application of ‘all’ to Smith’s belongings. But did I really mean to ask whether Smith has *all* his things? That would be a rather odd question. It’s more likely that I meant to ask whether he has all of a certain subset of his things, for example, whether he has all of the things he brought from home or acquired during his trip, and which he intends to bring back. One way to convey this question is by means of a narrower, explicit restriction. I might ask, “Do you have all the things that you brought from home or acquired during your trip and which you intend to bring back?” This formulation, however, is cumbersome. It is also unnecessary, for I am able to convey my intended question by means of an *implicit* restriction imposed by context. In the present case, facts that are well known to Smith and myself, for example, the fact that the only belongings he *can* bring on the plane are those that he brought from home or acquired during his trip, and the fact that he only *wants* to bring some of these—we may assume, for instance, that he does not wish to bring the leftovers from last night’s dinner—make the intended restriction on ‘all’ clear. In fact, I might have communicated the same question, had I placed no explicit restrictions and let context do *all* the work, that is, had I simply asked, “Do you have everything?”

Restrictions that have been imposed can subsequently be lifted. I can lift the restriction of ‘all’ to (some of) Smith’s belongings by asking a second question, e.g., whether the travelers on Smith’s flight (Smith included) have all *their* things. And, it seems, I could go further. I could lift *all* restrictions, and speak of absolutely all things. Waxing metaphysical, I might ask whether all things are self-identical, or actual, or concrete. In asking these questions, I take myself to be quantifying over absolutely all things whatsoever.

Like quantification in natural languages, quantification in formal languages may or may not be restricted. However, it is important to distinguish between two sorts of restrictions: *absolute restrictions* and *restrictions relative to a logical type*. Consider, for example, the standard view of second-order quantification as quantification over all classes (or properties) of those things in the range of the first-order quantifiers (individuals).¹ This view imposes absolute restrictions on both first and second-order quantifiers: second-order quantifiers are absolutely restricted to classes (or properties), while first order quantifiers are absolutely restricted to individuals. This view of second-order quantification corresponds to a general picture of typed logical languages, according to which quantifiers of different logical types range over entities of different logical categories. According to this picture, every typed quantifier is absolutely restricted. Absolutely unrestricted quantification is directly forbidden by the rules of logical syntax.

These rules do not forbid relatively unrestricted quantification, i.e., quantification in which a typed quantifier ranges over absolutely all entities of the corresponding logical category. Consider the language of first-order logic, in which the first-order quantifier, \forall^1 , ranges over all the individuals in the domain, D^1 . \forall^1 is unrestricted relative to its type if absolutely every individual is a member of D^1 . Similarly, in

¹See Enderton (2015) for a presentation of the standard view. See Quine (1986, 66–68) for discussion and criticism of the standard view.

the language of second-order logic, the second-order quantifier, \forall^2 , ranges over all the classes of individuals (properties of individuals) in the domain, D^2 , which is usually represented as the powerset of D^1 . \forall^2 is unrestricted relative to its type if absolutely every class of individuals (property of individuals) is a member of D^2 . Of course, this will be the case only if absolutely every individual is a member of D^1 . In general, for any $(n + 1)$ th order logic, the $(n + 1)$ th-order quantifier, \forall^{n+1} , is unrestricted relative to its type only if every entity of type n is a member of D^n .

Restrictions in a formal language are imposed by a model, $M = \langle I, D \rangle$. I is an interpretation that assigns meanings to the non-logical vocabulary (predicate and constant letters), and D is the domain of quantification. In a *standard model*, I assigns standard, intended meanings to non-logical vocabulary and D stands for a standard, intended (and usually restricted) domain. In the standard model for Peano arithmetic, for example, D is the (restricted) domain of natural numbers and I assigns standard, intended meanings to arithmetical vocabulary: the number 0 to ‘0’, and the functions Successor, Addition and Multiplication to ‘ $s(x)$ ’, ‘ $+(x, y)$ ’ and ‘ $\times(x, y)$ ’ respectively.

If the language is associated with a standard model, lifting restrictions may require changing languages. For example, lifting the restriction to natural numbers imposed by the standard model of Peano arithmetic may require switching from the language of Peano arithmetic to the more inclusive language of real analysis, whose domain (under the standard interpretation) includes the natural numbers as a proper subset. Similarly, one might lift the restrictions to propositions and possible worlds imposed by propositional and modal logics by switching to more expansive languages. Finally,

it seems that one might lift *all* restrictions by switching to an untyped language whose quantifiers are interpreted as ranging over absolutely everything.²

As a first pass, we might say that *absolutism* is the view that it is possible to lift all restrictions in a natural or (untyped) formal language; consequently, that it is possible to speak of absolutely everything. The alternative to absolutism is *restrictivism*, which we might describe as the view that it is impossible to lift all restrictions in a natural or formal language; consequently, that it is impossible to speak of absolutely everything. Related to absolutism and restrictivism *simpliciter* are absolutist views according to which it is possible to speak of everything within some domain, or of some kind, and restrictivist views according to which it is impossible to speak of everything within some domain, or of some kind. To illustrate, we might define a *numbers absolutist* as one who holds that it is possible to speak of every number, and a *numbers restrictivist* as one who denies this. Similarly, we might define a *propositions absolutist* as one who holds that it is possible to speak of every proposition and a *possible worlds absolutist* as one who holds that it is possible to speak of every possible world. In general, we might say that an *F*-absolutist is one who holds that it is possible to speak of every *F*, whereas an *F*-restrictivist is one who denies this. *F*-absolutism is weaker than, and entailed by, absolutism *simpliciter*. *F*-restrictivism is stronger than, and entails, restrictivism *simpliciter*. (Henceforth, ‘absolutism’ and ‘restrictivism’ are to be understood as absolutism *simpliciter* and restrictivism *simpliciter*, respectively.)

Restrictivism should be distinguished from the view that there is no universal kind, such as *thing*, or *object*, or *entity*. I call this view “sortalism” since it is generally accepted by sortal theorists.³ Restrictivism neither entails, nor is entailed by

²Unless otherwise noted, I use the terms ‘thing’, ‘object’, and ‘entity’ interchangeably. Thus ‘quantification over everything’ expresses the same thought as ‘quantification over every object’ and ‘quantification over every entity’.

³See, for example, Strawson (1959), Wiggins (1980), Hirsch (1982), and the discussion by Xu (1997, 368–369).

sortalism. Suppose that restrictivism is true. One might still hold that there is an absolutely general kind. Restrictivism only requires that it is impossible to quantify over all instances of this kind, presumably because of some expressive limitation endemic to logic and/or language. Next, suppose that sortalism is true. One might still hold that restrictivism is false. The reason is that the nonexistence of a universal kind does not entail the nonexistence of an absolutely general term. To illustrate, suppose, as seems plausible, that there is no kind *grue*. Nevertheless, it doesn't follow that '*x* is *grue*' (defined as '*x* is observed before some time, *t*, and is green, or else *x* is not so observed and is blue') fails to have an extension. One can deny that *grue* is a kind and still hold that the predicate '*x* is *grue*' is satisfied by all and only the things that are *grue*. Similarly, from the assumption that no absolutely general term denotes a universal kind, it doesn't follow that there are no absolutely general terms. One can deny that *thing* is a kind and still hold that the phrase 'all things' can be used in an absolutely unrestricted sense.⁴ Since this is all that is needed to make sense of the claim that it's possible to quantify over all things, sortalism does not entail restrictivism.

At present, absolutism is widely accepted as true, while restrictivism is widely rejected as inexpressible or self-defeating. This verdict is based on the observation that the restrictivist seems to be unable to express her own view—that it is impossible to speak of absolutely everything—without speaking of absolutely everything. David Lewis (1991, 68) makes this point quite forcefully when he writes that the restrictivist, “violates his own stricture in the very act of proclaiming it!” Since expressibility is closely tied to intelligibility, many philosophers have concluded that restrictivism is simply incoherent.⁵

⁴A sortal theorist who holds that reference is only possible relative to a sortal will reject this argument. Since 'thing' does not name a sortal, 'thing' fails to have *any* determinate reference.

⁵Among these are McGee (2000), Williamson (2003) and Shapiro (2003, 467).

This dissertation consists of four projects. The first (chapters 2 and 3) is a close examination of the three infamous paradoxes of set theory: Cantor's paradox, Russell's paradox and Burali-Forti's paradox. The second (chapter 4) is an articulation of the logical conception of *set* as the extension of a concept and an argument that by restricting quantification we may be able to secure a consistent set theory that preserves this concept. The third (chapter 5) is a defense against the charge that restrictivism is inexpressible. I propose a way of expressing restrictivism based on Bertrand Russell's notion of typical ambiguity that does not require absolutely unrestricted quantification. The fourth project (chapters 6 and 7) is an argument that popular absolutist strategies for blocking these paradoxes fail for two reasons: (a) the iterative concept of *set*, which motivates the non-quantificational restrictions needed to block the paradoxes in standard set theory is unsatisfactory and (b) absolutism leads to variants of these paradoxes outside of set theory to which the absolutist has no clear solution.

The plan for the rest of this chapter is as follows. In 1.2, I present a working definition of absolutism and I briefly consider what I take to be the primary motivation for adopting absolutism. In 1.3, I discuss two non-restrictivist arguments against absolutism based on the Quine/Putnam theses of *semantic indeterminacy* and *conceptual relativism* as well as a third argument against absolutism based on a (possibly Cantorian) view I call *mysticism*. In 1.4, I present a working definition of restrictivism that distinguishes it from each of these views and I discuss an argument for restrictivism based on the principle that quantification is always over a set or class. I then explain how restrictivism differs from Michael Dummett's thesis of *indefinite extensibility*, from Timothy Williamson's *generality relativism* and from Kit Fine's *limitivism*.

1.2 Absolutism

1.2.1 Defining absolutism

Above, I described *absolutism* as the view that it is possible to speak of absolutely everything. This description has the following shortcoming: it leaves room for a pseudo-absolutist view we might call “piecemeal absolutism”. According to piecemeal absolutism, you *can* speak of everything, just not all at once. To illustrate, assume (as seems reasonable) that everything is either abstract or concrete. The piecemeal absolutist might then say that it’s possible to speak of everything insofar as it’s possible to speak of all concrete entities, on the one hand, and it’s possible to speak of all abstract entities, on the other. But she will deny that it’s possible to speak of both all concrete entities and all abstract entities *at once*. Therefore, she will not accept the meaningfulness of utterances such as, “Everything is either concrete or abstract,” in which ‘everything’ purports to range simultaneously over all abstract entities and all concrete entities.⁶

Frege’s views of language and ontology may commit him to piecemeal absolutism. In “Function and Concept” (1891) and “Concept and Object” (1892b), Frege presents a logical type theory that assigns the referents of each syntactic type of a logical language to a unique logical category. According to Frege’s theory, singular expressions have syntactic type 0 and predicate expressions have syntactic type $n \geq 1$. The referents of singular expressions are “objects” and belong to the first logical category. The referents of predicate expressions are “concepts” and belong to different logical categories, one for each syntactic type of predicate expression. Predicate expressions that apply to objects have syntactic type 1. The referents of these predicate expressions—sometimes called first-order concepts—belong to the second logical category. Referents of predicate expressions of higher types (predicates that apply to

⁶For more details, see the discussion of mysticism below.

predicates) belong to higher logical categories. In this way, each entity in the universe is assigned to a unique logical category.

Piecemeal absolutism follows from Frege's prohibition on quantification across types. Unrestricted first-order quantification is restricted to quantification over absolutely all *objects*. Unrestricted second-order quantification is restricted to quantification over absolutely all *first-order concepts* (those concepts that are instantiated by objects). Unrestricted third-order quantification is restricted to quantification over absolutely all *second-order concepts* (those concepts that are instantiated by concepts that are themselves instantiated by objects). And so on. Therefore, while it's possible to quantify over all the *objects* there are and it's possible to quantify over all the concepts of a given category, it's impossible to quantify simultaneously over absolutely everything.

A better statement of absolutism, which closes this loop-hole, is that it is possible to speak of absolutely everything *at once*. Alternatively: it is possible for a *single* quantifier to range over absolutely everything. Call such a quantifier and the resulting quantification *absolutely unrestricted*. I define absolutism as the thesis that absolutely unrestricted quantification is possible, or coherent.

Absolutism Absolutely unrestricted quantification is possible, or coherent.

While absolutism does not make a claim about the expressive resources of any particular language (in particular, it does not claim that absolutely unrestricted quantification is possible in English, or first-order logic), it does make a claim about the possible expressive resources of language in general. Call a context or formal language in which quantification is absolutely unrestricted *maximal*. We might then say that absolutism is committed to the possibility of a maximal formal language or that absolutism is committed to the possibility of a maximal context for a natural language.

Absolutism is independent from particular ontological theses (claims about what there is). It does not tell us what things there are or how many things there are. Suppose you ask whether F s are included in a maximal context. If there are F s, the answer is yes. If there aren't F s, the answer is no. In either case, the answer is fixed by ontology, not absolutism.

1.2.2 Motivating absolutism

Why be an absolutist? One reason is that the truth conditions of certain utterances seem to require absolutely unrestricted quantification. Consider, for example, the atheist's claim:

- (1) There is no God.

Intuitively, (1) is true only if, for absolutely every x , it's not the case that x is God. Expressing this truth-condition requires absolutely unrestricted quantification. Philosophical contexts provide other examples. Consider the theses:

- (2) Everything is self-identical.
- (3) Nothing is abstract.
- (4) For any things, there exists a fusion of those things.

A philosopher who endorses any of these is committed to there being no counterexamples whatsoever. Consider the nominalist, who asserts (3). Suppose that sometime after making this assertion, he is convinced by some discovery or proof that there *are* abstract objects. It would be quite odd if he were to defend his original assertion of (3) on the grounds that all the things in some *restricted* domain (e.g., all of the things studied in physics) are concrete. Surely, he did not mean to be so cautious. It was his intention to make an *absolutely* general claim.

Timothy Williamson (2003, 435–444) suggests that absolutely unrestricted quantification may also be required to express the truth-conditions of statements that employ *restricted* quantification. Consider an utterance of the sentence

(5) There are no talking donkeys,

in which the quantifier ‘there are’ is plausibly understood as restricted to (absolutely all) donkeys. The absolutist might understand the truth condition for this utterance of (5) as given by

(6) $(\forall x)(\text{Donkey}(x) \rightarrow \neg[\text{Talks}(x)])$,

in which the quantifier \forall ranges over absolutely everything. On this interpretation, (5) is true iff absolutely everything is such that: if it is a donkey, then it doesn’t talk.

The restrictivist must find some other way to understand the truth-condition for this utterance of (5). Let D be a restricted domain variable interpreted according to context and let ‘ $\forall_D x$ ’ express unrestricted quantification relative to the domain D . So long as D includes absolutely all donkeys, the restrictivist might understand the truth condition for an utterance of (5) in which ‘there are’ is restricted to (absolutely all) donkeys as given by:

(7) $(\forall_D x)(\text{Donkey}(x) \rightarrow \neg[\text{Talks}(x)])$.

On this interpretation, an utterance of (5) is true iff *everything in D* is such that: if it is a donkey, then it doesn’t talk.

Of course, if there are donkeys outside of D , (7) will not be equivalent to (5). Suppose, for example, that D is the collection of all things on this planet. If there are talking donkeys on Mars, (5) will be false, but (7) will be true. Williamson argues that the restrictivist can only *select* a domain that she can *specify*. But it seems that the only way to specify a domain containing absolutely all donkeys requires quantifying over absolutely everything. To specify such a domain, one must be able

to say something like: ‘*Absolutely everything* is such that if it is a donkey, then it is in D ’.

Whether this argument is convincing depends on the relation between *using* a restriction and *specifying* a restriction. The restrictivist might argue that she can use a restricted domain containing absolutely all donkeys without having to specify this domain.

1.3 Non-restrictivist arguments against absolutism

Two arguments have been given against absolutism on the grounds that language, or reality, or both, lack the requisite determination for there to be a fact of the matter as to whether quantification is unrestricted. First, there is *the argument from semantic indeterminacy*. The semantic indeterminist uses the Skolem-Löwenheim theorem to argue that first-order theories are blind to the distinction between infinite domains of different sizes.⁷ He goes on to infer that there is no intelligible distinction to be made between first-order quantification over an all inclusive infinite domain (one that includes absolutely everything there is) and first-order quantification over a less than all inclusive infinite domain (one that leaves some things out).⁸

Second, there is *the argument from conceptual relativism*. The conceptual relativist holds that (a) the number of objects can only be determined *within* a conceptual framework, which fixes a determinate meaning for *object* and/or *exists* and (b) there is no unique conceptual framework that truly represents reality.⁹ It follows that while there is a fact of the matter as to whether a quantifier ranges over everything relative to a conceptual framework, there cannot be a fact of the matter as to whether a

⁷The theorem shows that if a first-order theory, T , is satisfied by an infinite model, then T is satisfied by infinite models of all sizes (no matter how large or how small).

⁸Putnam (1980) presents semantic indeterminism as a lesson of the Skolem-Löwenheim theorem. See also Quine (1969, 58–62).

⁹See Carnap (1950), Putnam (1977, 489–493), Putnam (1981, 22–48) and Putnam (1987).

quantifier ranges over everything *simpliciter*. An example of Putnam’s (1987, 18–19) illustrates these points. Imagine a world containing three mereological atoms. How many objects are there in this world? Relative to a conceptual framework that counts arbitrary mereological sums, the answer is seven (or eight, if we include a null-object). Relative to a conceptual framework that counts only atoms, the answer is three. But (we are invited to infer) neither of these frameworks can claim to be the true framework that captures the correct definition of *object*.

There is a third, more exotic non-restrictivist view, opposed to absolutism. According to this view, there are some peculiar things, about which it is impossible to speak, and therefore, over which, it is impossible to quantify. I call this *mysticism*.¹⁰ Suppose that *a* is a mystical object about which it is impossible to speak (never mind the fact that we just named it!). Since, according to mysticism, it is impossible to quantify over *a*, it is impossible to quantify over absolutely everything.

Some of Georg Cantor’s views on the infinite may suggest mysticism. In “Foundations of a General Theory of Manifolds” (1883) and his famous “Letter to Dedekind” (1899), Cantor makes a distinction between *transfinite* sets—totalities containing infinitely many members which can be well-ordered and therefore counted—and *absolutely infinite* or *inconsistent* totalities, which cannot be counted and therefore do not form sets. The latter include totalities such as “the absolutely infinite totality of [ordinal] numbers” ((Cantor, 1883) translated in Hallett (1984, 42)) and “the totality of everything thinkable” (Cantor, 1899). Cantor (1899) goes on to describe an absolutely infinite totality as one, “such that the assumption that *all* its elements are together leads to a contradiction, so that it is impossible to conceive of [it] as a unity.” If (a) we understand “totalities” in this passage as akin to singular objects and supplement Cantor’s description with the premises (b) that it’s only possible to

¹⁰Shapiro (2003, 467) mentions this view as a foil to restrictivism. He does not give it a name.

conceive of a totality as a singular object (“unity”) and (c) that it’s only possible to speak about what can be conceived, then if there are any absolutely infinite totalities, it follows that it’s impossible to speak of them (because they cannot be conceived). Since there are some things (absolutely infinite totalities) about which it is impossible to speak, mysticism is true.

Assumptions (a) and (b), however, are dubious. This leaves room for both the absolutist and restrictivist to understand Cantor’s remarks in ways that are friendly to their own views and which do not suggest mysticism.

The restrictivist can take Cantor’s remarks as support for her view by rejecting (a): totalities are *not* singular objects. She might then claim that in denying that an absolutely infinite totality can be conceived of as a unity, Cantor is not claiming that we cannot conceive of some sort of singular object. Instead, he is claiming that we cannot quantify over *all* the elements of very large pluralities. We cannot conceive of these pluralities as unities because they encompass too many things to talk about all at once.¹¹

The absolutist can make room for Cantor’s remarks by denying (b): it’s possible to conceive of totalities in *two* ways, as singular objects and as mere pluralities. One might conceive of a totality as a singular object, by conceiving of the set or the mereological sum that is made up of the members of the totality. But one might also conceive of a totality as a mere plurality by conceiving of all the things “in” the plurality without conceiving of any associated singular object that these things make up. To illustrate, consider the totality of natural numbers. I might conceive of this in the first way as a singular object, e.g., as the set that contains every natural number (and nothing else). But I might also conceive of this in the second way as

¹¹Though, see Parsons (1977) for an alternative modal interpretation of Cantor, according to which, while each of the elements of an absolutely infinite totality exists, it is impossible for all of them to exist *together*.

a mere plurality, simply by conceiving of all the natural numbers. On this view, Cantor’s claim that it is impossible to conceive of absolutely infinite totalities as unities expresses the thought that such totalities cannot be conceived in the first way. It doesn’t follow that they cannot be conceived of at all. It may still be possible to conceive of all the “members” of an absolutely infinite totality as a plurality. Hence, (c) no longer implies that we cannot speak of absolutely infinite totalities at all; only that we cannot speak of them as singular objects. If we have a properly plural conception of infinite totalities, then it is perfectly consistent with (c) that we can also speak of them as pluralities, using plural terms and plural quantification.

1.4 Restrictivism

1.4.1 Defining restrictivism

The theses of semantic indeterminism, conceptual relativism and mysticism are all highly implausible. Semantic indeterminism is implausible because it amounts to the denial that there can be any “intended” interpretation of a sentence (or theory) involving quantification over an infinite domain.¹² Conceptual relativism is implausible because it denies that there is any privileged division of reality into fundamental parts. And even if this is granted, it is contentious whether it follows from the fact that there is no privileged division of reality into fundamental parts that we cannot make sense of quantification over absolutely everything. Mysticism is the most implausible of all: one wonders what sort of odd qualities an object must have to make it repel any attempts at quantification.

I will not argue these points here. I will simply assume that if quantification is restricted, it is not because any of these theses are true. So what is it that might

¹²Adding that the speaker intended one interpretation rather than another amounts to adding “more theory” and is consequently subject to the same wide range of “unintended” interpretations as the original sentence.

explain the impossibility of unrestricted quantification? I think that the best answer appeals to a fundamental mismatch between the size of reality and the expressive strength of language and logic. In other words, restrictivism is true iff and because there are too many things for language or logic to capture all at once. I will say more about this when I discuss the challenge of expressing restrictivism; officially, I define restrictivism as the view that absolutely unrestricted quantification is impossible, or incoherent, but not because semantic indeterminism, or conceptual relativism, or mysticism is true.

Restrictivism Absolutely unrestricted quantification is impossible, or incoherent, but not because semantic indeterminism, or conceptual relativism, or mysticism is true.

In what follows, I will generally omit these qualifications and speak of restrictivism simply as the view that absolutely unrestricted quantification is impossible.

1.4.2 A quick argument for restrictivism?

Quantification is typically defined as taking place over a *domain*, which is typically identified with the *set* (or class, or collection) containing all and only the objects over which the quantifiers range. The domain fixes the interpretation of quantified statements. For example, by identifying the domain of quantification with the set of natural numbers, we determine that the quantifiers range over all and only the natural numbers. Consequently, any statement of the form ‘everything ϕ s’ is interpreted as “every natural number ϕ s.” These semantic facts suggest the following quick argument for restrictivism:

- P1. A quantifier ranges over all and only the objects in the domain of quantification.
- P2. Domains of quantification are sets (or classes or collections).
- P3. There is no universal set (or class).

C. Therefore restrictivism is true.

P2 and P3 jointly entail that there is no universal domain; whence, by P1, it follows that there is no absolutely unrestricted quantification.

Bertrand Russell (1908, 63) endorses something like this argument when he equates a collection's having "no total" with the impossibility of quantification over all its members.

When I say that a collection has no total, I mean that statements about *all* its members are nonsense.

If we interpret 'collection' as 'set' and identify a collection's having no total with its nonexistence, Russell's claim is that the nonexistence of a set is equivalent to the impossibility of (meaningful) quantification over all its members. If this claim were true, it would validate the inference from the nonexistence of a universal set—P3—to the truth of restrictivism—C. But what grounds the claim? Without the auxiliary premises P1 and P2, it's hard to see any clear logical connection between the nonexistence of a set and the impossibility of quantifying over its members. This suggests an implicit acceptance of P1 and P2.

Michael Dummett (1991) seems to endorse a similar argument. He takes the lesson of the set-theoretic paradoxes to be that unrestricted quantification over certain totalities, such as the totality of all sets (all cardinals, all ordinals, etc.), is impossible. However, these paradoxes do not directly target the possibility of quantification over these totalities, but rather the existence of particular sets (the universal set, the set of all cardinals, the set of all ordinals, etc.). Thus, Dummett seems to be inferring that there can be no quantification over all sets (all cardinals, all ordinals) because there can be no set of all sets (all cardinals, all ordinals). Again, this suggests an implicit acceptance of P1 and P2.

P1 and P2 jointly express a principle that Richard Cartwright (1994) has dubbed the All-In-One principle. Cartwright offers several statements of the principle.

To quantify over certain objects is to presuppose that those objects constitute a “collection,” or a “completed collection”—some one thing of which those objects are the members (1994, 7).

The values of the variables [of a first-order language] must be in, or belong to, some one thing (1994, 7).

Any objects that can be taken to be the values of the variables of a first-order language constitute a domain (1994, 17).

An arguably equivalent principle is the ‘domain principle’, which Michael Hallett (1984, 7) attributes to Cantor. Cantor’s formulation of the principle: “Each potential infinite . . . presupposes an actual infinite,” hardly sounds related to Cartwright’s statements of the All-in-One principle above; however, Cantor’s argument for the principle: that any infinitely variable quantity presupposes the existence of a “domain,” which he describes as “a definite, actually infinite series of values,” comes closer (cited in Hallett (1984, 25)). Graham Priest (2002, 125–126) argues for an even tighter resemblance when he writes that in modern semantics, Cantor’s argument for the domain principle amounts to the claim that:

For a sentence containing a variable to have a determinate meaning, the range of the quantifiers governing the variable . . . must be a determinate totality, a definite set.

And in later writing, Priest (2013, 1269) describes the domain principle, “in the context of modern logic,” as saying that a quantified statement, “has no determinate truth value unless there is a determinate collection over which the variable ranges.” Under this reading, Cantor’s domain principle is hardly discernible from Cartwright’s All-in-One principle.

In fact, referring to this as *the* All-in-One principle may be misleading: Cartwright suggests that in reality there are number of distinct, but closely related principles, each asserting the dependence of quantification on the existence of a particular kind of set-like entity.

According to one, the values of the variables of a first-order language must constitute a set; another requires only a class, perhaps ultimate [=proper]; still another, designed to accommodate talk of all classes, requires only a hyper-class; and so on.

Call the collection of principles named in this passage AOP. Specific principles will be individuated by subscripts. The first principle is AOP_{Set} , the second is $\text{AOP}_{\text{Class}}$, the third is $\text{AOP}_{\text{Hyper-Class}}$, and so on.

AOP_{Set} The values of the variables of a first-order language constitute a set.

AOP_{Class} The values of the variables of a first-order language constitute a set or a proper class.

AOP_{Hyper-Class} The values of the variables of a first-order language constitute a set, or a proper class or a hyper-class.

Next, consider a corresponding list of theories, ranked in order of increasing ontological complexity. The first is Zermelo-Fraenkel set theory (ZF), which countenances only sets. The second is Von Neumann-Bernays-Gödel set theory (NBG), which countenances both sets and proper classes. The third is hyper-class theory, which adds hyper-classes to NBG. (Hyper-classes are to proper classes as proper classes are to sets: while proper classes can be members of hyper-classes, hyper-classes are too large to be members of anything.) The list can be continued in this way indefinitely. Each theory from the list corresponds to an AOP principle. ZF corresponds to AOP_{Set} ; NBG corresponds to $\text{AOP}_{\text{Class}}$; Hyper-class theory corresponds to $\text{AOP}_{\text{Hyper-Class}}$ and so on.

We can now adapt the general argument for restrictivism stated above as a number of separate arguments, each applicable to one of these theories. These arguments share a single form: the first premise (which combines P1 and P2 above) is a statement that the appropriate AOP principle is true; the second premise denies the existence

of a universal entity of the appropriate kind (a universal set, class or hyperclass). To illustrate, consider the adaptation of the argument to ZF, which goes as follows:

P1. AOP_{Set} is true.

P2. There is no universal set

C. Therefore, it is impossible to quantify over all sets in ZF

Is this argument sound? I will argue that P2 is plausible but P1 is not.

There are at least three arguments that support P2. First, the existence of a universal set in ZF would contradict the axiom of foundation, since it would be self-membered. Second, the existence of a universal set in ZF would allow us to derive the existence of Russell's set from the subset axiom.¹³ Third, the existence of a universal set would contradict Cantor's theorem, since all its subsets would be members. Texts on ZF set theory commonly use one (or more) of these arguments as proofs that the universal set does not exist.¹⁴ Mutatis mutandis for NBG and hyper-class theory.

If one is determined to keep the universal set (and to reject P2), one must replace standard set theory with a nonstandard alternative, such as Quine's New Foundations (1937). To meet the first and second arguments, this replacement theory must be one in which the axioms of foundation and subset do not hold. To meet the third argument, this theory must also place restrictions on the powerset operation, limiting the number of subsets each set has. However, these restrictions may defeat their intended purpose: for if we restrict the powerset operation, it is doubtful whether any set can be a true universal set.

Turning to P1, I think it should be granted that the various AOP principles stand or fall together. Hence, P1 is true only if AOP is true, i.e., only if first-order

¹³*Proof:* The Subset axiom is: $(\forall z)(\exists y)(\forall x)(x \in y \leftrightarrow x \in z \wedge \phi x)$. Replace ' ϕ ' with ' $x \notin x$ ' and instantiate ' z ' to the universal set, u , to get: $(\exists y)(\forall x)(x \in y \leftrightarrow x \in u \wedge x \notin x)$. Since everything is a member of u , this is equivalent to $(\exists y)(\forall x)(x \in y \leftrightarrow x \notin x)$.

¹⁴See, for example, Enderton (1977, 22).

quantification, in general, presupposes the existence of a domain, taken to be a set, or class, or collection of some sort of all the objects “in” the domain. There are two arguments that might be adduced to support this. And I will argue that neither of them is very convincing.

The first appeals to how we speak about quantification. As I observed at the beginning of this section, we often describe quantification as taking place over a domain. Suppose, for example, that we are discussing quantification in the context of a first-order number theory. We might explain ourselves by uttering the following statement:

(8) Only natural numbers are in the domain of quantification.

Syntactically, ‘the domain of quantification’ occurs in (8) as a singular referring expression; semantically, its value is a particular object. If, as seems plausible, singular talk about a domain is acceptable whenever first-order quantification is employed, then the way we speak suggests that there is a collection of F s whenever there is first-order quantification over F s, just as AOP requires.

This Quinean argument is based on the ontological commitments of our assertions. The best known strategy for avoiding such commitments is to offer a paraphrase. One might execute this strategy by replacing the occurrence of the singular expression ‘the domain of quantification’ in (8) with a plural expression, such as ‘the things over which \forall ranges’. This replacement allows for a plural paraphrase of (8), such as:

(8P) Only natural numbers are among the things over which \forall ranges.¹⁵

The restrictivist may counter that even if it is possible to eliminate domain talk by systematically replacing statements such as (8) with statements such as (8P), doing

¹⁵The predicate ‘among’ is treated as a logical connective in plural logics. The adequacy of this sort of plural paraphrase of singular talk has been defended in detail by George Boolos (1998, ch. 4).

so requires abandoning certain pre-theoretic commitments to AOP. Thus, while it may be possible to reject AOP by adopting a paraphrase, doing so is revisionist.

I am not persuaded. Whether revisionism in the present case is objectionable depends on the degree to which we really took ourselves to be committed to the singular talk we use. Perhaps when the AOP-theorist speaks of domains, he takes this quite seriously. But his commitment comes from his prior attachment to AOP; not simply by his use of domain talk.

Suppose I have never thought about AOP but that I often engage in singular talk about domains of quantification. A colleague explains AOP to me and points out that my previous way of speaking carries commitment to AOP. I do not find AOP plausible, so I shift over to a plural paraphrase. To what extent does the adoption of this plural paraphrase signal a change in my original meaning? Not much, I claim. The reason is simply this: I never really *meant* to be taking a stand on AOP by my use of singular vocabulary. I was simply indifferent with respect to several ways in which what I said might be true. When I later come to consider AOP, I cease to be indifferent. Now I must settle on a more precise interpretation of my original comments; one that does not require the existence of singular domains. But in doing so, I have not changed my commitments. I have simply ceased to be indifferent about a matter on which I previously took no position.¹⁶

The second argument for AOP appeals to the formal representation of quantification in model theory. In model theory, quantification is handled by a set-theoretic semantics. According to the semantics, the range of quantification is a set. Therefore, quantification is committed to sets.

¹⁶See Eklund (2005, 565), who argues for *indifferentism*, according to which speakers are often, “indifferent with respect to certain aspects of the proposition literally expressed by the sentence assertively uttered”.

My response to this argument is that while it establishes that our formal representation of quantification is committed to sets, it does not establish that quantification itself (which is being formally represented) is so committed. The use of sets to represent domains of quantification provides no more evidence for the claim that first-order quantification is committed to sets, than does the use of variable assignment functions as semantic values of first-order variables for the claim that first-order quantification is committed to variable assignment functions, or the use of sets of n -tuples as semantic values for n -place predicates for the claim that predication is committed to sets of n -tuples. It simply does not follow from the fact that sets are used to represent domains of quantification that domains of quantification *really are* sets or that quantification is committed to sets or that the notion of quantification is only coherent relative to a set. In general, we must distinguish the commitments of a theory that *represents* quantification from the commitments of quantification itself.¹⁷ Agustin Rayo (2007, 431) makes a similar point regarding individual sentences:

One should generally distinguish between the ontological commitments carried by a sentence and the semantic machinery employed by a semantic theory assigning truth-conditions to that sentence.

I conclude that the arguments for AOP are unconvincing. As a result, the arguments considered in this section fail to establish restrictivism.

1.4.3 Indefinite extensibility

The thesis of restrictivism is reminiscent of Michael Dummett's thesis that certain concepts are *indefinitely extensible*. Dummett (1963, 1981, 1991) describes a concept

¹⁷In stressing the distinction between semantics and ontology, I don't mean to imply that semantic theory, in general, has *no* relation to what exists; only that no *particular* semantic theory is definitive in this regard. We are committed to both the truth and the meaningfulness of what we say. If the sentences we assert are meaningful only if their sub-sentential expressions have sets as their semantic values, then we *are* committed to sets. But conditions of meaningfulness are difficult to establish. And we should not be too hasty, even if we currently have no set-free ways of making sense of language, to infer a commitment to sets.

C as indefinitely extensible if given a “definite totality” of objects instantiating C (i.e., a definite totality of C s) and the ability to quantify over all elements in this totality, it is possible to introduce new C s which must lie outside the given totality. To illustrate, consider Dummett’s (1981, 532–533) argument that ‘object’ is indefinitely extensible.

The notion of ‘object’ ... has to be regarded as an indefinitely extensible one. If we have succeeded in forming some definite conception of a totality of objects, then we are able to introduce into the language quantification over this totality. But, by means of such quantification, together with certain term-forming operators which yield expressions for abstract objects, we are able to form new terms for objects which do not lie within the original totality.

If the fact that some concept C is indefinitely extensible entails that it is impossible to quantify over all C s, then the existence of indefinitely extensible concepts is a truth-maker for restrictivism.

I believe that this entailment holds, though it is not immediately evident that Dummett thinks so. The problem is that it would seem that the impossibility of quantification over all C s would be of the sort that would make quantification over all C s simply unintelligible. But Dummett flatly denies this. Speaking of the ordinals, i.e., the objects falling under the indefinitely extensible concept *ordinal number*, he writes: “It does not follow that quantification over the intuitive totality of all ordinals is unintelligible” (1991, 316–317). I argue that Dummett’s views are best understood as a sort of piece-meal absolutism (and consequently do count as restrictivist). The particular version of piece-meal absolutism he endorses allows for unrestricted quantification over every instance of an indefinitely extensible concept, but only in special cases where such quantification does not require quantification over all instances “at once.” It is these special cases that Dummett has in mind when he denies that unrestricted quantification over all ordinals is unintelligible.

Dummett’s views on quantification are based on the distinction between non-indefinitely extensible and indefinitely extensible concepts (described above) and a

corresponding distinction between definite and indefinite totalities. A definite totality is akin to a domain of quantification that encompasses all the instances of an ordinary, non-indefinitely extensible concept; an indefinite totality is akin to a domain of quantification that encompasses all the instances of an indefinitely extensible concept. Dummett claims that whereas all syntactically well-formed statements purporting to quantify over a definite totality have determinate truth-values, and therefore always succeed in quantifying over the entire definite totality; syntactically well-formed statements that purport to quantify over an indefinite totality have determinate truth-values and therefore succeed in quantifying over the entire totality, only if their interpretations under suitable assignments of definite sub-totalities have a single, determinate truth-value. In this way, we can make sense of unrestricted quantification over an indefinite totality in terms of quantification over its arbitrary sub-totalities because these can be considered in isolation, without looking at all instances of the indefinitely extensible concept “at once.” Thus, Dummett seems to endorse piece-meal absolutism.

He offers ‘every ordinal has a successor’ as an example of a statement that is “true of all ordinals whatever,” because it is “true in *any* definite totality of ordinals” (1991, 316). But other statements cannot be so construed. The statement ‘there is an ordinal number of all the ordinals’ seems to be an example. For (a) this demands that we consider *all* ordinals together in a single totality, but “no definite totality comprises everything intuitively recognizable as an ordinal number.”

There are two ways in which indefinite extensibility goes beyond the core restrictivist thesis. First, it is committed to concepts and totalities and to the related notions of ‘definite concept’ and ‘definite totality’. Second, indefinite extensibility is intricately tied to the possibility of introducing new terms. This suggests an essential tie to a sort of proof-theoretic modality: a concept is indefinitely extensible if given quantification over any elements instantiating the concept it is possible to *prove* the

existence of an object instantiating the concept that is not among these. It is not entirely clear how seriously to treat this sort of modal talk. One might argue that it is primarily heuristic. If this is right, then Dummett's claim that we can introduce new *C*s that are not among any given *C*s may simply be a colorful way of saying that given any totality of *C*s there *is* another, more inclusive, totality of *C*s. This "reductionist" view of Dummett's modal talk fits nicely with some of the things he says about indefinitely extensible concepts. He writes, for example, that indefinitely extensible concepts have an "increasing sequence of extensions," (Dummett, 1991, 317) each more inclusive than those preceding. This may be taken to suggest a restrictionist picture, according to which for any interpretation of the quantified expression "all *C*s", there is another, more inclusive interpretation. On the other hand, the fact that he calls the sequence "increasing," seems to lend support to a non-reductionist, dynamic picture, according to which new extensions are always being added by some sort of process.

Philosophers after Dummett have taken the notion of indefinite extensibility to capture other special primitive modalities. For example, Linnebo (2010, 2013) argues that the indefinite extensibility of 'set' captures the special mathematical modality involved in the claim that a set is merely potential relative to its members. Alternatively, Uzquiano (2015) argues that the indefinite extensibility of 'set' captures a special linguistic modality involved in the claim that given any interpretation of the predicate 'set', it is possible to provide a more inclusive interpretation.¹⁸ Finally, it's possible to understand indefinite extensibility in an ontological way, according to which reality itself is "incomplete" or "unfinished." Joshua Spencer (2012) seems to be advocating such a view when he argues that "there are no things that are all things." His argument is a *reductio*, couched entirely in terms of ontology, not lan-

¹⁸See also Williamson (1998) and Studd (2013).

guage. From the assumption that there are some things that are all the things, a contradiction follows. Hence (90):

Although the claim that there are some things that any things whatsoever are amongst seems intuitively plausible, I believe this thesis must be rejected.

Steven Yablo (2006, 177–178) seems to have a similar view in mind when he writes that what gives rise to the set-theoretic paradoxes is not the principle of naive comprehension (according to which every property defines *the set* of all things that have that property) but the principle that for every property *P* there exist *all the things* that have that property.

In particular, it is not the case that there are some things comprising all and only the self-identical things.¹⁹

Restrictivism is distinct from all these views. Of course, it has a modal element: according to restrictivism, unrestricted quantification is *impossible*. But it does not require that concepts such as ‘object’ and ‘set’ exhibit any special proof-theoretic or mathematical modalities. Nor does it require that the corresponding predicates exhibit any special linguistic modalities of reinterpretation or that ontology is incomplete or unfinished. These views may entail restrictivism. If it is always possible to prove the existence of objects outside any given domain of quantification, or if it is always possible to extend our language so as to quantify over more things, then it is also impossible to quantify over absolutely everything and so restrictivism is true. There is a sense in which restrictivism may also follow from the thesis that ontology is incomplete or unfinished. For, if there are no things that are all the things, then it would seem to be impossible to quantify over all the things. You can’t quantify over what isn’t there! On the other hand, it would seem that on this view, the inability to quantify over all things is not due to any failing of language or logic; rather, it

¹⁹Spencer and Yablo express their views in terms of plural quantification, which I discuss in chapters 4 and 6.

is a reflection of the incompleteness of ontology. In this sense, the ontological view is not properly restrictivist. The heart of restrictivism is a picture of the world on which ontology outstrips the resources of language and logic. The restrictivist is not an ontological sceptic. In this sense, he does not deny there are some things that are all the things. What he denies that it is possible to quantify over all the things that there are.

1.4.4 Generality relativism and limitivism

In the recent literature, a variety of labels have been used for theses about the possibility/impossibility of unrestricted quantification. In a particularly influential paper, Timothy Williamson (2003) uses the labels ‘generality absolutism’ and ‘generality relativism’ for the views that it is possible/impossible to quantify over absolutely everything. In another influential paper, Kit Fine (2006b, 21) calls the same views ‘universalism’ and ‘limitivism’. As far as I can tell, Williamson’s ‘generality absolutism’ and Fine’s ‘universalism’ have the same meaning as my ‘absolutism’. However, both Williamson’s ‘generality relativism’ and Fine’s ‘limitivism’ differ in meaning from my ‘restrictivism’. Williamson and Fine treat the views expressed by their labels as logical contradictories of generality absolutism. This makes semantic indeterminism, conceptual relativism and mysticism come out as relativist views (for Williamson) and limitivist views (for Fine). None of these counts as restrictivist according to my understanding of restrictivism.

In addition, the particular limitivist view that Fine adopts involves a modal element, common to discussions of indefinite extensibility, that is absent from my conception of restrictivism. The modal limitivist holds that any use of quantification *can* be shown to be non-absolute. In fact, Fine distinguishes two limitivist positions, which he calls ‘restrictionism’ and ‘expansionism’. The restrictionist holds that quantification is non-absolute since it can always be extended by lifting restrictions

on the interpretation of the quantifier (Fine, 2006b, 35), whereas the expansionist holds that quantification is non-absolute since it can always be extended by postulating new objects (Fine, 2006b, 37–41). The difference between Fine’s restrictionism and my restrictivism is perhaps not fundamental. One might plausibly reduce the modality of restrictionism to the existence of infinitely many possible interpretations of the quantifier, each restricted relative to others. On this view, the modal claim that a restriction on the quantifier *can* be lifted is understood as the quantificational claim that there *is* a possible interpretation of the quantifier which is less restrictive. However, Fine’s preferred view is expansionism, for which the modality is a primitive mathematical one.²⁰ In this respect, Fine’s expansionism is quite different from my understanding of restrictivism.

²⁰I discuss modal views in more detail in ch. 6.

CHAPTER 2

THREE SET-THEORETIC PARADOXES

In this chapter, I discuss the three set-theoretic paradoxes. 2.1 is dedicated to basic terminology and definitions. I introduce the necessary set-theoretic vocabulary and define a number of basic operations, relations and functions on sets. In 2.2–2.4, I articulate the additional conceptual background, existential assumptions and auxiliary proofs needed to derive each of the traditional set-theoretic paradoxes of Cantor, Burali-Forti and Russell. I then distinguish several forms that these arguments may take. In 2.5, I ask whether they are better viewed as genuine paradoxes or as *reductio* proofs and I discuss answers to this question given by Cantor, Burali-Forti and Russell.

It will be helpful to keep the ultimate goal in mind: an argument from the paradoxes to restrictivism. This ‘argument from paradox’ takes the paradoxes as evidence that it is impossible to quantify over all sets. Since the impossibility of quantifying over all sets implies the impossibility of quantifying over absolutely everything, the conclusion is that restrictivism is true.¹ To formulate the argument, I define ‘set absolutism’ and ‘set restrictivism’ as:

¹Michael Glanzberg (2006) understands the paradoxes in this way:

So far, we have observed that, given any quantifier domain, it is possible to build an object which does not fall under that domain, via directions provided by familiar paradoxes. . . . Call this the *argument from paradox*. What does this argument really show? Though the issue is contentious, my starting point . . . is that the argument shows no quantifier can range over ‘absolutely everything’.

Set Absolutism Quantification over absolutely all sets is possible, or coherent.

Set Restrictivism Quantification over absolutely all sets is impossible, or incoherent.

The argument can then be presented as follows:

The argument from paradox

P1. Set absolutism leads to paradox.

C1. Therefore, set restrictivism is true.

P2. Set restrictivism entails restrictivism.

C2. Therefore, restrictivism is true.

In rejecting all non-absolutist views other than restrictivism, I am assuming that set restrictivism is the only viable alternative to set absolutism. Given this assumption, C1 follows from P1. P2 is true as it follows from the definitions of ‘set restrictivism’ and ‘restrictivism’. The action lies with P1. I will provide an (indirect) argument for P1 in chapter 6, in which I attempt to show that the best absolutist strategy for making set absolutism consistent fails. But first, we need to get a handle on the paradoxes themselves. This requires some basic set-theoretic terminology, which I introduce below. (Readers familiar with set theory may want to skip to the next section.)

2.1 Basic terminology and definitions

Letters from the end of the alphabet (x , y , z , X , Y and Z), possibly with numerical subscripts, are used as variables for sets. Letters from the beginning of the alphabet (a , b , c , A , B and C), possibly with numerical subscripts, are used

as names for sets. As a visual aid, lower-case letters will be used for sets that are elements or subsets of sets denoted by capital letters. The capital letter ‘ R ’, possibly with numerical subscripts, is used as a variable for relations and the letters ‘ f ’, ‘ g ’, ‘ h ’, possibly with numerical subscripts, are used as variables for functions. I use the same letters, with non-numerical subscripts, as names for functions. Whenever possible, I will choose subscripts that suggest the particular function that is named. Thus, ‘ f_s ’ is used to name a particular successor function, ‘ $g_=_$ ’ to name a particular identity function, and ‘ g_c ’ to name a particular choice function.

I make use of the following basic relations and operations on sets:

X is a **subset** of Y , $X \subseteq Y$, if $x \in X$ implies $x \in Y$:

$$X \subseteq Y =_{df} (\forall x)(x \in X \rightarrow x \in Y)$$

The notation ‘ $X \subset Y$ ’ (X is a *proper subset* of Y) is used for $X \subseteq Y \wedge X \neq Y$.

The **powerset** of X , $\mathcal{P}X$, is the set of all subsets of X :

$$\mathcal{P}X =_{df} \{x \mid x \subseteq X\}$$

The **union** of X , $\bigcup X$, is the set of all members of members of X :

$$\bigcup X =_{df} \{x \mid (\exists y)(y \in X \wedge x \in y)\}$$

The notation ‘ $X_1 \cup X_2 \cup \dots \cup X_n$ ’ is used for $\bigcup\{X_1, X_2, \dots, X_n\}$.

The **relative complement** of X with respect to Y , $X - Y$, is the set of all $x \in X$ such that $x \notin Y$:

$$X - Y =_{df} \{x \mid x \in X \wedge x \notin Y\}.$$

The expression ‘ $R(x, y)$ ’ is used whenever x bears the relation R to y . The *domain* of R , $\text{dom}(R)$, is the set of all x such that $(\exists y)(R(x, y))$. The *range* of R , $\text{ran}(R)$, is the set of all y such that $(\exists x)(R(x, y))$. R is a relation *on* X if $\text{dom}(R) \subseteq X$ and $\text{ran}(R) \subseteq X$. R is a relation *from* X *to* Y if $\text{dom}(R) \subseteq X$ and $\text{ran}(R) \subseteq Y$.

A relation, R , from X to Y is a *function*, written ' $f : X \rightarrow Y$ ' if (i) $R(x, y_1)$ and $R(x, y_2)$ implies $y_1 = y_2$ and (ii) $\text{dom}(R) = X$. A condition, ϕ , is *functional* on X if for all $x \in X$: $\phi(x, y_1)$ and $\phi(x, y_2)$ implies $y_1 = y_2$. Intuitively, a function is a mapping from every object in $\text{dom}(f)$ to a unique object in $\text{ran}(f)$. The expression ' $f(x) = y$ ' is used whenever f maps x to y . For any $Z \subset \text{dom}(f)$, the result of replacing $\text{dom}(f)$ with Z is a new function, $f|_Z : Z \rightarrow Y$, called *the restriction of f to Z* . A relation, R , from X to Y is a *partial function* if condition (i) holds but $\text{dom}(R) \subset X$. We say that R is *not defined* for any $x \in X - \text{dom}(R)$. An example of a partial function is the relation of identity $f_ =$ from the set $\{1, 2\}$ to the set $\{1\}$. This relation is not defined for 2, since $2 \in \{1, 2\} - \text{dom}(f_ =)$.

A function $f : X \rightarrow Y$ is a *surjection* iff $\text{ran}(f) = Y$. A function $f : X \rightarrow Y$ is an *injection* iff for all $x_1, x_2 \in X$, $f(x_1) = f(x_2)$ implies $x_1 = x_2$. Any injection, $f : X \rightarrow Y$, has an inverse, $f^{-1} : Y \rightarrow X$, where $f^{-1}(f(x)) = x$. A function $f : X \rightarrow Y$ is a *bijection* iff f is both a surjection and an injection.²

Sets are either *pure* or *impure*. To define these terms, it is convenient to use the familiar metaphor of set formation, according to which sets are formed by collecting their elements. The pure sets are those sets that include the empty set and all the sets that can be formed out of the empty set. Thus, the empty set is a pure set; the collection of the empty set (its singleton) is a pure set; all collections of these sets are pure sets, and so on. The impure sets are all those sets that can be formed out of objects that include non-sets. Thus, the collection of Socrates (his singleton) is an impure set; any collection that includes Socrates or his singleton is an impure set; any collection that includes these collections is an impure set, and so on.

²Another name for injection is 'one-one function'. Another name for bijection is 'one-to-one correspondence'.

2.2 Cantor's paradox

Cantor's paradox is a consequence of Cantor's theorem, according to which the cardinality of any set is less than the cardinality of its powerset.

Cantor's theorem For any set, X , $\text{Card}(X) < \text{Card}(\mathcal{P}X)$.

It is assumed that every set has a cardinality and that every set has a powerset.

The cardinality of a set, $\text{Card}(X)$, is a measure of its size and is defined by:

D1. $\text{Card}(X) = \text{Card}(Y) =_{df}$ there is a bijection between X and Y .

D2. $\text{Card}(X) \leq \text{Card}(Y) =_{df}$ there is an injection from X to Y .

The relation ' $<$ ' is defined on cardinalities in terms of '=' and ' \leq ':

D3. $\text{Card}(X) < \text{Card}(Y) =_{df}$ $\text{Card}(X) \leq \text{Card}(Y)$ and $\text{Card}(X) \neq \text{Card}(Y)$.

In 2.2.1, I run through a proof of Cantor's theorem. In 2.2.2, I show how to apply Cantor's theorem to several existence assumptions to derive the paradox in a number of different ways.

2.2.1 Proof of Cantor's theorem

To prove Cantor's theorem, it suffices to show that for an arbitrary set A , (i) there is an injection from A to $\mathcal{P}A$, and, (ii) there is no bijection between A and $\mathcal{P}A$.

(i) Define $f_s : A \rightarrow \mathcal{P}A$ by $f_s(x) = \{x\}$. Since distinct objects have distinct singletons, f_s is an injection.

(ii) To prove (ii), it suffices to show that no function $f : A \rightarrow \mathcal{P}A$ is a surjection. The key insight behind Cantor's proof was the discovery of a procedure (diagonalization), by means of which we can use any $f : A \rightarrow \mathcal{P}A$ to define a set w such that: $w \in \mathcal{P}A$ but $w \notin \text{ran}(f)$.³ This shows that f is not a surjection.

³The procedure is called diagonalization because w can be represented geometrically as lying along the diagonal of a matrix for f . For a clear explanation, see Klement (2010, 16–18).

Let f_a be an arbitrary function from A to $\mathcal{P}A$ and let ‘ x ’ range over A . Define w as the set of all x such that $x \notin f_a(x)$:

$$(1.1) \quad (\forall x)(x \in w \leftrightarrow x \notin f_a(x)).^4$$

Since $w \subseteq A$, $w \in \mathcal{P}A$. Assume for reductio that $w \in \text{ran}(f_a)$. It follows that there is some x —call this x_w —such that $f_a(x_w) = w$. As an instance of (1.1), we get

$$(1.2) \quad x_w \in w \leftrightarrow x_w \notin f_a(x_w).$$

Since $f_a(x_w) = w$, (1.2) is equivalent to the contradictory

$$(1.3) \quad x_w \in w \leftrightarrow x_w \notin w.$$

Discharging our reductio assumption, $w \notin \text{ran}(f_a)$. So f_a is not a surjection. Since f_a was an arbitrary function, no function from A to $\mathcal{P}A$ is a surjection.

We now apply D1–D3 to reach the conclusion. Since $f_s : A \rightarrow \mathcal{P}A$ is an injection, it follows by D2 that $\text{Card}(A) \leq \text{Card}(\mathcal{P}A)$. Since no function from A to $\mathcal{P}A$ is a surjection, it follows by D1 that $\text{Card}(A) \neq \text{Card}(\mathcal{P}A)$. Therefore, by D3, $\text{Card}(A) < \text{Card}(\mathcal{P}A)$.

2.2.2 Derivations of Cantor’s paradox

Applying Cantor’s theorem to any of the following existence assumptions generates a contradiction.

A1. There exists a set of all *things*, U .

A2. There exists a set of all *sets*, V .

A3. There exists a set of all *pure sets*, V_P .

A4. There exists a set of all *cardinal numbers*, K .

⁴To illustrate, suppose $f_a(x_1) = \{x_1\}$ and $f_a(x_2) = \{x_1, x_3\}$. Then $x_1 \notin w$ since $x_1 \in f_a(x_1)$, but $x_2 \in w$ since $x_2 \notin f_a(x_2)$.

Speaking somewhat loosely, we might refer to the family of contradictions so generated as ‘Cantor’s paradox’. But there are two reasons to be cautious.

First, this way of speaking departs somewhat from ordinary usage: it’s common for authors to refer to a particular member of this family as ‘Cantor’s paradox’. To take a somewhat random sampling, Russell (1903, 101); (1906, 138) identifies Cantor’s paradox with the derivation of the contradiction that results from A1. In their *Encyclopedia Britannica* entry on set theory, Robert Stoll and Herbert Enderton (2016) identify Cantor’s paradox with the derivation of the contradiction that results from A2. Graham Priest (2002, 128) identifies Cantor’s paradox with the derivation of the contradiction that results from A3. And Stewart Shapiro and Crispin Wright (2006, 257) identify Cantor’s paradox with the derivation of the contradiction that results from A4.

Of course, the mere fact that these authors speak in these ways hardly shows that they are engaged in a real dispute about the true identity of Cantor’s paradox. A more plausible explanation for these diverse identifications is simply that in order to present the paradox, one of A1–A4 must be selected, though any will do, and, as a result, these authors have made their particular selections either entirely arbitrarily or purely out of pedagogical preference.

But derivations of Cantor’s paradox differ in more ways than their starting assumptions. They also differ in the identity of the contradictions they generate and in the intermediary steps leading up to these contradictions. Taken together, these differences warrant a division into several argument types. I distinguish three, each of which generates contradictions from some of A1–A4, and each of which has been identified with Cantor’s paradox. I refer to these as ‘Cantor I’, ‘Cantor II’ and ‘Cantor III’. Cantor I generates contradictions from each of A1–A3. Cantor II generates contradictions from each of A1 and A2. Cantor III generates a contradiction from A4.

I will not attempt to justify my division of Cantor’s paradox into three argument types by presenting and defending particular identity criteria for argument types. Instead, I base my division on differences in the particular contradiction (or contradictions) derived or in the steps used to derive this contradiction that strike me as “significant enough.” Often, arguments from different types differ in both ways. Thus, it is a consequence of Cantor I, but not of Cantor II, that $\text{Card}(V_P) = \text{Card}(\mathcal{P}V_P)$ (which contradicts Cantor’s theorem). In addition, these paradoxes derive their contradictions in different ways. Cantor I proceeds by showing that the cardinal number of particular sets is the same as the cardinal number of their powersets; whereas Cantor II proceeds by showing something stronger: that the cardinal number of particular sets is the largest possible cardinal number. However, I do not count every difference between arguments as constituting a difference in argument type. Thus, I distinguish two *variants* (or versions) of Cantor I. The decision of when to call two arguments variants of the same type and when to call them two distinct types is somewhat arbitrary; it’s always possible to look at things with a finer microscope and see differences in type where before one had only noticed variants of the same kind.

2.2.2.1 Cantor I

Cantor I consists of a proof that the sets U , V and V_P have the same size as their powersets. This contradicts Cantor’s theorem. *Overview:* Let ‘ X ’ stand for any of the sets U , V , or V_P . The argument proceeds in three steps:

- (i) The singleton function, $f_s : X \rightarrow \mathcal{P}X$ is an injection.
- (ii) The identity function $g_{=} : \mathcal{P}X \rightarrow X$ is an injection.
- (iii) It follows from (i) and (ii) that there is a bijection between X and $\mathcal{P}X$. By D1, $\text{Card}(X) = \text{Card}(\mathcal{P}X)$. This contradicts Cantor’s theorem, which, together with D3, implies $\text{Card}(X) \neq \text{Card}(\mathcal{P}X)$.

(i) has already been established (see 2.2.1). It remains to prove (ii) and (iii).

(ii) may seem obvious, for, it is trivial that any identity function $f : X \rightarrow Y$ is an injection. (Given that $f(x) = x$, it follows immediately that if $x_1 \neq x_2$, then $f(x_1) \neq f(x_2)$.) However this is not enough to establish (ii). We also need to show that $g_{=}$ is defined for all $x \in \mathcal{P}X$. This requires showing that $\mathcal{P}X \subseteq X$. We do this by considering each possible value of ‘ X ’. $\mathcal{P}U \subseteq U$ since every subset of U is a thing and therefore a member of U . $\mathcal{P}V \subseteq V$ since every subset of V is a set and therefore a member of V . Finally, $\mathcal{P}V_P \subseteq V_P$ since every subset of V_P is a pure set and therefore a member of V_P . Notice, however, that these cases are exceptional and do not hold for smaller sets.

(iii) requires the Schröder-Bernstein theorem (SBT), which was independently discovered by Ernst Schröder and Felix Bernstein in 1898:

SBT. If there is an injection, $f : X \rightarrow Y$, and an injection, $g : Y \rightarrow X$, then there is a bijection between X and Y .

In terms of cardinalities, SBT tells us that if $\text{Card}(X) \leq \text{Card}(Y)$ and $\text{Card}(Y) \leq \text{Card}(X)$, then $\text{Card}(X) = \text{Card}(Y)$. In other words, the $<$ -relation on cardinalities behaves as we’d naturally expect.

The preceding discussion suggests formulating Cantor’s paradox I as follows. By Cantor’s theorem,

$$(2.1) \text{ Card}(X) < \text{Card}(\mathcal{P}X).$$

Consequently, by D3,

$$(2.2) \text{ Card}(X) \neq \text{Card}(\mathcal{P}X).$$

However, by SBT, there is a bijection between X and $\mathcal{P}X$. So, by D1,

$$(2.3) \text{ Card}(X) = \text{Card}(\mathcal{P}X).$$

Contradiction.

A more common, though less rigorous, presentation of Cantor's paradox I consists of a proof that the sets U , V and V_P are at least as large as their powersets, which contradicts Cantor's theorem. *Overview:* Let ' X ' stand for any of the sets U , V or V_P . As before, the fact that $\mathcal{P}X \subseteq X$ establishes that $g_{=} : \mathcal{P}X \rightarrow X$ is an injection. However, instead of using this (together with the injective function $f_s : X \rightarrow \mathcal{P}X$) to show that there is a bijection between X and $\mathcal{P}X$, this presentation of the argument infers directly that $\text{Card}(X) \not\leq \text{Card}(\mathcal{P}X)$. This contradicts Cantor's theorem, according to which $\text{Card}(X) < \text{Card}(\mathcal{P}X)$.

Shaughan Lavine (1994, 61–62) presents the paradox in this way (for $X = U$):

The class of classes can be no larger than the class of individuals, since it is contained in the class of individuals. But the class of classes is the class of all subclasses of the class of individuals, and so Cantor's diagonal argument shows it to be larger than the class of individuals.

Lavine's presentation might be formulated as follows. Since $\mathcal{P}X \subseteq X$, $g_{=} : \mathcal{P}X \rightarrow X$ is an injection. Therefore, by D1,

$$(3.1) \text{Card}(X) \geq \text{Card}(\mathcal{P}X).$$

But (3.1) implies

$$(3.2) \text{Card}(X) \not\leq \text{Card}(\mathcal{P}X),$$

which contradicts Cantor's theorem, according to which

$$(3.3) \text{Card}(X) < \text{Card}(\mathcal{P}X).$$

Given the intuitive meaning of $<$, the step from (3.1) to (3.2) seems to be straightforward and even analytic. But notice that this step doesn't follow from the definitions D1–D3, which define $<$ on cardinalities. In fact, it requires SBT in a proof such as the following:

Proof: Given (3.1), suppose $\neg(3.2)$ for reductio. Applying SBT to (3.1) and $\neg(3.2)$, it follows that $\text{Card}(X) = \text{Card}(\mathcal{P}(X))$. This is impossible, since by $\neg(3.2)$ and D3, $\text{Card}(X) \neq \text{Card}(\mathcal{P}(X))$. Therefore, $\text{Card}(X) \not\leq \text{Card}(\mathcal{P}X)$.

2.2.2.2 Cantor II

Cantor II begins by showing that X has the greatest cardinal number when X is U or V . This contradicts Cantor's theorem, according to which $\mathcal{P}X$ has a greater cardinal number. Russell (1906, 138) presents the paradox in this form for $X = U$.

The cardinal contradiction is simply this: Cantor has a proof that there is no greatest cardinal, and yet there are properties (such as ' $x = x$ ') which belong to *all* entities. Hence the cardinal number of entities having such a property must be the greatest of cardinal numbers. Hence a contradiction.⁵

Russell's version of the argument may be formulated as follows. Suppose $X = U$ and let κ be any cardinal number. It follows that there is some set, A , such that

$$(4.1) \quad \kappa = \text{Card}(A).$$

Since $A \subseteq U$, the identify function, $g_{=} : A \rightarrow U$, is an injection. By D2, $\kappa \leq \text{Card}(U)$.

Generalizing,

$$(4.2) \quad (\forall \kappa)(\kappa \leq \text{Card}(U)).$$

Therefore,

$$(4.3) \quad \text{Card}(\mathcal{P}U) \leq \text{Card}(U).$$

By SBT,

$$(4.4) \quad \text{Card}(\mathcal{P}U) \not\leq \text{Card}(U).$$

This contradicts Cantor's theorem, according to which,

$$(4.5) \quad \text{Card}(\mathcal{P}U) > \text{Card}(U).$$

If X is V , the proof that X has the greatest cardinal number is different. Suppose first that $X = V$. As before, given an arbitrary cardinal number κ , it follows that

⁵See also Russell (1903, 362) and Fraenkel et al. (1973, 7).

there is some set, A , such that $\kappa = \text{Card}(A)$. However, we cannot assume $A \subseteq V$, since A may contain non-sets as members. To get around this problem, define B as the set of all and only singletons of members of A . Since no two objects have the same singleton, the singleton function, $f_s : A \rightarrow B$, is an injection. Since no two singletons are singletons of the same object, f 's inverse, $f^{-1} : B \rightarrow A$, is an injection. By SBT, there is a bijection between A and B . Hence, by D1, $\kappa = \text{Card}(B)$. Since every member of B is a set, $B \subseteq V$. By D2, $\kappa \leq \text{Card}(V)$. The argument then proceeds as above.

Cantor II doesn't go through when $X = V_P$ because we cannot satisfactorily prove that V_P has the greatest cardinal number. To do this, we'd need to rule out the possibility that the set of individuals, $U - V$, has a greater cardinal number. The most direct way of proving this is to come up with an injective function from $U - V$ to V_P . In the standard Zermelo-Fraenkel set theory with individuals (urelemente)—ZFCU—injective functions can be defined from arbitrary sets of individuals to pure sets.

Proof sketch: Let x be an arbitrary set of individuals in ZFCU. There is a well-ordering of x , by the well-ordering theorem. This can be used to define a bijection between x and an ordinal, y , which is a pure set. (In ZFCU, the ordinals are defined as pure sets.)

However, this proof goes through only because, in ZF set theories, there can be no set with at least as many members as there are pure sets. In any ZF set theory, a collection with as many members as there are pure sets is not a set, but a proper class. And we cannot derive a contradiction from the premise that the proper class of all pure sets has the greatest cardinal number, since proper classes do not have "power classes". Of course, the rationale for this restriction in ZFCU is, in large part, to block the derivation of Cantor's paradox.

2.2.2.3 Cantor III

Cantor III derives the contradiction from A4 by using K to construct a distinct set, which is then shown to have the greatest cardinal number. This version of the paradox is based on Shapiro and Wright's (2006, 257) presentation of the argument that *cardinal number* is an indefinitely extensible concept.

Let C be a collection of cardinal numbers. Let C' be the union of the result of replacing each $\kappa \in C$ with a set of a size κ . The collection of subsets of C' is larger than any cardinal in C . So *cardinal number* is indefinitely extensible.

If we begin with the assumption that there is a set of all cardinal numbers, i.e., there is some $C = K$, this argument generates a contradiction.

Let g_c be a choice function mapping every κ to an arbitrary set of cardinality κ . It follows that

$$(5.1) \quad (\forall x)(x \in \text{ran}(g_c) \leftrightarrow (\exists \kappa)(x = g_c(\kappa))).$$

C' is then defined as $\bigcup \text{ran}(g_c)$ which is defined by:

$$(5.2) \quad (\forall x)(x \in \bigcup \text{ran}(g_c) \leftrightarrow (\exists \kappa)(x \in g_c(\kappa))).$$

$\bigcup \text{ran}(g_c)$ has the greatest cardinal. For let κ_1 be any cardinal number and take $g_c(\kappa_1)$. From (5.2), it follows that

$$(5.3) \quad g_c(\kappa_1) \subseteq \bigcup \text{ran}(g_c).$$

Since $\text{Card}(g_c(\kappa_1)) = \kappa_1$, it follows by D2 that

$$(5.4) \quad \text{Card}(\bigcup \text{ran}(g_c)) \geq \kappa_1.$$

Generalizing,

$$(5.5) \quad (\forall \kappa)(\text{Card}(\bigcup \text{ran}(g_c)) \geq \kappa), \text{ which (by SBT) implies}$$

$$(5.6) \quad (\forall \kappa)(\text{Card}(\bigcup \text{ran}(g_c)) \not\leq \kappa).$$

This contradicts Cantor's theorem, according to which

$$(5.7) \quad \text{Card}(\bigcup \text{ran}(g_c)) < \text{Card}(\mathcal{P} \bigcup \text{ran}(g_c)).$$

2.3 Burali-Forti's paradox

Burali-Forti's paradox (1897) is the result of applying certain general facts about ordinal numbers, well-ordering relations and well-ordered sets, to the set of all ordinal numbers. In this section, I introduce the conceptual machinery needed to define the ordinals as order types of well-ordering relations. I will then use this definition, in conjunction with two theorems about ordinals, to derive the paradox.

Let R be a relation defined on the set A . R is a *well-ordering* of A iff the following conditions are met:

- R is *transitive*.
- R satisfies *trichotomy*, i.e., for any $x, y \in A$ exactly one of the three alternatives:

$$(i) xRy, \quad (ii) x = y, \quad (iii) yRx,$$

holds.

- Every non-empty subset of A has a *least* element under R .

A set, X , is well-ordered iff there is a relation R that is a well-ordering of X . Since one of the simplest and most familiar well-ordering relations is the relation $<$, defined on the natural numbers, it is standard to use the symbol ' $<_j$ ' for well-ordering relations in general. I write ' $<_A$ ' to denote the well-ordering relation $<$ on the set A and ' $(A, <_A)$ ' to denote the set A when it is well-ordered by $<_A$. For any element $x \in (A, <_A)$: (i) x is *maximal* if for every $z \in (A, <_A)$, $z \leq_A x$; (ii) x is *greatest* if x is maximal and nothing else is; (iii) x is *minimal* if for every $z \in (A, <_A)$, $z \geq_A x$ and (iv) x is *least* if x is minimal and nothing else is.

There are three immediate consequences of these definitions that are worthy of note. The first is that every well-ordered set contains a *least* element. The second

is that any non-maximal element of a well-ordered set has a unique successor.⁶ The third is that well-ordered sets may contain non-minimal elements having no immediate predecessors. These are “limit points.” (In the ordered set of natural numbers $\{1 < 3 < 5, \dots, 2 < 4 < 6, \dots\}$, in which the odd numbers (from least to greatest) precede the even numbers (from least to greatest), 2 is a limit point.)

In standard set theories, the sequence of ordinal numbers is identified with a particular sequence of well-ordered sets. According to one standard identification, 0 is defined as the empty set, \emptyset , and the successor of each natural number, $s(n)$, is defined as the union of n with its singleton: $s(n) = n \cup \{n\}$.⁷ According to another standard identification, the successor of each natural number is defined as its singleton: $s(n) = \{n\}$. Identifications such as these are mathematically advantageous, however, philosophically they appear arbitrary. This arbitrariness cannot be avoided by selecting the identification that is the most useful: there are many such sequences that are useful and none of these stands out as the sequence that is most useful (Benacerraf, 1965). One way to avoid making an arbitrary identity is to hold that ordinal numbers are *sui generis*.⁸ While I am sympathetic to this approach, I will officially go no further than to define the ordinal numbers as measures (order-types) of the “length” of well-ordered sets. This leaves it open what particular entities play this role; consequently, whether the ordinals are *sui generis* or whether they are identical to some other entities.⁹

⁶*Proof:* Suppose x is a non-maximal element of the well-ordered set $(A, <_A)$. By the definition of ‘maximal’, there is some $y \in (A, <_A)$ such that $y >_A x$. Define $B \subseteq A$ as the set of all $z \in (A, <_A)$ such that $z >_A x$. Since $<_A$ is a well-ordering, B has a least element. This is the successor of x .

⁷This is the so-called von Neumann sequence: $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots$, named after John von Neumann.

⁸This point is made by Resnik (1980, 231) and criticized by Balaguer (1998, 78–80). Another way to avoid making an arbitrary identity is to adopt structuralism. This seems to be Benacerraf’s preference (Benacerraf, 1965, 290–294).

⁹Though, any argument for the latter must address the arbitrariness worry.

Every well-ordered set, $(X, <_X)$, has an ordinal number, written ‘ $\text{Ord}(X, <_X)$ ’. $\text{Ord}(X, <_X)$ is a measure of X ’s length under the well-ordering relation $<_X$. Ordinals stand to set length as cardinal numbers stand to set size. To illustrate, consider the (unordered) set:

{Aristotle, Plato, Socrates}.

If we order this by the relation, *Born Earlier Than*, which we may symbolize as $<_{BET}$, we get the well-ordered set:

{Socrates $<_{BET}$ Plato $<_{BET}$ Aristotle}.

To calculate the length of this set, we take its elements *in order*. The *first* element is Socrates. The *second* element is Plato. The *third* and *last* element is Aristotle. Thus, its length is the ordinal number 3. Note that the ordinal number of {Socrates $<_{BET}$ Plato $<_{BET}$ Aristotle} is the same as the cardinal number of {Aristotle, Plato, Socrates}. In general, for finite sets, length and size coincide. This does not hold for infinite sets.

Two ordinals are identical when the corresponding well-ordered sets are related by an order-preserving bijection, called an *isomorphism*.

D4. $f : (A, <_A) \rightarrow (B, <_B)$ is an isomorphism between $(A, <_A)$ and $(B, <_B) =_{df}$ (i) f is a bijection and (ii) f is order-preserving, i.e., $x <_A y$ implies $f(x) <_B f(y)$.

Using ‘ \cong ’ for ‘is isomorphic to’, ordinal identity is defined by:

D5. $\text{Ord}(A, <_A) = \text{Ord}(B, <_B) =_{df} (A, <_A) \cong (B, <_B)$.

To define $<$ on the ordinals, I use the notion of an *initial segment*.

D6. If $x \in (A, <_A)$, then $\{z \in A \mid z <_A x\}$ is an initial segment of $(A, <_A)$, written ‘ $\text{seg}(x)$ ’.

I then define $<$ on the ordinals by:

D7. $\text{Ord}(A, <_A) < \text{Ord}(B, <_B) =_{df}$ there is a $y \in B$ such that $(A, <_A) \cong \text{seg}(y)$.¹⁰

Following standard convention, I use lowercase Greek letters as variables for ordinals. In addition, for any ordinal, α , I let $\boldsymbol{\alpha}$ (in boldface) denote $\text{seg}(\alpha)$:

D8. $\boldsymbol{\alpha} =_{df} \text{seg}(\alpha) = \{\beta \mid \beta < \alpha\}$.

Two important theorems that can be proved from D4–D8 are:

BF1. Any set of ordinals is well-ordered by $<$.

BF2. Every ordinal, $\alpha = \text{Ord}(\boldsymbol{\alpha})$.

I provide my own proofs in the appendix to this chapter. One reason for doing this is that standard proofs of BF1 and BF2 presuppose an (arbitrary) identification of ordinal numbers with particular well-ordered sets, which I find objectionable. Another reason is to minimize the sort of intellectual dissatisfaction that may result from simply being *told* that a certain proposition can be (or has been) proved. This is of particular importance when the proposition enters into the derivation of a paradox.

In the present case, dissatisfaction can arise in two ways. First, insofar as one is unfamiliar with the proofs, one may come to doubt the validity of the inferences from D4–D8 to BF1 and BF2. As a result, one may be led to interpret the derivation of the resulting contradiction as a simple *reductio* of either (or both) of these theorems instead of as constituting a genuine paradox. Second, one might tend towards an opposite extreme and think of BF1 and BF2 as parts of the definition of ‘ordinal number’. This may cause one to blame the Burali-Forti contradiction on an incoherent concept. For example, if in addition to playing the role of measuring well-ordered sets, the ordinals are *defined* as being themselves well-ordered and each ordinal is *defined* as the ordinal of the set of all lesser ordinals—which is, in essence, what is done

¹⁰To distinguish the ordinal relation $<$ from the cardinal relation $<$, we might write the former $<_O$ and the latter $<_K$. In what follows, context should be enough to clear up any ambiguity.

in standard ZF set theory (Jech, 2003, 19)—then it is hardly surprising to discover that the assumption that there is a set of all ordinals leads to a contradiction. Give this assumption, it is a direct implication of the definition of ‘ordinal number’ that a single ordinal exists, which is the measure of two non-isomorphic sets. This result is more surprising if ordinals are simply defined as order types of well-ordered sets. One begins with the assumption that there is a set of all ordinals. One then proves that this set is well-ordered (BF1) and therefore has an ordinal number and one then proves that this same ordinal number is the ordinal number of all lesser ordinals (BF2). Whereas the former definition seems to build the Burali-Forti contradiction into the very notion of an ordinal number, the latter definition does not.

2.3.1 Derivations of Burali-Forti’s paradox

Applying BF1–BF2 to the existence assumption:

A5. There exists a set of all the ordinal numbers,

generates a contradiction. As I noted above in the discussion of Cantor’s paradox, there are several distinct argument types, each of which generates a contradiction from A5, and each of each has been identified with Burali-Forti’s paradox. All these, however, share a common core. This core consists of employing BF1–BF2 to prove (a) that the set of all ordinal numbers is well-ordered by $<$ and therefore has an ordinal number, $\Omega = \text{Ord}(O, <)$, and (b) that $\Omega = \text{Ord}(\Omega)$. BF1–BF2 play a role in derivations of Burali-Forti’s paradox similar to that which Cantor’s theorem plays in derivations of Cantor’s paradox. Just as Cantor’s theorem implies that the cardinal number of any set, $\text{Card}(X)$, is less than some cardinal number, viz., the cardinal number of X ’s powerset, BF1–BF2 imply that the ordinal number of any well-ordered set, $\text{Ord}(X, <_X)$, is less than some ordinal number, viz., the ordinal of the set formed by adding $\text{Ord}(X, <_X)$ to the set of all smaller ordinals.

(a) and (b) can be proved as follows:

- (a) Given A5, there is a set of all ordinals, O . By BF1, this set is well-ordered by $<$. By D5, it has an ordinal number, $\Omega = \text{Ord}(O, <)$.
- (b) Instantiating BF2 to Ω gives us $\Omega = \text{Ord}(\Omega)$.

At this point, the derivations diverge. To avoid confusion, I label these ‘Burali Forti I’, ‘Burali-Forti II’ and ‘Burali-Forti III’. Below, I consider each of these in turn.

2.3.1.1 Burali-Forti I

Burali-Forti I uses (a) and (b) to derive the proposition $\Omega < \Omega$, which is inconsistent with BF1. Cantor (1899, 115) presents the paradox in this way.

Since $[(O, <)]$ is a well-ordered set, there would correspond to it a number $[\Omega]$ greater than all numbers of the system $[(O, <)]$; but the number $[\Omega]$ also occurs in the system $[(O, <)]$, because this system contains *all* numbers; $[\Omega]$ would thus be greater than $[\Omega]$, which is a contradiction.¹¹

Cantor’s argument can be reconstructed as follows:

- (i) It begins by establishing (a): “Since $[(O, <)]$ is a well-ordered set, there would correspond to it a number $[\Omega]$.”
- (ii) The argument then asserts that Ω is greater than any ordinal: “[Ω is] greater than all numbers of the system $[(O, <)]$.” *Proof:* By (a) and (b), $\text{Ord}(\Omega) = \text{Ord}(O, <)$. It follows that $\Omega = (O, <)$. Since Ω is greater than all members of Ω , Ω is greater than all members of $(O, <)$.
- (iii) Finally, the argument uses the fact that Ω is an ordinal to reach the conclusion that $\Omega < \Omega$: “But the number $[\Omega]$ also occurs in the system $[(O, <)]$, because this system contains *all* numbers; $[\Omega]$ would thus be greater than $[\Omega]$, which is a contradiction.” *Proof:* Ω is greater than all ordinals (by (ii)). But Ω is an ordinal. Therefore, $\Omega < \Omega$.

¹¹Cantor assumes that 0 is not an ordinal. Thus, instead of $(O, <)$, he uses the class containing 0 as its least element followed by all the ordinals (beginning with 1). As nothing of importance hinges on this difference, I have replaced Cantor’s class with $(O, <)$ throughout.

Cantor calls $\Omega < \Omega$, “a contradiction.” While this statement doesn’t have the logical form of a contradiction, it can be proved to be impossible. Applying D7, ‘ $\Omega < \Omega$ ’ means that $(O, <)$ is isomorphic to one of its initial segments. This is easily proved to be impossible, given the definitions for ‘isomorphism’ and ‘initial segment’ (D4 and D7).¹² Alternatively, we can derive an explicit contradiction by applying BF1, as follows:

(iv) *Proof:* By BF1, every ordinal satisfies trichotomy. Since Ω is an ordinal number, Ω satisfies trichotomy. Therefore, $\Omega \not< \Omega$.

Putting this all together, Burali-Forti I might be formulated as follows:

By (a) and (b),

(6.1) Ω is greater than any $\alpha \in (O, <)$.

But Ω is an ordinal. Therefore,

(6.2) $\Omega < \Omega$.

This contradicts BF1, according to which

(6.3) $\Omega \not< \Omega$.

A slight variant of this argument begins in the same way, by arguing that Ω is greater than any $\alpha \in (O, <)$, but then proceeds to derive the contradiction $\Omega \in (O, <) \wedge \Omega \notin (O, <)$. This is how Russell (1906, 141) describes the paradox:

Burali-Forti’s contradiction may be stated, after some modification, as follows. If u is any segment of the series of ordinals in order of magnitude, the ordinal number of u is greater than any member of u , and is, in fact, the immediate successor of $u \dots$ But now consider the whole series of ordinal numbers. This is well ordered, and therefore should have an ordinal number. This must be an ordinal number, and yet must be greater than any ordinal number. Hence it both is, and is not, an ordinal number, which is a contradiction.

¹²See the appendix, Lemma 2.

Russell's description suggests the following variant of (6.1)–(6.3):

By (a) and (b),

(7.1) Ω is greater than any $\alpha \in (O, <)$.

So,

(7.2) $\Omega \notin (O, <)$.

But Ω is an ordinal and $(O, <)$ is the set of *all* ordinals. So,

(7.3) $\Omega \in (O, <)$.

Contradiction.

2.3.1.2 Burali-Forti II

Burali-Forti II derives a contradiction by using Ω to define a new ordinal, $\Omega + 1$, which is shown to be both greater than and less than Ω . This is, in fact, how Burali-Forti originally formulated the paradox.¹³ A recent presentation of this argument is given by Geoffrey Hellman (2011, 631):

Often the Burali-Forti paradox is referred to as the paradox of ‘the largest ordinal’, which goes as follows: Let $[O]$ be the class of all (von Neuman, say) ordinals. Then, since Ω represents the order-type of the well-ordering $<$ on ordinals ... Ω itself qualifies as an ordinal. But then it has a successor, $\Omega + 1$, which is an ordinal and so must occur as a member of $[O]$, by definition of the latter as the class of *all* ordinals. But then we have that $\Omega + 1 < \Omega < \Omega + 1$, a contradiction.

Burali-Forti II has been criticized on the grounds that once Ω is defined as the ordinal of *all* ordinals there cannot be any ordinals left to add to it to generate a new ordinal. The mathematicians W. H. Young and G. C. Young (quoted by Irving Copi in (Copi, 1958, 284)) voice this criticism, when they write:

¹³See Copi (1958, 281) for a presentation of Burali-Forti's version.

When you have taken in your mind a set of type Ω , you have taken everything, nothing remains to give a new element, and the whole of the reasoning is mere confusion.

The problem with this line of reasoning is not the claim that “nothing remains to give a new element.” It is rather the inference from this to the denial that $\Omega + 1$ exists. For we can prove that if $(O, <)$ exists, then it must contain an element, $\Omega + 1$, which is both less than and greater than Ω . Nevertheless, we can avoid the appearance of conjuring up an ordinal that cannot possibly exist (given our assumption that there is a set of all ordinals) by focusing on the set $\mathbf{\Omega}$. We know that when you have taken *this* set, you have *not* taken everything. Ω remains to give a new element. Therefore, the result of adding Ω to $\mathbf{\Omega}$ is a new set, $\mathbf{\Omega} \cup \{\Omega\}$, consisting of all the members of $\mathbf{\Omega}$ with Ω stuck on at the end. Burali-Forti II can be formulated as follows:

Since $\Omega \notin \mathbf{\Omega}$,

$$(8.1) \text{ Ord}(\mathbf{\Omega}) < \text{Ord}(\mathbf{\Omega} \cup \{\Omega\}).$$

By (b), $\text{Ord}(\mathbf{\Omega}) = \Omega$. Consequently, $\text{Ord}(\mathbf{\Omega} \cup \{\Omega\}) = \Omega + 1$. Substituting into (8.1),

$$(8.2) \Omega < \Omega + 1.$$

$\Omega + 1$ is an ordinal. So, $\Omega + 1 \in (O, <)$. For the reasons given above (in step (ii) of Burali-Forti I), Ω is greater than any member of $(O, <)$. Therefore,

$$(8.3) \Omega + 1 < \Omega.$$

By (8.2) and (8.3),

$$(8.4) \Omega + 1 < \Omega < \Omega + 1$$

By BF1, $<$ is transitive, so

$$(8.5) \Omega + 1 < \Omega + 1.$$

Also, by BF1, $<$ satisfies trichotomy, so

$$(8.6) \Omega + 1 \not< \Omega + 1.$$

Contradiction.

2.3.1.3 Burali-Forti III

In both Burali-Forti I and II, the argument that Ω is greater than any $\alpha \in (O, <)$ works through the premise that $\mathbf{\Omega} = (O, <)$. The inconsistency of this identity with D8, according to which $\mathbf{\Omega}$ is the sequence of all ordinals *less* than Ω , raises the worry that the derivation of a contradiction which follows is invalid: perhaps it is merely the result of an incoherent assumption. However, this is not the case. What this observation shows is that Ω , if it existed, would have to be the ordinal of two sequences of different lengths: $\mathbf{\Omega}$, i.e., the sequence of all ordinals less than Ω , and $(O, <)$, i.e., the sequence of all ordinals (including Ω). Nevertheless, the sense that these proofs go awry at this step may be assuaged by deriving the contradiction in the following alternate way, which I label ‘Burali-Forti III’.

By (a) and (b),

$$(9.1) \text{ Ord}(\mathbf{\Omega}) = \Omega = \text{Ord}(O, <).$$

By D6, $\mathbf{\Omega}$ is an initial segment of $(O, <)$. Consequently, by D7,

$$(9.2) \text{ Ord}(\mathbf{\Omega}) < \text{Ord}(O, <).$$

It follows from (9.1) and (9.2) that

$$(9.3) \Omega < \Omega.$$

This contradicts BF1, according to which

$$(9.4) \Omega \not< \Omega.$$

2.4 Russell’s Paradox

Russell’s paradox (1901) consists in deriving a contradictory conclusion from the existence assumption:

A6. There exists a set of all the non-self-membered sets.

2.4.1 Derivation of Russell's Paradox

The derivation of Russell's paradox is straightforward and requires no additional conceptual background, special existential assumptions, or auxiliary proofs. According to A6, there exists the set, R , of all non-self-membered sets:

$$(10.1) \quad R = \{x \mid x \notin x\}$$

Suppose $R \in R$. It follows that R is a non-self-membered set. So

$$(10.2) \quad R \notin R.$$

From (10.1) and (10.2) it follows that:

$$(10.3) \quad R \in R.$$

Contradiction.

2.5 Genuine paradoxes or mere reductios?

Like a paradox, a reductio proof consists in the derivation of a contradiction from a set of premises. The difference is that a reductio has a single target, a suspect assumption, which is presumed guilty until proven innocent. Of course, in nearly all cases, the suspect assumption is not the only premise upon which the derivation of the contradiction depends. There are almost always other premises. But these are not under suspicion. The reasons why may vary. Perhaps they have been previously proved. Perhaps they are taken to be self-evident. Or, perhaps they are simply widely accepted. Whatever the reason, when a contradiction results, it is clear that the target premise is the one to be rejected.

A paradox, on the other hand, has no single target. Multiple premises are presumed innocent, usually because each has been previously accepted as highly plausible in its own right. Since there is no obvious scapegoat, the discovery that these premises are inconsistent places one in a delicate situation. To avoid the contradiction, some

premise must be rejected, but with no good reason for picking one premise rather than another, any choice is arbitrary. To escape this arbitrariness, one might be led to question the argument's validity, or even consider revising the underlying logic or language.¹⁴

Are the so-called paradoxes of set theory genuine paradoxes? Or mere reductios? To decide this question, I introduce the notion of a *natural conception of set*. By a natural conception, I mean (very roughly) an independently plausible view about the nature of sets that is explanatory in the sense that it motivates the basic principles of set theory, and in particular, grounds the conditions under which sets are said to exist. Conditions for naturalness are discussed by Michael Potter (1993, 178–179) and George Boolos (1998, 14–18, 89–90). Of particular interest are Russell's (1959, 79–80) autobiographical remarks on three requisites that had guided his search for a “wholly satisfying” solution to the paradoxes. The first two conditions do not have much to do with naturalness. These are “that the contradictions should disappear” and that “the solution should leave intact as much of mathematics as possible.” The third, however, gets to the very heart of this notion. Russell writes:

The third, which is difficult to state precisely, was that the solution should, on reflection, appeal to what may be called ‘logical common sense’—i.e., that it should seem, in the end, just what one ought to have expected all along.

What Russell has in mind is further illustrated when he considers those philosophers who being contented with “logical dexterity” have not regarded this criterion “as essential”.

Professor Quine, for example, has produced systems which I admire greatly on account of their skill, but which I cannot feel to be satisfactory because they seem to be created *ad hoc* and not to be such as even the cleverest logician would have thought of if he had not known of the contradictions.

¹⁴See Moore and Garciadiego (1981, 321) for similar remarks on the difference between contradiction and paradox.

The key idea is the requirement that a satisfactory solution to the paradoxes must be *independently motivated*. It must provide an account of sets that blocks the paradoxes, without being motivated by the need to block the paradoxes. This is what prevents the account from being ad hoc and one that not even the cleverest logician would have thought of if he had not known of the contradictions.

I claim that a conception of set is the more natural the more closely it conforms to the following conditions:

- Condition 1: It answers questions about the existence of particular sets by providing a uniform and believable account of the existence conditions for all sets.
- Condition 2: It is capable of being readily understood without prior training in set theory. The conception ought to motivate the formal theory, not the other way round.
- Condition 3: It does not contain ad hoc devices or restrictions designed principally to avoid paradox.
- Condition 4: Our confidence in the consistency of the formal theory that a conception of *set* motivates does not rest solely on the fact that no contradiction has been proved in the theory, but also on an intuitive grasp of the corresponding concept of *set*.

The question of whether a contradiction in set theory constitutes a genuine paradox is then answered relative to a conception of set. In the preceding three sections, I've shown a number of ways in which contradictions follow from the assumptions A1–A5 that particular “contradiction-inducing” sets exist. I contend that these contradictions constitute genuine paradoxes relative to a natural conception of set, C , only if C is committed to set theoretic principles which:

- (a) entail the existence of the contradiction-inducing sets, and,
- (b) motivate the set-theoretic principles needed to derive the contradictions.

When confronted with a genuine paradox in set theory, it may be unclear how to go about restoring consistency. One way is to reject the underlying conception of *set*; however, given its naturalness, we should be loathe to do this. But if we wish to keep the conception, we cannot very easily reject the contradiction-inducing sets, or throw out some of the set-theoretic principles so as to block the derivations, for if the paradox is genuine, the conception of *set* is committed to both.

Of course, one might avoid the contradictions by abandoning sets altogether. But, given their widespread acceptance in mathematics and logic and their remarkable utility, very few would consider such a draconian solution. Alternatively, one might adopt a pragmatic approach, which seeks to preserve a theory of sets without any underlying conception of *set*. I find this to be highly unsatisfying. Without a conception of *set*, we are in no position to understand what it is that set theory studies, or what makes it true. In particular, we are unable to explain, in any satisfying way, why there exist only those sets that our theory says exist; hence, we cannot explain why the contradiction-inducing sets do not exist. These shortcomings of the pragmatic strategy are nicely illustrated by Hermann Weyl's (1949, 231) remark:

The attitude is frankly pragmatic; one cures the visible symptoms [of the paradoxes] but neither diagnoses nor attacks the underlying disease.

Weyl's remark underscores the sense of dissatisfaction that results from an inability to provide an explanation: the pragmatist knows she must restrict set theory in some way for the sake of consistency; but insofar as consistency is her sole motivation, we are left dissatisfied, for we have no explanation. A similar sentiment is expressed by Dummett (1991, 316), when he writes that blocking the paradoxes by simply prohibiting those entities judged to be immediately responsible for the contradictions, "is to wield the big stick, not to offer an explanation."

Of course, if our only choice is between a consistent set theory with no explanation and an inconsistent set theory or no set theory at all, then pragmatism may be the best option. For some time after the discovery of the paradoxes, it was the general consensus amongst philosophers that we faced such an unattractive choice: it seemed that there was only one natural conception of *set*—the logical conception of *set*—and that this was committed to principles satisfying (a) and (b).¹⁵ (I discuss this conception in some detail in chapter 4.) However, most philosophers today believe they have found an alternative conception of *set* which delivers both consistency and explanation. This is the *iterative conception*, which, they claim, is natural without succumbing to (a) and (b). On their view, while the contradictions constitute genuine paradoxes relative to the original, logical conception of *set*, they do not constitute genuine paradoxes relative to the iterative conception of *set*. As Gödel (1947, 518) famously observed, relative to the iterative conception of *set*, the contradictions “cause no trouble at all.” Of course, standard texts in contemporary set theory include derivations of contradictions from A1–A5; but insofar as contemporary set theory can be motivated by the iterative conception of *set*, these derivations can be understood as harmless reductio proofs. For the iterative conception does not entail the existence of these contradiction-inducing sets. My own view is that the iterative conception fails to satisfy the conditions on naturalness; in particular, it fails to explain why the contradiction-inducing sets don’t exist, but I will take this up in chapter 6.

2.5.1 Cantor, Burali-Forti and Russell

What did Cantor, Burali-Forti and Russell think of the contradictions they had discovered? Did they view these as genuine paradoxes? The very short answer is that Russell did, while Cantor and Burali-Forti did not.

¹⁵See Potter (1993) for some citations from philosophers attesting to this view.

Burali-Forti published his contradiction in 1897 as part of a reductio proof against BF1; more specifically, as part of a proof that the ordinals were not connected under $<$, i.e., there are pairs of distinct ordinals, α and β , such that neither $\alpha < \beta$ nor $\beta < \alpha$.¹⁶ However, this was inconsistent with Cantor's (1897) proof that every proper subset of the ordinals is well-ordered, and later, with Ernest Zermelo's (1904) proof that every set could be well-ordered.

Cantor also thought of his contradiction as a reductio proof, but he placed the blame on the existence assumptions A1–A4. As I noted in chapter 1, he distinguished *absolutely infinite* totalities from smaller transfinite (and finite) totalities and he argued that only the latter formed sets. Totalities such as the totality of everything (A1), the totalities of all pure and impure sets (A2–A3) and the totality of all cardinal numbers (A4) were “absolutely infinite” and therefore too big to form sets. Cantor thought of Burali-Forti's contradiction in the same way: as a reductio proof that the totality of all ordinals (A5) was absolutely infinite and therefore did not form a set. Von Neumann (1925) later encapsulated Cantor's distinction in terms of the set/proper class distinction. Von Neumann's sets correspond to Cantor's transfinite and finite totalities; his proper classes correspond to Cantor's absolutely infinite totalities. Unlike sets, proper classes cannot be members of other classes. This restriction on membership blocks the paradoxes.¹⁷ The plausibility of Cantor's position depends on the plausibility of the limitation of size doctrine, which I will discuss later.

Russell took a different line. Following Peano and Frege, he thought of sets as the logical extensions of properties, or, as he put it, propositional functions. Given the

¹⁶See Copi (1958). Russell (1903) briefly considers the related view that O was not well-ordered even though all its proper subsets were.

¹⁷Cantor's paradox is blocked since proper classes do not have powerclasses (any proper class would be a member of its powerclass, which is impossible). Burali-Forti's paradox is blocked since proper classes do not have ordinals (any proper class would be a member of its ordinal, which is impossible). Russell's paradox is blocked since there is no class of all non-self-membered proper classes (if there was, it would contain proper classes as members, which is impossible).

existence of the properties *thing*, *set*, *pure set*, *cardinal number* and *ordinal number*, the logical conception is committed to the existence of their extensions, i.e., the sets U , V , V_P , K and O . Russell's commitment to these sets and his awareness of Cantor's paradox led him to question the veracity of Cantor's theorem. This, in turn, led him to the discovery of his own contradiction, which he regarded as a genuine paradox.¹⁸

Recall that part (ii) of Cantor's theorem consists of a reductio argument that no function from A to $\mathcal{P}A$ is a surjection. Russell argued that this reasoning fails for certain large sets, such as U , V and V_P . In a 1901 draft of "Principles of Mathematics", he showed how to define surjective functions from these sets onto their powersets. Consider the case in which $A = U$.¹⁹ Russell's proposal was to define the function $f_r : U \rightarrow \mathcal{P}U$ as follows:

For all $x \in U$, $f_r(x) = \{x\}$ if x is not a set and $f_r(x) = x$ otherwise.

Since every subset of U is a member of U , and is consequently mapped to itself by f_r , $\text{ran}(f_r) = \mathcal{P}U$. So f_r is a surjection. Now consider the set, $w_r = \{x \in U \mid x \notin f_r(x)\}$, defined by Cantor's method of diagonalization on U with respect to f_r . Clearly, $w_r \subseteq U$. So, $w_r \in \text{ran}(f_r)$. But this is impossible if the proof of Cantor's theorem is sound, for, as we've seen (2.2.1), this proof purports to show that no $f : X \rightarrow \mathcal{P}X$ can be a surjection on the grounds that diagonalization on X with respect to f always defines a set $w \notin \text{ran}(f)$.

In the same draft, Russell goes on to observe that $w_r = R$, the set of all non-self-membered sets. In other words, the result of applying diagonalization to Russell's function f_r is the Russell set, R . To see this, it will help to recall Cantor's recipe for diagonalization. Given any function, $f : X \rightarrow \mathcal{P}X$, the diagonal set w is defined by:

¹⁸Russell reports the role Cantor's argument played in the discovery of his paradox in several places. Among these are: Russell (1903, 101, 362), Russell (1919, 136) and Russell (1959, 75–76).

¹⁹In the draft, Russell confines himself to the case in which $A = V$, but the same reasoning applied in this case is easily extended to U and V_P . See Coffa (1979) for discussion.

(11.1) $(\forall x)(x \in w \leftrightarrow x \notin f(x))$, where ‘ x ’ ranges over X .

To define the diagonal set w_r , replace f with f_r in (11.1). Given the definition of f_r , if x is not a set, then $x \in f_r(x)$, and so $x \notin w_r$. On the other hand, if x is a set, then $x \in f_r(x) \leftrightarrow x \in x$, and so $x \in w_r \leftrightarrow x \notin x$. Consequently, w_r ’s members are all and only the non-self-membered sets. So $w_r = R$. If we let ‘ x_s ’ be the restriction of ‘ x ’ to sets, then we can represent this by:

(11.2) $(\forall x_s)(x_s \in w_r \leftrightarrow x_s \notin x_s)$

Putting ‘ R ’ in for ‘ w_r ’ and instantiating ‘ x_s ’ to R delivers the contradictory

(11.3) $R \in R \leftrightarrow R \notin R$.

Recognizing this contradiction, Russell goes on to remark that

The procedure [using f_r to define the diagonal set R on U] is, in this case, impossible; for if we apply it to R itself, we find that R is a [member of] $f_r(R)$, and therefore not a [member of] R ; but from the definition, R should be a [member of] R .²⁰

Russell’s point seems to be that diagonalization must fail with respect to f_r . For if it succeeded, it would introduce the set R and this leads to the contradictory (11.3). Note the contrast between this case and the ordinary cases of Cantor’s theorem. In the ordinary cases, we use diagonalization on a set X with respect to f to define the set w and then infer by reductio that $w \notin \text{ran}(f)$; or equivalently, that $x_w \notin \text{dom}(f)$. This option is not available when we use diagonalization on U with respect to f_r to define the set R . In this case, we cannot infer by reductio that $R \notin \text{ran}(f_r)$; or equivalently, that $R \notin \text{dom}(f_r)$. Since $\text{dom}(f_r) = U$ and U contains absolutely everything, $R \in \text{dom}(f_r)$ if R exists at all.

Surely, Russell is correct that diagonalization is impossible here. But what explains the impossibility? Is it a problem with diagonalization, or with something else?

²⁰As noted in the previous footnote, Russell defines R on V , not U .

Later that same year, Russell seems to have changed his mind, however, and decided that the fault was not with diagonalization. In a 1901 letter, he writes

I thought that I could refute Cantor; now I see that he is irrefutable.²¹

But at this point he was unable to offer an alternative. His discovery of ‘the contradiction’—as he called it—left him bewildered.

Today, we would likely reply that Russell’s mistake was his endorsement of a conception of *set* that was committed to individual sets such as U , V and V_P . The solution, we might add, is to replace this conception with an alternative conception of *set*, which does not have these entailments. But when Russell first discovered the paradox, there was no alternative conception available. It seemed that the only conception of *set* was also inconsistent. This explains Russell’s sense of paradox. In Russell (1903, xv–xvi) he expressed his state of mind as follows:

In the case of [sets], I must confess, I have failed to perceive any concept fulfilling the conditions requisite for the notion of [*set*]. And the contradiction . . . proves that something is amiss, but what this is I have hitherto failed to discover.²²

At the risk of oversimplification, we might say that he remained in this predicament until 1908, by which time he had worked out a type-theoretic solution (originally proposed in Russell (1903)). During the period from 1901–1908, however, it seems likely that Russell viewed the contradictions of set theory as constituting genuine paradoxes.

2.5.2 Paradoxes of Logic and Paradoxes of Set Theory

We’ve seen two ways of deriving Russell’s paradox. According to the first, the contradiction follows directly from A6. According to the second, it follows from A1 and diagonalization. The direct version involves a single ideological assumption

²¹Translated in Coffa (1979, 37).

²²I have replaced ‘class’ with ‘set’ throughout.

(in addition to some elementary logic): *set-membership*. In his presentation of set theory, Michael Potter (2004, 25) notes that the derivation of the contradiction from the assumption that the Russell set exists (A6) requires *none* of the axioms of set theory.

The first thing to notice about this result is that we have proved it before stating *any* axioms for our theory.

The fact that the contradiction can be derived so directly, without appeal to additional mathematical or set theoretic machinery, bears on the degree to which it strikes us as purely *logical*. Of course, to show that the contradiction is a logical *paradox*, and not merely a *reductio*, the assumption that R exists must itself be logically well-motivated. This is a question I will take up when I discuss the logical conception of *set* in chapter 4.

Like Russell's paradox, Cantor's paradox and Burali-Forti's paradox require the existence of particular sets (U , V , or V_P in the case of Cantor's paradox; O , in the case of Burali-Forti's paradox). But these paradoxes require more. Each makes use of additional ideology, existential assumptions and auxiliary theorems that do not appear to be part of *logic*.

All derivations of Cantor's paradox depend on Cantor's theorem, which involves the functional notions of *surjection*, *injection* and *bijection*; as well as the set-theoretic notions of *cardinality* (defined by D1–D3) and *powerset*²³ These latter two notions go hand in hand with the existential assumptions that every set has a cardinal number and that every set has a powerset. All derivations of Cantor's paradox also depend on SBT (p. 30), which, in spite of its apparent obviousness, is quite difficult to prove. Cantor's paradox III requires additional principles for set formation corresponding to the familiar ZFC axioms of Union, Replacement and Choice.

²³See pp. 31–33.

All derivations of Burali-Forti's paradox involve the notions of *well-ordering*, *isomorphism* (defined by D4) and *ordinal number* (defined by D5–D7).²⁴ These notions go hand in hand with the existential assumption that every well-ordered set has an ordinal number. All derivations of Burali-Forti's paradox also involve the auxiliary theorems BF1 and BF2. As I show below in the appendix, the proofs of these theorems require additional lemmas.

I conclude that even if we accept a logical conception of sets, according to which set-membership is a logical notion and set existence is secured by logic alone, the paradoxes of Cantor and Burali-Forti are paradoxes of set theory, not logic.

2.6 Appendix

The following proofs are based largely on Jech (2003, 18-19) and Potter (2004, 180-181).

Lemma 1. **If f is an isomorphism on the well-ordered set $(A, <_A)$, then $(\forall x \in A)(f(x) \geq_A x)$**

Proof: Let f be an isomorphism on the well-ordered set $(A, <_A)$ and suppose for reductio that there exists $x \in A$ such that $f(x) <_A x$. Let x_0 be the least such x and let $f(x_0) = y$. Since $y < x_0$, $f(y) \geq y$. Therefore, $f(y) \geq f(x_0)$, which is impossible, since f is order-preserving.

Lemma 2. **No well-ordered set is isomorphic to one of its initial segments**

Proof: Suppose for reductio that f is an isomorphism on $(A, <_A)$ and that there is some $x \in A$ such that $\text{ran}(f) = \text{seg}(x)$. It follows that $f(x) <_A x$, which contradicts Lemma 1.

²⁴See pp. 42–44.

Lemma 3. For any ordinal α , $(\alpha, <)$ is a well-ordered set

Proof: Since α is an ordinal, there is a well-ordered set, $(A, <_A)$, such that $\alpha = \text{Ord}(A, <_A)$. For each $x \in A$, let $f(x) = \text{Ord}(\text{seg}(x))$. I will show that $(\alpha, <)$ is a well-ordered set by showing two things:

1. $\text{ran}(f) = \alpha$.
2. $f : (A, <_A) \rightarrow (\text{ran}(f), <)$ is an isomorphism.
 - 1a. First, I show that $y \in \text{ran}(f)$ implies $y \in \alpha$. Suppose $y \in \text{ran}(f)$. Then, for some $x \in A$, $y = f(x) = \text{Ord}(\text{seg}(x))$. Since y is the ordinal of an initial segment of $(A, <_A)$ and $\alpha = \text{Ord}(A, <_A)$, it follows by D7 that $y < \alpha$. Hence, by D8, $y \in \alpha$.
 - 1b. Next, I show that $y \in \alpha$ implies $y \in \text{ran}(f)$. Suppose $y \in \alpha$. By D8, y is an ordinal and $y < \alpha$. Hence, there is a well-ordered set, $(B, <_B)$, such that $y = \text{Ord}(B, <_B)$ and $\text{Ord}(B, <_B) < \text{Ord}(A, <_A)$. By D7, there exists some $x \in A$ such that $(B, <_B) \cong \text{seg}(x)$. Let this be a . Then $f(a) = \text{Ord}(\text{seg}(a)) = y$. So, $y \in \text{ran}(f)$.
 - 1c. It follows from 1a and 1b that $\text{ran}(f) = \alpha$.
 - 2a. First, I show that f is a bijection. f is a surjection by definition. To show that f is an injection, I need to show that for any $x_1 \neq x_2 \in A$, $f(x_1) \neq f(x_2)$. Since $(A, <_A)$ is well-ordered, either $x_1 <_A x_2$ or $x_2 <_A x_1$. By D5, if $x_1 <_A x_2$, then $\text{seg}(x_1)$ is an initial segment of $\text{seg}(x_2)$. Similarly, if $x_2 <_A x_1$, then $\text{seg}(x_2)$ is an initial segment of $\text{seg}(x_1)$. In either case, $\text{seg}(x_1) \not\cong \text{seg}(x_2)$ by Lemma 2. By D4, $\text{Ord}(\text{seg}(x_1)) \neq \text{Ord}(\text{seg}(x_2))$. Therefore, $f(x_1) \neq f(x_2)$.
 - 2b. To show that f is order-preserving, I need to show that for any $x_1 <_A x_2$, $f(x_1) < f(x_2)$. Suppose $x_1 <_A x_2$. It follows that $\text{seg}(x_1)$ is an initial segment of $\text{seg}(x_2)$. Hence, by D6, $f(x_1) < f(x_2)$.

2c. It follows from 2a and 2b that $f : (A, <_A) \rightarrow (\text{ran}(f), <)$ is an isomorphism.

By 1c and 2c, $(A, <_A) \cong (\alpha, <)$. Hence, $(\alpha, <)$ is a well-ordered set.

BF1. Any nonempty set of ordinals is well-ordered by \jmath

Proof: Let D be a nonempty set of ordinals. We need to show three things:

1. \jmath is transitive on D .
 2. \jmath satisfies trichotomy on D .
 3. Any nonempty subset $X \subseteq D$ has a least element.
- 1a. Suppose $\alpha, \beta, \gamma \in D$ and $\alpha < \beta < \gamma$. We need to show $\alpha < \gamma$. Since α, β and γ are ordinals, there are well-ordered sets, $(A, <_A)$, $(B, <_B)$ and $(C, <_C)$ such that $\alpha = \text{Ord}(A, <_A)$, $\beta = \text{Ord}(B, <_B)$ and $\gamma = \text{Ord}(C, <_C)$.
 - 1b. Since $\alpha < \beta$, it follows by D7 that there is some $y \in B$ such that $A \cong \text{seg}(y)$. Similarly, since $\beta < \gamma$, it follows by D7 that there is some $z \in C$ such that $B \cong \text{seg}(z)$.
 - 1c. Let f be an isomorphism between B and $\text{seg}(z)$. The restriction of f to $\text{seg}(y)$ is an isomorphism between $\text{seg}(y)$ and an initial segment of $\text{seg}(z)$. Since $A \cong \text{seg}(y)$, there is an isomorphism between A and an initial segment of $\text{seg}(z)$. Since $\text{seg}(z)$ is an initial segment of C , there is an isomorphism between A and an initial segment of C . Therefore, by D7, $\alpha < \gamma$.
- 2a. Suppose $\alpha, \beta \in D$. We need to show exactly one of the three alternatives:
(i) $\alpha < \beta$, (ii) $\alpha = \beta$, (iii) $\beta < \alpha$,
holds.

2b. Since α and β are ordinals, there are well-ordered sets, $(A, <_A)$ and $(B, <_B)$ such that $\alpha = \text{Ord}(A, <_A)$ and $\beta = \text{Ord}(B, <_B)$. Therefore, (i)–(iii) are equivalent, respectively, to:

$$(i^*) \quad \text{Ord}(A, <_A) < \text{Ord}(B, <_B),$$

$$(ii^*) \quad \text{Ord}(A, <_A) = \text{Ord}(B, <_B),$$

$$(iii^*) \quad \text{Ord}(B, <_B) < \text{Ord}(A, <_A).$$

Applying D5 and D7, it follows that (i^*) – (iii^*) are equivalent, respectively, to:

$$(i^{**}) \quad (\exists y \in B)((A, <_A) \cong \text{seg}(y)),$$

$$(ii^{**}) \quad (A, <_A) \cong (B, <_B),$$

$$(iii^{**}) \quad (\exists x \in A)((B, <_B) \cong \text{seg}(x)).$$

Therefore, (i)–(iii) are equivalent, respectively, to (i^{**}) – (iii^{**}) .

2c. Let $f : A_f \rightarrow B_f$ be defined as the set of all ordered pairs $(x, y) \in A \times B$ such that $\text{seg}(x) \cong \text{seg}(y)$. It follows that $A_f = \{x \in A \mid (\exists y \in B)(\text{seg}(x) \cong \text{seg}(y))\}$ and $B_f = \{y \in B \mid (\exists x \in A)(\text{seg}(x) \cong \text{seg}(y))\}$. I will show that f is an isomorphism.

First, I show that f is a bijection. For reductio, suppose f is not a bijection. Then either f is not a surjection or f is not an injection. However, f must be a surjection, since $B_f = \text{ran}(f)$. Suppose, therefore, that f is not an injection. It follows that there exist elements $x_1, x_2 \in \text{dom}(f)$ such that $x_1 \neq x_2$ and $f(x_1) = f(x_2)$. By definition of f , this means that there are two distinct initial segments of A each isomorphic to the same initial segment of B . Consequently, these initial segments of A are isomorphic to one another. This contradicts Lemma 2.

Second, I show that f is order-preserving. Suppose $x_1, x_2 \in \text{dom}(f)$ and $x_2 <_A x_1$. I need to show that $f(x_2) < f(x_1)$.

By definition of f , there is an isomorphism, h , between $\text{seg}(x_1)$ and $\text{seg}(f(x_1))$. Since $x_2 \in \text{dom}(h)$ and $\text{ran}(h) = \text{seg}(f(x_1))$, $h(x_2) <_B f(x_1)$. Furthermore, the restriction of h to $\text{seg}(x_2)$, $h|_{\text{seg}(x_2)}$, defines an isomorphism between $\text{seg}(x_2)$ and $\text{seg}(h(x_2))$. Consequently, by definition of f , $f(x_2) = h(x_2)$. Therefore, $f(x_2) <_B f(x_1)$. This establishes that f is order-preserving.

2d. If $\text{dom}(f) = A$ and $\text{ran}(f) = B$, then case (ii**) holds.

2e. If $\text{ran}(f) \neq B$, then $\text{ran}(f)$ is an initial segment of B , $\text{dom}(f) = A$ and consequently, case (i**) holds.

First, I show that $\text{ran}(f)$ is an initial segment of B . By the definition of f , $\text{ran}(f) \subseteq B$. By assumption, $\text{ran}(f) \neq B$. Consequently, $\text{ran}(f) \subset B$, and $B - \text{ran}(f) \neq \emptyset$. Since B is well-ordered, $B - \text{ran}(f)$ has a least element, y_0 . By the definition of f , if $y_1 <_B y_2$ and $y_2 \in \text{ran}(f)$, then $y_1 \in \text{ran}(f)$. Therefore, $\text{ran}(f) = \text{seg}(y_0)$.

Next, I show that $\text{dom}(f) = A$. Assume for reductio that $\text{dom}(f) \neq A$. By the definition of f , $\text{dom}(f) \subseteq A$. Consequently, $\text{dom}(f) \subset A$, and $A - \text{dom}(f) \neq \emptyset$. Since A is well-ordered, $A - \text{dom}(f)$ has a least element, x_0 , and $\text{dom}(f) = \text{seg}(x_0)$. Since f is an isomorphism between $\text{dom}(f)$ and $\text{ran}(f)$, $\text{seg}(x_0) \cong \text{seg}(y_0)$. But then, by definition of f , $x_0 \in \text{dom}(f)$. Since $x_0 \in \text{dom}(f) \wedge x_0 \notin \text{dom}(f)$, $\text{dom}(f) = A$.

To show that case (i**) holds, note that f is an isomorphism between $\text{dom}(f)$ and $\text{ran}(f)$. Since $\text{dom}(f) = (A, <_A)$ and $\text{ran}(f) = \text{seg}(y_0)$, $(A, <_A) \cong \text{seg}(y_0)$. Therefore, $(\exists y \in B)((A, <_A) \cong \text{seg}(y))$. So case (i**) holds.

2f. If $\text{dom}(f) \neq A$, then $\text{dom}(f)$ is an initial segment of A , $\text{ran}(f) = B$ and consequently, case (iii**) holds. The argument parallels 2e. First, use the assumption that $\text{dom}(f) \neq A$ together with the fact that $\text{dom}(f) \subseteq A$ to prove that $A - \text{dom}(f)$ has a least element, x_0 , and that $\text{dom}(f) = \text{seg}(x_0)$. Next, use

a reductio proof to show that $\text{ran}(f) = B$. Since f is an isomorphism between $\text{dom}(f)$ and $\text{ran}(f)$, $(B, <_B) \cong \text{seg}(x_o)$. Therefore, $(\exists x \in A)((B, <_B) \cong \text{seg}(x))$. So case (iii**) holds.

2g. It's evident that at least one of the possibilities:

(iv) $\text{dom}(f) = A$ and $\text{ran}(f) = B$, (v) $\text{ran}(f) \neq B$, (vi) $\text{dom}(f) \neq A$,

obtains. Hence at least one of (i**)–(iii**) holds. It remains to show that at most one of (i**)–(iii**) holds. I will show first that (ii**) is incompatible with either (i**) or (iii**). Then I will argue that (i*) and (iii**) are mutually exclusive.

(ii**) is incompatible with (i**), for if both (ii**) and (i**) hold, then $(B, <_B)$ is isomorphic to one of its initial segments, contrary to Lemma 2. Similarly, (ii**) is incompatible with (iii**), for if both (ii**) and (iii**) hold, then $(A, <_A)$ is isomorphic to one of its initial segments.

To show that (i**) and (iii**) are mutually exclusive, suppose for reductio that both (i**) and (iii**) hold. By (i**), there is an isomorphism, f , between $(A, <_A)$ and $\text{seg}(y)$. By (iii**), there is an isomorphism, h , between $(B, <_B)$ and $\text{seg}(x)$. The restriction of h to $\text{seg}(y)$, $h \upharpoonright_{\text{seg}(y)}$, is an isomorphism between $\text{seg}(y)$ and $\text{seg}(x_1)$, where $x_1 <_A x$. Therefore, the composition of $h \upharpoonright_{\text{seg}(y)}$ and f , $h \upharpoonright_{\text{seg}(y)} \circ f$, is an isomorphism between $(A, <_A)$ and $\text{seg}(x_1)$, which contradicts Lemma 2.

2h. Since (by 2g) exactly one of the possibilities (iv)–(vi) obtains, it follows (by 2d–2f) that exactly one of the cases (i**)–(iii**) holds. Since these are equivalent, respectively, to (i)–(iii), exactly one of the alternatives (i)–(iii) holds.

3a. Let A be an arbitrary subset of D . Take an arbitrary $\alpha \in A$. If α is not least, then $\alpha \cap A \neq \emptyset$. By Lemma 3, α is well-ordered. Since $\alpha \cap A \subseteq \alpha$, $\alpha \cap A$ has a least element. This must also be the least element of A .

BF2. For any ordinal α , $\text{Ord}(\alpha, <) = \alpha$

Proof: Since α is an ordinal, there is a well-ordered set, $(A, <_A)$, such that $\alpha = \text{Ord}(A, <_A)$. Define f on A as in Lemma 3. Since $(A, <_A) \cong (\alpha, <)$, $\text{Ord}(\alpha, <) = \text{Ord}(A, <_A) = \alpha$.

CHAPTER 3

RUSSELL'S SCHEMA

In this chapter, I discuss Russell's 1906 diagnosis of the paradoxes as instances of a single logical form, which has subsequently been called 'Russell's schema'. Russell's schema suggests two ways of blocking the paradoxes. The first, which I call No Set, seeks to block the paradoxes by denying that the contradiction-inducing sets exist. The second, which I call No Function, seeks to block the paradoxes by denying that certain functions needed to derive the contradictions can be defined on the contradiction-inducing sets. After presenting Russell's schema in 3.1, I discuss Russell's notion of a self-reproductive class (which he introduced in the same paper as his schema). I show how to extend this notion to properties and I argue that the No Set reply amounts to the claim that only non self-reproductive properties define sets. In 3.2, I show how each of the three set theoretic paradoxes can be put into the form of Russell's schema. In 3.3, I discuss Russell's 1908 vicious circle principle and I show that it can be used to motivate both No Set and No Function. In 3.4, I argue that No Set provides a more uniform solution to the paradoxes and is therefore preferable to No Function.

3.1 Russell's schema

Russell (1906, 142) wrote that all three contradictions could be described as instances of the following "generalization":

Given a property ϕ and a function f , such that, if ϕ belongs to all the members of u , $f(u)$ always exists, has the property ϕ , and is not a member of u ; then the supposition that there is a class w of all terms having the

property ϕ and that $f(w)$ exists leads to the conclusion that $f(w)$ both has and has not the property ϕ .

These remarks describe each paradox as the product of two assumptions (specified in the consequent, after the semicolon) given certain background conditions (specified in the antecedent, preceding the semicolon). The first assumption concerns the existence of a contradiction-inducing set (or class), w , defined as the set “of all terms having the property ϕ .” In other words, $w = \{x \mid \phi(x)\}$. The second assumption concerns the existence of a function, f , defined on w . The background conditions specify two conditions governing the behavior of f on arbitrary subsets $u \subseteq w$: (i) $f(u)$ “is not a member of u ,” and (ii) $f(u)$ “has the property ϕ .” Since w is the set of all ϕ s, we can express (ii) as ‘ $f(u) \in w$ ’. Substituting ‘ w ’ for ‘ u ’ in (i) and (ii) then delivers the contradictory $f(w) \notin w \wedge f(w) \in w$.

Graham Priest (1994, 27–28), (2002, 129) refers to this generalized account of the paradoxes as ‘Russell’s schema’, which he expresses as follows:

Russell’s Schema. The paradoxes result from the following two assumptions:

- A1. There exists the set, w , of all terms having the property ϕ .
- A2. There exists a function, f , defined on every set, u , of ϕ s such that:
 - (i) $f(u) \notin u$
 - (ii) $f(u) \in w$.

In fact, Priest (1994), (2002, 133–136) goes further, arguing that Russell’s schema is an instance of a more complex schema, “the inclosure schema,” which applies to both set theoretic *and* semantic paradoxes, e.g., the Liar paradox, König’s paradox, Berry’s paradox, and Richard’s paradox. While I will not discuss this more complex schema, I will adopt some of Priest’s “inclosure schema” terminology. Following Priest, I will refer to A2(i) as ‘Transcendence’ because A2(i) says that $f(u)$ falls outside of, or *transcends* u and I will refer to A2(ii) as ‘Closure’ because A2(ii) says that w is *closed*

under $f(u)$. The combination of A1 and A2 leads to a contradiction when w is put in for u . Under this substitution, Transcendence, $f(w) \notin w$, contradicts Closure, $f(w) \in w$.

Russell's schema suggests two replies to the set theoretic paradoxes:

No Set: Deny A1, that is, deny that there is a set, w , of all terms having the property ϕ .

No Function: Deny A2, that is, deny that there is a function, f , defined on every set, u , of ϕ s, satisfying Transcendence and Closure.

One of the strengths of Russell's schema is its generality. Once we have settled on one reply (either No Set or No Function), the schema tells us under what general conditions a set (or function) fails to exist. These conditions are sufficiently general to block not just the three paradoxes I've been focusing on, but *any* set-theoretic paradox of the requisite form. According to No Set, the conditions under which a set, w , of all terms having the property ϕ , fails to exist are: *whenever there is (or would be) a function, f , defined on every set of ϕ s (including w), satisfying Transcendence and Closure.* According to No Function, the conditions under which a function, f , defined on every set of ϕ s, satisfying Transcendence and Closure, fails to exist are: *whenever there is a set, w , of all terms having the property ϕ .*

3.1.1 Self-reproductive properties

No Set tells us that some properties are without a set of *all* their instances and, in this sense, fail to define a set. No Set also tells us the conditions under which a property, ϕ , fails to define a set: whenever there is (or would be) a function, f , defined on every set of ϕ s (including the set of all ϕ s), satisfying Transcendence and Closure. This observation, however, does not tell us what sort of relation f bears to ϕ . When f is present, is its existence entailed by the intrinsic nature of ϕ ? If so,

then we might be able to attribute the paradoxes to a special kind of ‘contradiction-inducing property’. On the other hand, if f ’s existence depends on something else, then it would seem that whether a property leads to a contradiction will depend on things external to the property itself.

Russell (1906, 144) attributes the contradictions to “self-reproductive processes or [sets].”¹

The above considerations point to the conclusion that the contradictions result from the fact that, according to current logical assumptions, there are what we may call *self-reproductive* processes and [sets]. That is, there are some properties such that, given any [set] of terms all having such a property, we can always define a new term also having the property in question. Hence we can never collect *all* the terms having the said property into a whole; because, whenever we hope we have them all, the collection which we have immediately proceeds to generate a new term also having the said property.

There is some unclarity in this passage regarding the proper subject of the predicate ‘self-reproductive’. Russell refers to processes and sets; but I think he is better understood as attributing the label directly to properties. We might then say that a property, ϕ , is self-reproductive, if given any arbitrary set of ϕ s, S , it’s possible to define a new $\phi \notin S$.²

The definition of this ϕ may be viewed as a process that “expands” the set S . This process satisfies Transcendence (because it involves defining some $\phi \notin S$) and Closure (because this ϕ would belong to the set of all ϕ s, if such a set existed). In light of this observation, we might read Russell as advocating a No Set solution to the paradoxes, according to which the set of all ϕ s does not exist whenever ϕ is self-reproductive. I call this ‘the doctrine of self-reproductive properties’ (SRP).

¹I replace ‘class’ with ‘set’ throughout.

²Note that there can be no direct correspondence between self-reproductive properties and self-reproductive sets; for if ϕ is a self-reproductive property, there is no corresponding set of all ϕ s. However, we apply the label ‘self-reproductive’ to arbitrary sets of ϕ s according to the rule that a set of ϕ s, S_1 , is self-reproductive if given any set of ϕ s, S_2 , it’s possible to define a new $\phi \notin S_2$.

SRP If ϕ is self-reproductive, then there is no set of all ϕ s.

Indeed, I claim that Russell’s schema is properly understood as a representation of the logical profile of self-reproductive properties—they give rise to expansive definitions or processes—and that Russell (1906) is properly interpreted as attributing the paradoxes to the (false) assumption that all properties define sets. Finally, I claim that the feature of being self-reproductive is plausibly viewed as an intrinsic feature of properties and as the feature in virtue of which a property is contradiction-inducing.

3.2 Putting the paradoxes into the form of Russell’s schema

In this section, I show how each of the three set-theoretic paradoxes can be translated into the form of Russell’s schema. The translation is easiest for Russell’s paradox, which I present first. I will then present translations of the other two paradoxes.

3.2.1 Russell’s schema: Russell’s paradox

Let ‘ x ’ range over sets. Define ϕ as the property *being a non-self-membered set* and define w as the set of all ϕ s.

$$\text{D1. } \phi(x) =_{df} x \notin x;$$

$$w =_{df} \{x \mid \phi(x)\}$$

By D1, $w = \{x \mid x \notin x\} = R$. Let u be any subset of R and let $f(u) = u$. We need to show that $f(u)$ satisfies Transcendence (i.e., $f(u) \notin u$), and Closure (i.e., $f(u) \in R$). The argument goes as follows:

From the assumption $f(u) \in u$, it follows that $f(u) \notin u$. This establishes Transcendence:

$$(1.1) \quad f(u) \notin u.$$

Because u is a set, $f(u)$ is a non-self-membered set. So, $f(u)$ satisfies Closure:

$$(1.2) \quad f(u) \in R.$$

Putting R in for u in (1.1) and (1.2) delivers the contradictory

$$(1.3) \quad f(R) \notin R \wedge f(R) \in R.$$

3.2.2 Russell's schema: Cantor's paradox

I present two translations for Cantor's paradox.

The first derives a contradiction from the assumption that any of the sets U , V or V_P exists. In this respect, it resembles Cantor I (chap. 2, pp. 36–38); however, it differs from Cantor's paradox I both in the identity of the particular contradiction that is subsequently derived and in the derivation itself. Recall that for $X = U, V$ or V_P , Cantor's paradox I shows that there is a bijection between X and $\text{Card}(X)$, and consequently that $\text{Card}(X) = \text{Card}(\mathcal{P}X)$.³ This contradicts Cantor's theorem. To get things into the form of Russell's schema, we want a contradiction of the form $f(w) \notin w \wedge f(w) \in w$. We can put Cantor's paradox into this form by defining w as the diagonal set on X with respect to an arbitrary bijection, g , between X and $\mathcal{P}X$. Define f as the inverse of g so that $g(x_w) = w$ and $f(w) = x_w$. The reductio argument from Cantor's theorem can then be applied to g to derive the contradictory proposition that $f(w)$ both is and is not a member of w .

The second translation applies to arguments such as Cantor I (chap. 2, p. 41), which are based on the existence of K .

³Alternatively, that there is an injection from $\mathcal{P}X$ to X and consequently that $\text{Card}(X) \not\leq \text{Card}\mathcal{P}X$. See chap. 2, p. 38.

3.2.2.1 First translation

Let ‘ x ’ range over X , which is interpreted as either U , V or V_P . Let $g : X \rightarrow \mathcal{P}(X)$ be an arbitrary bijection from X onto $\mathcal{P}(X)$.⁴ Define $f : \mathcal{P}(X) \rightarrow X$ as the inverse of g .

$$\text{D2. } f =_{df} g^{-1}$$

It follows by D2 that $f(g(x)) = x$. Define ϕ as the property x has when $x \notin g(x)$ and define w as the set of all ϕ s.

$$\text{D3. } \phi(x) =_{df} x \notin g(x);$$

$$w =_{df} \{x \mid \phi(x)\}$$

The argument goes as follows.

Let u be any subset of w . Since $w \subseteq X$, $u \subseteq X$. So $u \in \mathcal{P}(X)$. Since g is a bijection, there is some x, x_u , such that, $g(x_u) = u$. By D2, it follows that $f(u) = x_u$. For conditional proof, suppose $f(u) \in u$. Since $u \subseteq w$, $f(u) \in w$. By D3, $w = \{x \mid x \notin g(x)\}$. Therefore, $f(u) \notin g(f(u))$. But $g(f(u)) = g(x_u) = u$. Discharging the assumption, $f(u) \in u$ implies $f(u) \notin u$. This establishes Transcendence:

$$(2.1) \quad f(u) \notin u.$$

Since $u = g(f(u))$ and $f(u) = x_u$, (2.1) is equivalent to: $x_u \notin g(x_u)$. It follows that $x_u \in w$. This establishes Closure:

$$(2.2) \quad f(u) \in w.$$

Putting in w for u in (2.1) and (2.2) delivers the contradictory:

$$(2.3) \quad f(w) \notin w \wedge f(w) \in w.$$

⁴Recall that by application of SBT it is possible to prove that there must be a bijection $g(x)$.

3.2.2.2 Second translation

Let ‘ x ’ range over K and let ‘ u ’ range over $\mathcal{P}K$. Define ϕ as the property *being a cardinal number* and define w as the set of all ϕ s:

D4. $\phi(x) =_{df} x$ is a cardinal;

$$w =_{df} \{x \mid \phi(x)\}$$

By D4, $w = \{x \mid x \text{ is a cardinal}\} = K$. The translation follows Cantor’s paradox III quite closely.

Let g_c be a choice function on K mapping every $\kappa \in K$ to an arbitrary set of cardinality κ . Given any $u \subseteq K$, let $g_c(u)$ denote the set formed by replacing every $\kappa \in u$ with $g_c(\kappa)$. Define $f : \mathcal{P}K \rightarrow K$ as the function that maps every u to the cardinality of the powerset of the union of $g_c(u)$.⁵

D5. $f(u) =_{df} \text{Card}(\mathcal{P} \bigcup g_c(u))$

The argument goes as follows.

Take an arbitrary u . For every $\kappa \in u$, the union of the κ -sized sets selected by g_c , $\bigcup g_c(u)$, has at least κ -many members. Hence, its cardinal number is at least as great as the greatest $\kappa \in u$. In other words, for every $\kappa \in u$, $\text{Card}(\bigcup g_c(u)) \geq \kappa$. By Cantor’s theorem, $\text{Card}(\mathcal{P} \bigcup g_c(u)) > \text{Card}(\bigcup g_c(u))$. Therefore, for every $\kappa \in u$, $\text{Card}(\mathcal{P} \bigcup g_c(u)) > \kappa$. But $\text{Card}(\mathcal{P} \bigcup g_c(u)) = f(u)$, by D5. This establishes Transcendence:

$$(3.1) \quad f(u) \notin u.$$

Since $f(u)$ is a cardinal, $f(u)$ satisfies Closure:

⁵You can think of f as a step-by-step plan for constructing a cardinal number, $f(u)$, out of u . The first step is to replace every $\kappa \in u$ with a κ -sized set. This is what $g_c(u)$ does. The second step is to form a new set out of all the members of these κ -sized sets. This is what $\bigcup g_c(u)$ does. The third step is to form the powerset of $\bigcup g_c(u)$. This is what $\mathcal{P} \bigcup g_c(u)$ does. The final step is to take the cardinality of $\mathcal{P} \bigcup g_c(u)$. This is what $\text{Card}(\mathcal{P} \bigcup g_c(u))$ does.

$$(3.2) \quad f(u) \in K.$$

Putting K in for u in (3.1) and (3.2) delivers the contradictory:

$$(3.3) \quad f(K) \notin K \wedge f(K) \in K.$$

3.2.3 Russell's schema: Burali-Forti's paradox

The translation for Burali-Forti's paradox most closely resembles the structure of Burali-Forti I (chap. 2, pp. 47–49), and in particular, Russell's variant. Let ' x ' range over O and let ' u ' range over \mathcal{PO} . Define ϕ as the property *being an ordinal number* and define w as the set of all ϕ s:

$$D6. \quad \phi(x) =_{df} x \text{ is an ordinal};$$

$$w =_{df} \{x \mid \phi(x)\}$$

By D6, $w = \{x \mid x \text{ is an ordinal}\} = O$.

Let g be the function on \mathcal{PO} that maps every u to the initial segment of ordinals defined by u . (By 'the initial segment of ordinals defined by u ', I mean the set of all ordinals less than or equal to every $x \in u$.) Define $f : \mathcal{PO} \rightarrow O$ as the function that maps every u to the ordinal number of $g(u)$.

$$D7. \quad f(u) =_{df} \text{Ord}(g(u))$$

The argument goes as follows.

By BF2, $f(u)$ is greater than any $x \in u$. This establishes Transcendence:

$$(4.1) \quad f(u) \notin u.$$

Since $f(u)$ is an ordinal, $f(u)$ satisfies Closure:

$$(4.2) \quad f(u) \in O.$$

Putting O in for u in (4.1) and (4.2) delivers the contradictory:

$$(4.3) \quad f(O) \notin O \wedge f(O) \in O.$$

3.3 The vicious circle principle

Two years after introducing his schema, Russell (1908, 63) argued that A1 and A2 violate a general principle of logic, which he expressed as follows:

Whatever involves *all* of a collection must not be one of the collection; or conversely: If, provided a certain collection had a total, it would have members only definable in terms of that total, then the said collection has no total.

Just what Russell means by *involvement in a collection* is not perfectly clear. Under one plausible interpretation, something is involved in a collection iff it is defined by quantification over that collection. This is supported by Russell’s “converse” statement of the principle, which makes explicit reference to definition. If this interpretation is correct, then Russell’s proposal is to ban all impredicative definitions. (An impredicative definition of a term is a definition that involves quantification over any collection to which the term belongs.) Such a ban, which may be motivated by the thought that impredicative definitions are viciously circular, is commonly known as *the vicious circle principle*.

Let C be any collection, τ any term, and $\text{Def}(\tau)$ a definition of τ . (I make no distinction here between collections and sets.) The vicious circle principle (VCP) may then be formulated as:

VCP $\text{Def}(\tau)$ is viciously circular (and therefore, illegitimate) iff (i) $\text{Def}(\tau)$ involves quantification over C and (ii) τ belongs to C .

VCP motivates a No Set response to Russell’s paradox and the first translation of Cantor’s paradox. $\text{Def}(R)$ violates VCP since (i) it involves quantification over the collection of all non-self-membered sets and (ii) R belongs to this collection. $\text{Def}(U)$ violates VCP since (i) it involves quantification over the collection of all things and (ii) U belongs to this collection. Similar remarks apply to $\text{Def}(V)$ and $\text{Def}(V_P)$.

Whether VCP motivates a No Set response to the second translation of Cantor’s paradox and to Burali-Forti’s paradox depends on the identities of cardinal and

ordinal numbers. Both $\text{Def}(K)$ and $\text{Def}(O)$ satisfy condition (i): $\text{Def}(K)$ involves quantification over the collection of all cardinals and $\text{Def}(O)$ involves quantification over the collection of all ordinals. If cardinals and ordinals are identified with sets of cardinals and sets of ordinals, so that K is a cardinal number and O is an ordinal number, $\text{Def}(K)$ and $\text{Def}(O)$ will also satisfy condition (ii): K will belong to the collection of all cardinal numbers and O will belong to the collection of all ordinal numbers. Hence, $\text{Def}(K)$ and $\text{Def}(O)$ will violate VCP. On the other hand, if cardinals and ordinals are identified with different sets or if they are sui generis, so that K is not a cardinal number and O is not an ordinal number, then K and O will not satisfy condition (ii). Consequently, $\text{Def}(K)$ and $\text{Def}(O)$ will not violate VCP.

VCP provides more unilateral support for a No Function response. In 3.2, I defined four functions—one for Russell’s paradox, two for Cantor’s paradox and one for Burali-Forti’s paradox—and I have shown that each satisfies Transcendence and Closure if it is defined for every $u \subseteq w$. Here is a list of all four functions, which I now label f_1 – f_4 :

(f_1) Russell’s paradox: $f_1(u) = u$, where $u \subseteq R$

(f_2) Cantor’s paradox: $f_2(u) = x_u$, where $u \subset w \subseteq X$, $X = U, V$, or V_P , $g(x_u) = u$ and g is an arbitrary bijection from X to $\mathcal{P}X$

(f_3) Cantor’s paradox: $f_3(u) = \text{Card}(\mathcal{P} \cup g_c(u))$, where $u \subseteq K$, g_c is a choice function mapping every $\kappa \in K$ to an arbitrary set of cardinality κ and $g_c(u)$ is the set formed by replacing every $\kappa \in u$ with $g_c(\kappa)$.

(f_4) Burali-Forti’s paradox: $f_4(u) = \text{Ord}(g(u))$, where $u \subseteq O$ and g maps u to the initial segment of ordinals defined by u .

Since the contradiction results when $u = w$, it follows that someone who adopts No Function must deny that f_1 – f_4 are defined when $u = w$, i.e., she must deny that

$f(w)$ exists. We've seen that $w = R$ in the case of Russell's paradox, that $w = \{x \in X \mid x \notin g(x)\}$ in the case of the first translation of Cantor's paradox, that $w = K$ in the case of the second translation of Cantor's paradox and that $w = O$ in the case of Burali-Forti's paradox. Consequently, someone who adopts No Function must deny that each of the following exists:

- $f_1(R)$,
- $f_2(w)$, where $w = \{x \in X \mid x \notin g(x)\}$ and $X = U, V$, or V_P ,
- $f_3(K)$,
- $f_4(O)$.

Russell's vicious circle principle supports the claim that none of these values exists, since each of f_1 – f_4 violates the principle under the arguments R , w , K and O .

In the case of Russell's paradox, $\text{Def}(f(R))$ violates VCP since $f(R) = R$ and $\text{Def}(R)$ violates VCP. The definitions used in the other paradoxes— $\text{Def}(f(w))$, $\text{Def}(f(K))$ and $\text{Def}(f(O))$ —can be shown to satisfy conditions (i) and (ii) and consequently to violate VCP as well. Each of these three definitions satisfies condition (i), for $f(w)$, $f(K)$ and $f(O)$ are defined as functions of the sets w , K and O and consequently are defined in terms of quantification over w , K and O . To show that each of these definitions satisfies condition (ii), it suffices to show that $f(w)$, $f(K)$ and $f(O)$ belong to w , K and O . We already did this, in the translations above, when we proved that each of these functions satisfies Closure ((2.2), (3.2) and (4.2)). Since $\text{Def}(f(w))$, $\text{Def}(f(K))$ and $\text{Def}(f(O))$ satisfy conditions (i) and (ii), these definitions violate VCP. Consequently, VCP motivates a No Function response to all three of the set-theoretic paradoxes.

The plausibility this lends to No Function depends on the prior plausibility of VCP. Unfortunately for VCP, mathematics and logic are rife with impredicativity.

In blocking all impredicative definitions, VCP imposes restrictions that are so strict, they would block many otherwise benign definitions and cripple much of mathematics. For this reason, it seems that VCP misdiagnoses the problem. Note that Russell’s schema is less restrictive. Any function satisfying Closure is impredicative: in each case, $f(w)$ is defined by quantifying over w and, by Closure, $f(w) \in w$. But this is allowed by Russell’s schema, so long as f does not also satisfy Transcendence. Unless it is combined with Transcendence, this sort of circularity is benign, according to Russell’s schema.

3.4 Logical relations between No Function and No Set

In this section, I consider some of the logical relations between No Function and No Set. Say that No Function is ‘logically equivalent’ to No Set if the question of whether $f(w)$ exists determines the question of whether w exists and that No Function is ‘logically independent’ from No Set if the question of whether $f(w)$ exists does not determine the question of whether w exists. I argue that No Function is equivalent to No Set with respect to f_1 and f_2 and that No Function is independent from No Set with respect to f_3 and f_4 .

3.4.1 f_1 and f_2 : No Function is equivalent to No Set

In the case of Russell’s paradox, $f_1(R) = R$. Hence, the No Function response—that $f_1(R)$ does not exist—is indistinguishable from the No Set response—that R does not exist.⁶

The case of the first translation of Cantor’s paradox is similar, though the point is not as immediate. According to the No Function response, $f_2(w)$ does not exist. Given the assumption that the bijection $g : X \rightarrow \mathcal{P}X$ exists, it seems that the only way that $f_2(w)$ could fail to exist would be if w was not a permissible argument for

⁶As Russell himself observes: 1906, 143.

f_2 . This would be the case only if $w \not\subseteq X$. But it is evident from the definition of w (see D3, p. 75) that $w \subset X$. Thus, we are led back to question whether there is a bijective function $g : X \rightarrow \mathcal{P}X$. I showed earlier (chapter 2, pp. 33, 36–37) that we can define injections from X into $\mathcal{P}X$ and from $\mathcal{P}X$ into X ; hence, by SBT, we know that there is a bijection, g , if X and $\mathcal{P}X$ exist. The only remaining option, then, is to deny that X (and consequently $\mathcal{P}X$) exists. ‘ X ’ stands for U , V or V_P . So, it seems that the only option is to deny that these sets exist. The No Function response is therefore equivalent to the No Set response.

3.4.2 f_3 and f_4 : No Function is independent from No Set

Things change when we come to the second translation of Cantor’s paradox and to Burali-Forti’s paradox. In both of these cases, it’s possible to consistently hold that the contradiction-inducing sets K and O exist and then to block the contradiction “at the next step,” by denying that the functions f_3 and f_4 are well-defined for the arguments K and O .

Consider, first, the second translation of Cantor’s paradox. Recall that f_3 maps u to the cardinal number of $\mathcal{P} \bigcup g_c(u)$, where $u \subseteq K$ and g_c is a choice function on $\mathcal{P}K$ that replaces every $\kappa \in u$ with a set of cardinality κ . f_3 may fail to be well-defined when $u = K$ for two reasons:

R1. One of the principles for set formation used to construct $\mathcal{P} \bigcup g_c(K)$ from K fails.

R2. $\mathcal{P} \bigcup g_c(K)$ has no cardinal number because it is too big.

Let’s briefly consider the plausibility and implications of each.

The set formation principles relevant to R1 are: Choice and Replacement (used to construct $g_c(K)$ out of K); Union (used to construct $\bigcup g_c(K)$ out of $g_c(K)$); and Powerset (used to construct $\mathcal{P} \bigcup g_c(K)$ out of $\bigcup g_c(K)$). While any of these *might* be

rejected, an attractive response along this line would have to provide some account for *why* the selected principle fails.

If R2 obtains (and R1 does not), there is good reason to infer that not only $\mathcal{P}\bigcup g_c(K)$, but also K has no cardinal number. For if R2 is true, then this is either because $\bigcup g_c(K)$ has no cardinal number (because it is too big also) or because the powerset operation has taken us from a set with a cardinal number ($\bigcup g_c(K)$) to a set that is too big to have a cardinal number ($\mathcal{P}\bigcup g_c(K)$). Similarly, if $\bigcup g_c(K)$ has no cardinal number, then this is either because $g_c(K)$ has no cardinal number or because the union operation has taken us from a set with a cardinal number ($g_c(K)$) to a set that is too big to have a cardinal number ($\bigcup g_c(K)$). Finally, if $g_c(K)$ has no cardinal number, then it must also be the case that K has no cardinal number, for the operation of replacing each $\kappa \in K$ with a κ -sized set produces a set no bigger than K . Now, it seems arbitrary for the powerset operation or the union operation to take us from sets with cardinal numbers to sets that are too big to have cardinal numbers. So, the possibility that R2 is the reason why f_3 fails to be well-defined for $u = K$ reduces to the possibility that K has no cardinal number.

Next, consider Burali-Forti's paradox. Recall that f_4 maps u to the ordinal number of $g(u)$, where $g(u)$ is the initial segment of ordinals defined by u . Like f_3 , there are two reasons why f_4 may fail to be well-defined when $u = O$:

R3. O fails to define an initial segment of ordinals.

R4. O does not have an ordinal number.

Note however, that R3 implies that O does not exist; for the initial segment of ordinals defined by O (i.e., the set of all ordinals less than or equal to any $x \in O$), if it exists, *is* O . So, if R3 is true, then O does not exist. This means that No Function is equivalent to No Set, for No Set blocks Burali-Forti's paradox by denying that O exists. On the other hand, R4 is compatible with the existence of O . So, if R4 is true,

then No Function and No Set are independent. R4 invites the question: why does O not have an ordinal number? Cantor's (1897) proof that O is well-ordered rules out one possible answer, viz., that the ordinals are not well-ordered. Perhaps the only alternative is that sets such as O lack ordinal numbers because they are "too long."⁷

R4 has some interesting historical precedent in the work of one of Russell's students, Phillip Jourdain. Jourdain thought that No Function afforded the best defense of Russell's 1903 set theory, which was committed to the contradiction-inducing set O . He claimed that this was no mark against Russell's theory, since it was not the existence of O that was to blame for Burali-Forti's contradiction, but rather the assumption that the functions $\text{Ord}(x)$ and $\text{Card}(x)$ were defined on O . For this reason, Jourdain argued, O should be called 'absolutely infinite', not 'inconsistent'.

The series of all ordinal numbers may, it seems to me, properly be called an "*absolutely*" infinite series. For, if a well-ordered series has a type, it is, in a certain sense, completed; while the above series [O] cannot, as is shown by Burali-Forti's contradiction, have a type.

This seems to be the most promising way of regarding Burali-Forti's contradiction, and the words "absolutely infinite" seem preferable to the equivalent word "inconsistent," which I, in common with Cantor, have used hitherto; because an "inconsistent" aggregate is not itself contradictory (it exists in the mathematical sense of the word), but a cardinal number or type of it does not exist.⁸

Unfortunately, Jourdain's suggestion was of little help: Russell's theory was also committed to $\text{Card}(K)$ and $\text{Ord}(O)$.⁹ Of course, these commitments are the result of Russell's identification of the cardinal and ordinal numbers with particular sets. They might be avoided by another theory, in which cardinals and ordinals are sui generis.

⁷Recall that length is a matter of well-ordering relations. See chapter 2 for discussion.

⁸Jourdain (1905, 54)

⁹These are identified (respectively) with the set of all sets equinumerous to K and the set of all sets isomorphic to O and the existence of these sets is entailed by Russell's comprehension principle. See Hallett (1984, 179–180).

3.4.3 A uniform solution

It is natural, once one comes to think of the paradoxes as sharing a common form, to look for a uniform solution to all three. No Set is uniform. No Function is not. No Function is a mix of two diagnoses. It blocks Russell's paradox and the first translation of Cantor's paradox in one way: by denying the existence of the sets R , U , V and V_P ; it blocks the second translation of Cantor's paradox and Burali-Forti's paradox in another way: by denying the existence of $f_3(K)$ and $f_4(O)$ but *not* the existence of the sets K and O . For this reason, No Set is preferable to No Function. However, a satisfactory solution requires more than a diagnosis, it requires also an explanation of why the diagnosis is true. In the present case, if we are to embrace the No Set diagnosis, it remains to explain why the contradiction-inducing sets do not exist. I will return to the question of explanation in chapters 4 and 6.

CHAPTER 4

THE LOGICAL CONCEPTION OF *SET*

In this chapter, I consider the *logical* conception of *set*, according to which a set is the extension of a concept (or property). This conception is called “logical” because at the time it was developed in the 19th century, it was widely held that logic was the study of concept extensions, which were viewed as logical objects. Thus, syllogistic logic, which comprised much of logic in the 19th century, was commonly explicated in terms of relations of inclusion and exclusion between concept extensions. For example, the Barbara syllogism—“all As are Bs, all Bs are Cs; therefore all As are Cs”—was explicated as the claim that the extension of the concept C includes the extension of the concept A (“all As are Cs”) if the extension of the concept C includes the extension of the concept B (“all Bs are Cs”) and the extension of the concept B includes the extension of the concept A (“all As are Bs”). The logical conception of *set* is also sometimes called “naive” in light of the relative ease by which it can be shown to engender contradictions. Arguably, its most influential proponent was Gottlob Frege, who argued that arithmetic could be reduced to a formal theory of sets as extensions of concepts (Frege, 1884, 1893, 1903).

Frege’s views will play a significant role in this chapter, which is organized as follows. In 4.1 and 4.2, I present a broadly Fregean account of concepts and their extensions. In 4.3, I discuss the relation of priority between concepts and concept extensions and I contrast this with the relation of priority between the members of a collection and the collections they form. (The iterative conception of *set*, which I discuss in chapter 6, is based on the identification of sets with collections.) In

4.4, I formulate a set theory based on the logical conception of *set* and I show this theory to be inconsistent by deriving Russell’s paradox. In 4.5 and 4.6, I present and evaluate two “fixes” that purport to restore consistency without abandoning the logical conception of *set*. These are the limitation of size doctrine and first-order set restrictivism.

4.1 Concepts

The Fregean account of concepts is based, in large part, on the notion of a mathematical function. In mathematics, expressions such as ‘ $y = x^2$ ’ stand for functions that take numbers as argument (in the ‘ x ’-place) and yield numbers as value (in the ‘ y ’-place).¹ Frege (1891) argues that ordinary predicate expressions are similar to these mathematical expressions: they too stand for functions, but instead of taking numbers to numbers, these functions take all objects (including numbers) to truth-values. Such functions are commonly called “truth-functions.” Frege’s notion of a concept is the notion of a truth-function defined on absolutely all objects that is denoted by a (possible) predicate expression. To illustrate, consider the predicate expression ‘is a horse’. According to the Fregean account, this stands for a particular truth-function from absolutely all objects to truth-values, viz., the concept *horse*. Pursuing the analogy with mathematics, if we add an “argument variable”, we can express this concept by means of the more mathematical-looking expressions: ‘ x is a horse’ or ‘horse(x)’. If we add a “value variable”, we can express this concept as we did the quadratic function above, by means of an equation: ‘ $y = \text{horse}(x)$ ’. This equation represents the concept as the truth-function that takes objects as argument (in the ‘ x ’-place) and yields a truth-value (in the ‘ y ’-place): this truth-value is True whenever x is a horse; otherwise it is False.

¹Frege (1891, 133) suggests replacing ‘ x ’ with a blank space: ‘ $y = ()^2$ ’ to make it clear that ‘ x ’ indicates the function’s argument position and not any particular object.

The identification of Fregean concepts with truth-functions provides an extensional account of meaning for predicate expressions. (An account of ‘meaning’ is *extensional* if it assigns the same meaning to predicate expressions that are satisfied by all the same objects, i.e., if it assigns the same meaning to *coextensional* predicate expressions.) Such an account is suitable to mathematics, in which *intensional* distinctions—distinctions in meaning between coextensional expressions—are generally ignored. In most mathematical contexts, for example, any differences between expressions that denote the same thing, such as the same number or the same function, are ignored. The reason is that in these contexts, one is concerned with the identity of the number or function in question, not with how it is represented. Because ‘4’ denotes the same number as ‘ 2^2 ’, it makes no (mathematical) difference whether ‘4’ or ‘ 2^2 ’ is substituted for ‘ x ’ in a given formula. Similarly, because ‘ $y = x^2$ ’ and ‘ $y - x^2 = 0$ ’ denote the same function, it makes no difference which expression we use. In other contexts, intensional distinctions are important and cannot be ignored. For example, in the interests of developing a satisfactory analysis of natural language, one must distinguish between the meaning of coextensional expressions such as ‘has a kidney’ and ‘has a heart’. Of course, truth-functions are inadequate for such a project. If concepts are to serve as meanings in this context, they would have to be more finely individuated, like properties “in the abundant sense” (Lewis, 1983, 10–19). In what follows, I reserve the term ‘concept’ for the coarser-grained notion of a truth-function.

In the philosophy of language today, it is common to mark the distinction between truth-functional meaning and non-truth-functional meaning as the distinction between ‘extensional meaning’ and ‘intensional meaning’. Frege (1892c) recognized such a distinction, but he expressed it in different terms. He distinguished the *reference* (Bedeutung) of an expression—which for predicates was always a concept and thus truth-functional—from its *sense* (Sinn)—which was not truth-functional.

Whereas we would say that the predicate expressions ‘has a kidney’ and ‘has a heart’ have different intensional meanings but the same extensional meaning, Frege would say that they have different senses but the same reference, i.e., they refer to the same concept. It is likely that this difference is only terminological, not substantive. There is a second difference, however, between Frege’s views of concepts and contemporary conceptions of meaning that is quite substantive. Frege held that concepts could only be referred to indirectly, by predicate expressions occurring predicatively. This led to the baffling conclusion that while the occurrence of ‘is a horse’ in the sentence ‘Ed is a horse’ refers to the concept *horse*; the singular expression ‘the concept *horse*’ does not (Frege, 1892b, 184–185). The tenability of Frege’s position and his reasons for adopting it have been discussed in a number of places.² I will not follow Frege here. I assume that it is possible to name concepts by singular expressions such as ‘the concept *horse*’. In this respect, my account of concepts is unFregean.

Two additional notes about concepts are in order. First, to contemporary ears, talk of concepts is likely to suggest something mental or linguistic. Fregean concepts are neither mental nor linguistic; their existence depends neither on our thoughts, nor on the expressive resources of any actual language. In these respects, they are fully objective. It may be more common to use the term ‘property’ or simply ‘truth-function’ to convey the sort of objectivity that is intended. The second note concerns a distinction between determinate and indeterminate concepts. Say that a concept, F , is *determinate* iff for any object, x (or objects the xx) it is determinate whether or not x (or the xx) fall(s) under F . Similarly, say that a concept, F , is *indeterminate* iff there is at least one object, x (or objects the xx) such that it is indeterminate whether or not x (or the xx) fall(s) under F . If the predicate expressions ‘cloud(x)’ and ‘pile(xx)’ refer unambiguously to the concepts *cloud* and *pile*, then it would seem

²See Dummett (1981), Parsons (1986) and Hale and Wright (2001, Essay 3).

to follow that these concepts are indeterminate. For it is often indeterminate whether some portion of sky is a cloud or whether some grains of sand constitute a pile. Since concepts are objective, it follows that belief in indeterminate concepts is committed to ontological vagueness.

Frege held that mathematics and logic were precise sciences and consequently that they required determinate concepts. There are a number of passages in which he makes this point. For example, in (1884, 41), he argues that subjective definitions of the number one in terms of the indeterminate (or “hazy”) concept *being thought of as one (or indivisible)* are unacceptable since they fail to uniquely determine whether the number of some thing (or things) is (or are) one. This sort of indeterminacy is unacceptable in arithmetic, “which bases its claim to fame precisely on being as definite and accurate as possible.” In (1903, 259) he argues that logic also demands determinate concepts.

A definition of a concept (of a possible predicate) must be complete; it must unambiguously determine, as regards any object, whether or not it falls under the concept (whether or not the predicate is truly ascribable to it). Thus there must not be any object as regards which the definition leaves in doubt whether it falls under the concept . . . We may express this metaphorically as follows: the concept must have a sharp boundary.

Indeed, he goes on to claim that the law of excluded middle is to be understood as the demand that, in logic, every concept must be determinate.

It is unclear whether Frege meant to deny the existence of indeterminate concepts entirely or only to deny their admissibility to the fields of precise sciences such as logic and mathematics.³ Whatever Frege’s intention might have been, the undesirability of ontological vagueness may be reason enough to deny the existence of indeterminate concepts. Someone who does so can provide a linguistic account of indeterminacy, according to which the indeterminacy we see in cloud-talk and pile-talk

³See van Heijenoort (1986) for a defense of the former interpretation; see Ruffino (2003) for a defense of the latter interpretation.

is the result of indeterminacy in the predicates ‘cloud(x)’ and ‘pile(xx)’.⁴ Say that a *predicate* is determinate if it unambiguously denotes a single concept; otherwise, it is indeterminate. Equipped with these notions, we can attribute the indeterminacy of the predicate ‘cloud(x)’ to the fact that it simultaneously expresses multiple, but individually determinate cloud concepts, each differing from the others with respect to the precise division of the sky into individual clouds. Similarly, we can attribute the indeterminacy of the predicate ‘pile(xx)’ to the fact it simultaneously expresses multiple, but individually determinate pile concepts, each differing from the others with respect to the minimum number of grains necessary to make a pile. In what follows, I will assume that all concepts are determinate.

4.2 Extensions

Like his account of concepts, Frege’s account of their extensions is based on an analogy to mathematics. Frege (1891, 134) notes that any numeric function can be geometrically represented in the familiar way as a set of points, “that presents itself to intuition (in ordinary cases) as a curve.” To illustrate, he asks us to consider a particular quadratic function: $y = x^2 - 4x$. This function is represented geometrically as an upwards parabolic curve, with a vertex at $(2, -4)$ that intersects the x -axis at $(0, 0)$ and $(4, 0)$. Frege (1891, 135) calls this curve “the graph” (Wertverlauf) of the function. His idea is to provide an account of extensions by generalizing on the notion of a geometric graph. This generalization is a sort of corollary to the previous generalization of the notion of a numeric function, according to which not only mathematical expressions (such as ‘ $y = x^2$ ’) but also non-mathematical predicate expressions (such as ‘ x is a horse’) refer to functions. The corollary is that not only mathematical functions (such as the quadratic function expressed by ‘ $y = x^2$ ’), but

⁴See Lewis (1993) for an account of indeterminacy along these lines. See Ruffino (2003, 272–275) for an argument that this account fits nicely with some of Frege’s own remarks on indeterminacy.

also non-mathematical functions, in particular, truth-functions (such as the truth-function expressed by ‘ x is a horse’), determine graphs. Since concepts are truth-functions, every concept has a graph. The graph of a concept is its extension.

We can designate as an extension the [graph] of a function whose value for every argument is a truth-value (Frege, 1891, 139).

But just what is a graph, in this generalized sense? In the mathematical case, graphs of simple numeric functions can be visualized or imagined as curves; however, it’s likely that many (if not most) numeric functions exceed human powers of visualization and imagination. For this reason, it seems that an adequate account of even purely geometric graphs should not be tied to such human cognitive capacities. A more promising account of geometric graphs is one that is based on the role they play of representing functions. A geometric graph represents a function by representing two things. First, each point along the graph represents a unique pair of numbers, $\langle n, f(n) \rangle$, which is determined by the function. The first member of this pair, n , is an argument for the function; the second member, $f(n)$, is the value of the function for this argument. Second, any two points along the graph, representing the pairs $\langle n_1, f(n_1) \rangle$, $\langle n_2, f(n_2) \rangle$, are ordered by the relation $<$ defined on the arguments n_1 , n_2 . This ordering is needed to fix the graph’s shape. Putting these two observations together, we might say that a geometric graph is an entity, consisting of a series of argument-value pairs, ordered under the relation $<$ defined on the arguments, which represents a function. Alternate translations of ‘Wertverlauf’, as ‘course-of-values’ or ‘value-range’ may suggest this view more vividly than ‘graph’.

Extending this account to graphs in general, it would follow that the graph of a concept, i.e., an extension, is also an entity, consisting of a series of argument-value pairs, ordered under a relation R defined on the arguments. The arguments are objects; the values are truth-values determined by the corresponding concept. Since concepts are determinate, this series is ‘total’, meaning that absolutely every object is included as the argument of one pair. According to the present account,

then, the extension of the concept *horse* consists in the total ordered series of the argument-value pairs $\langle x, \textit{horse}(x) \rangle$, where $\textit{horse}(x)$ is the value True if x is a horse and the value False otherwise.

However, the analogy between geometric graphs and concept graphs is imperfect. For amongst objects in general there is no analogue of the mathematical relation $<$ to provide a natural ordering. Consequently, any ordering of objects in general—which is necessary to generate an ordered *series* of argument-value pairs—would be arbitrary. Such arbitrariness is unacceptable, since it would make the identity of concept extensions arbitrary as well. To illustrate, the identity of the extension of the concept *horse* would vary, depending on the relative ordering of horses and non-horses. The extension that would result from placing every pair containing a horse (as its first member) before every pair containing a non-horse (as its first member) would differ from the extension that would result from placing every pair containing a non-horse (as its first member) before every pair containing a horse (as its first member) and also from the extension that would result from alternating pairs that contained horses (as their first member) with pairs that did not. Since there is no non-arbitrary ordering of objects in general, there is no unique, total series of argument-value pairs that can (non-arbitrarily) be identified with the extension of the concept *horse*. This result is completely general. There is no unique, total series of argument-value pairs that can (non-arbitrarily) be identified with the extension of any concept. It follows that if extensions are graphs at all, they are unordered graphs, i.e., entities consisting of argument-value pairs that are unordered relative to one another. I will refer to such graphs as ‘mappings’. A mapping represents a function, but it does not represent it as having any determinate “shape”. You might think of a mapping as the result of abstracting away the shape from a geometric graph. I submit that the Fregean notion of a concept extension is the notion of a total mapping.

4.2.1 Involvement and membership

Since extensions are total mappings, they “involve” every object in the sense that every object is included in the mapping. This notion of involvement in an extension should be distinguished from the notion of membership in an extension. Let ‘ $\text{Ext}(F)$ ’ denote the extension of the concept F . While every object is involved in $\text{Ext}(F)$, only those objects that fall under F are members of $\text{Ext}(F)$.

Involvement in an extension is defined by:

D1. x is *involved* in $\text{Ext}(F) =_{df}$ x is an argument for F .

Since $\text{Ext}(\textit{horse})$ is a total mapping, absolutely every object is involved in $\text{Ext}(\textit{horse})$, for absolutely every object is an argument of the function ‘ $\textit{horse}(x)$ ’.

The notion of *membership* in an extension is defined by:

D2. x is a *member* of $\text{Ext}(F) =_{df}$ Fx .

While absolutely every object is involved in $\text{Ext}(\textit{horse})$, only horses are members of $\text{Ext}(\textit{horse})$. (In what follows, I use the usual membership sign, ‘ \in ’ for membership in an extension.)

4.2.2 Identity

The identity conditions for extensions are suggested by another analogy between mathematical and non-mathematical language. Frege (1891, 135) notes that numeric expressions which return the same values for the same arguments—what we might call “coextensional numeric expressions”—have the same graph.⁵ By analogy, it seems to follow that coextensional predicate expressions must have the same extension. However, there is no need to resort to analogy at this point: the definitions of concepts as truth-functions and extensions as mappings are enough to fix identity conditions

⁵Frege’s example involves two algebraic equations for the same parabolic curve.

for extensions. On the one hand, if two concepts are identical, then their extensions are identical (since otherwise, there would be identical concepts whose mappings differed with respect to least one argument-value pair, which is impossible, given the assumption that concepts are determinate truth-functions). On the other hand, if two extensions are identical, then their concepts are identical (since otherwise, there would be distinct concepts that determined all the same argument-value pairs, which is impossible given the assumption that concepts are determinate truth-functions). Putting these results together, two concept extensions are identical iff the corresponding concepts are identical.

The most direct way of expressing this is as follows:

$$(1) (\forall F)(\forall G)[\text{Ext}(F) = \text{Ext}(G) \leftrightarrow F = G]$$

However, we can express virtually the same thought without speaking of identity between concepts so directly. Say that two concepts F and G are coextensional if they have the same values for the same arguments, i.e., if $(\forall x)(Fx \leftrightarrow Gx)$. Given the extensional identity conditions for concepts, we know that $F = G$ iff they are coextensional. Hence, we might substitute the relation of coextensionality between concepts for the relation of identity between concepts in (1) and thereby express the identity conditions for extensions as:

$$(2) (\forall F)(\forall G)[\text{Ext}(F) = \text{Ext}(G) \leftrightarrow (\forall x)(Fx \leftrightarrow Gx)].$$

In fact, this is just what Frege did. The reason why and the relation between (1) and (2) is worthy of note. Recall Frege's view that concepts can only be referred to by predicate expressions occurring predicatively. A consequence is that concept terms cannot flank the identity sign. Hence, Frege could not express the identity conditions for extensions directly, by means of (1); instead he used the relation of

coextensionality between concepts as a sort of stand-in for first-order identity and expressed the identity conditions between concept extensions by means of (2).⁶

4.3 The priority of concepts to extensions

The Fregean view of extensions as mappings determined by concepts suggests that concepts are prior to their extensions. The sort of priority that is at play here may be of any of three types (Burge, 1984, 285). First, concepts may be *semantically prior* to extensions if the definitions of extensions involve reference to or quantification over concepts (and the converse does not hold: i.e., the definitions of concepts do not involve reference to or quantification over extensions). Second, concepts may be *epistemically prior* to extensions if extensions are properly conceived as belonging to concepts, as being the extensions *of* concepts (and the converse does not hold: i.e., concepts are not properly conceived as belonging to extensions). Third, concepts may be *ontologically prior* to extensions if the identity and existence of extensions is *grounded* in the identity and existence of concepts (and the converse does not hold: i.e., the identity and existence of concepts is not grounded in the identity and existence of concepts).⁷

Semantic priority is supported by the following asymmetry of reference. On the one hand, expressions that refer to individual concepts, e.g., ‘the concept *horse*’, or simply ‘*horse*’, do not involve any obvious reference to extensions. However, the converse does not seem to hold. Expressions that refer to individual extensions, e.g., ‘the extension of the concept *horse*’, or ‘Ext(*horse*)’, do involve obvious reference

⁶Frege (1892a, 175–176) writes that it is coextensionality that “we have in mind” when we say things such as “the concept *F* is identical to the concept *G*.”

⁷Klement (2012, 154–155) argues for the stronger view that extensions *are* concepts, conceived of as objects. A similar view may have been held earlier by Wright (1983, 19): “Extensions now emerge simply as concepts objectified.” See also Cocchiarella (1987).

to concepts, for they contain expressions such as ‘the concept *horse* and ‘*horse*’ as constituent subexpressions.

An argument for epistemic priority is suggested by Frege’s views on the relation between statements of identity between extensions and statements of coextensionality between concepts, corresponding to the left and right-hand sides of (2) above.⁸ Frege held that the two sides of (2) were semantically, but not epistemically equivalent. Statements of identity between extensions were semantically equivalent to statements of coextensionality between concepts, since each had the same truth condition⁹; however, the latter were epistemically prior, since extensions were only known in terms of the relation of coextensionality on concepts (Frege, 1884, 73–79). Consequently, (2) functions as an epistemic principle for expanding our conceptual repertoire. Principles that play this role are commonly called abstraction principles. They work by employing an equivalence relation between *previously understood* entities to define identity statements involving a *new* type of singular term and thereby to introduce a *new* kind of object. In the case of (2), the previously understood entities are concepts, the equivalence relation between them is coextensionality, and the identity statements that are defined are statements of the form ‘ $\text{Ext}(F) = \text{Ext}(G)$ ’. The new singular terms are ‘ $\text{Ext}(F)$ ’ and ‘ $\text{Ext}(G)$ ’, which refer to a new kind of object: extensions.

The ontological priority of concepts to extensions may be suggested by the observation that an extension, being a total mapping, is nothing over above its associated concept and the objects to which this concept applies or does not apply. However, this observation is unhelpful, for to the extent that it supports the thesis that concepts are prior to their extensions, it also supports the thesis that all objects are prior to extensions. This latter thesis has the absurd consequence that extensions are prior to

⁸For a development and defense of Frege’s views, see Wright (1983, ch. 1), Hale (1988) and Hale and Wright (2001, Introduction, Essays 5, 6).

⁹See Hale and Wright (2001, Essay 4) for discussion of several candidate notions of same truth-condition in this context.

themselves. It is interesting to note that when Frege speaks of the priority relation between concepts and their extensions in ontological terms he speaks of the extension of a concept as depending solely on the concept and *not* on objects.

The extension of a concept is constituted in being, not by the individuals, but by the concept itself (Frege, 1895, 224–225).

The extension of a concept simply has its being in the concept, not in the objects which belong to it (Frege, 1906, 183).

Since, for Frege, the concept does not depend on objects, his view avoids the absurd consequence that extensions are prior to themselves.

Whether we think that concepts are ontologically prior to extensions may depend on how we understand the statements (1) and (2) which express the relation between the two. It's possible to view these statements merely semantically, or epistemically, as defining identity for extensions in terms of the semantically, or epistemically prior relation of identity (or coextensionality) between concepts. But we might go further and claim that this relation of identity (or coextensionality) between concepts is what grounds the relation of identity between extensions. Under this interpretation, (1) and (2) express an ontological priority of concepts to their extensions.

Insofar as concepts are prior to extensions in some, or all, of these senses, the logical conception provides a “top-down” view of sets according to which concepts—which one may picture as being situated above the domain of objects in a Platonic heaven—come first, and their extensions—which one may picture as the shadows cast by the concepts above—come second. Contemporary philosophers of set theory are accustomed to work the other way round: to take the members of a set as basic and to then tell a story about how the sets are formed out of their elements by the process of forming collections. Various images are used to illustrate the process of collection formation: one forms a collection by coralling its elements, or lassoing them (Kripke), or thinking of them as one (Cantor), or enclosing them in a sack (Dedekind). This

conception (which I discuss in detail in ch. 6) is “bottom-up”, insofar as the elements of collections are prior to any collections they form.

To avoid confusion, a brief note on my use of ‘collection’ is in order. Unless otherwise noted, I use this term in a rather technical sense to refer to those objects that answer to the bottom-up conception of *set*. Collections are singular, non-mereological entities whose identities are determined by the identity of their members, but which are something over and above their members in the sense of being formed out of them. Although technical, this use is not unfamiliar. For example, Charles Parsons (1997, 174) uses the term in this way when he writes:

By a collection I mean an object that consists of its elements; it is typically thought of as in some way “formed” from them ... Thought of as a collection, a set is formed from objects that are already “available” or “given”.

This use of ‘collection’ is much narrower than the ordinary meaning, which allows for uses of ‘collection’ that refer to entities that are not taken to be something over and above their members, for example, uses of ‘collection’ to refer to mere pluralities, such as “the collection of planks from Theseus’s original ship.” Ordinary speech also allows for uses of ‘collection’ that refer to entities that are something over and above their members but which are not formed out of their members, for example, uses of ‘collection’ to refer to extensions. Thus, George Boolos (1998, 14), in his exposition of the logical, top-down conception of *set* as an extension, in a paper in which distinguishes this conception from the iterative, bottom-up conception of *set*, uses ‘collection’ as a synonym for extensions. Neither of these uses of ‘collection’ are permitted under my restriction.

4.3.1 Extensions as pluralities

I have said that extensions are (total) mappings. This leaves many questions about the nature of extensions unanswered. In particular, we would like to know

more about the sort of relation that an extension (mapping) bears to the argument-value pairs that constitute it. Since the current project is to provide a conception of *set*, we cannot elucidate this relation in set-theoretic terms, i.e., we cannot identify extensions with *sets* of argument-value pairs.¹⁰ We cannot, for example, conceive of the extension of the concept *horse* as the set of argument-value pairs $\{\langle \text{Bucephalus}, \text{True} \rangle, \langle \text{Socrates}, \text{False} \rangle, \dots\}$. But how else can we conceive of it?

Nor can we say that extensions are *collections* of argument-value pairs. If extensions are collections, then the relation of membership in an extension must be a determinate of the determinable relation of membership in a collection; but it seems that this cannot be the case, for the latter relation is irreflexive and so the notion of a self-membered collection is incoherent; but the notion of a self-membered extension is not. According to D2,

$$\text{Ext}(F) \text{ is a self-member} =_{df} F(\text{Ext}(F)).$$

In other words, to say that an extension is a self-member is simply to say that it falls under the concept of which it is the extension. There is nothing particularly puzzling about this. There are even positive examples of extensions meeting this description. Consider, for example, the extension of the concept *self-identical*. Certainly, $\text{Ext}(\textit{self-identical})$ —if it exists at all—is self-identical. It follows by D2 that $\text{Ext}(\textit{self-identical})$ is a self-member. Similarly with $\text{Ext}(\textit{non-philosopher})$, which is a non-philosopher and $\text{Ext}(\textit{abstract})$, which is abstract. So it seems that the relation of membership in an extension is not irreflexive. Since the relations of membership in an extension and membership in a collection differ in this way, the former cannot be a determinate of the latter and so we cannot identify extensions with collections. One might try identifying extensions with some other sort of singular entity. In addition to being something other than a collection, this entity must satisfy the following desiderata:

¹⁰Michael Beaney (1997, 135) writes that Frege’s notion of an extension is that of “a set of *pairings* of arguments with values.”

(i) it must be one object that unifies (in some way) its many argument-value pairs; (ii) its identity must be uniquely determined by these argument-value pairs, and not by any other things (as is the case with mereological compositions, which are identical to the sum of their parts under any decomposition); (iii) finally, it must be understood independently of set theory. I am doubtful whether there is any type of non-collection satisfying (i)–(iii).

In an intriguing discussion, Max Black (1971, 622–624) considers an analysis of ‘set’ as the extension of a concept, defined as the property of *being one of the objects that falls under it*. Black ends up rejecting this proposal on account of two objections. The first is that it presupposes “our prior possession of the concept of a set” (624) and is therefore unacceptable as an analysis of ‘set’. The second is that the plural phrase ‘being one of the objects that falls under F ’ “reduces to vacuity” in cases where F is uninstantiated. Black’s first objection stems from the fact that he regards sets as mere vehicles of plural reference so that the conception of *set* is exhausted by the rules governing “set talk” (635). This claim, I think, should be resisted. Black’s second objection is based on the more plausible claim that empty plural referring expressions are meaningless. It follows that we cannot make sense of the property of being one of the objects that falls under an uninstantiated concept. This problem does not arise if extensions are mappings, however, for the property of *being one of the argument-value pairs determined by F* is never uninstantiated (assuming there is at least one object). The plural phrase ‘*being one of the argument-value pairs determined by F* ’ refers to just as many pairs for uninstantiated concepts—one pair for every object—as it does for universal concepts. For these reasons, I think that Black’s proposed analysis survives his objections. However, there is a simpler alternative: instead of identifying $\text{Ext}(F)$ with the property of being one of the argument-value pairs determined by F , identify $\text{Ext}(F)$ with the argument-value pairs themselves.

According to this proposal, extensions are *pluralities* of argument-value pairs, so that, for example, the extension of the concept *horse* is the plurality of pairs: $\langle \text{Bucephalus}, \text{True} \rangle, \langle \text{Socrates}, \text{False} \rangle$, etc. Over the past twenty years or so, it has become quite common to enrich classical first-order logic by adding devices for plural reference and quantification. The result is a *plural logic*, which contains plural variables ($'xx'$, $'yy'$) bound by the familiar quantifiers \forall and \exists and a logical predicate, \prec (among). Read $'\exists xx \dots'$ as 'there are some $xx \dots$ '. Read $'\forall xx'$ as 'for any $xx \dots$ '. Read $'xx \prec yy'$ as 'the xx are among the yy '.¹¹

In what follows I will employ these devices freely in philosophical regimentations of plural English claims. Note that while the term 'plurality' is a convenient singularizing term, pluralities are not singular entities. Thus, singular expressions such as 'the extension of F ' and 'the plurality of F 's argument-value pairs' must be understood as stand-ins for plural terms, such as 'the argument-value pairs of F '. Though it may be tedious, we can in general always dispense with reference to pluralities of objects by referring plurally to the objects themselves—for example, we can dispense with reference to a plurality of F s by referring directly to the F s—and we can always dispense with quantification over pluralities of objects by plurally quantifying over the objects themselves—for example, we can dispense with quantification over every plurality of F s in 'every plurality of F s is finite' by means of a paraphrase such as 'any F s are finite'.¹²

¹¹Nonlogical plural terms (e.g., 'the Egyptians') and predicates (e.g., 'built a pyramid') may be added. Boolos (1984) provides one of the earliest presentations and defenses of plural logic. More recent formulations can be found in Burgess and Rosen (1997), Rayo (2002), Yi (2005, 2006), McKay (2006), Oliver and Smiley (2013) and ?. Officially, \prec takes only singular terms in its first argument place and plural terms in its second argument place. Read $'x \prec yy'$ as ' x is one of (among) the yy '. (Unofficially, \prec can take plural terms in its first argument place. Read $'xx \prec yy'$ as an abbreviation for $'(\forall z)(z \prec xx \rightarrow z \prec yy)'$.)

¹²Like a singular quantifier, a plural quantifier ranges over objects taken one-at-a-time. But it also ranges over objects taken two-at-a-time, and three-at-a-time, and so on. A sentence of the form 'some things are F ' is true iff at least some individual or individuals are F . The crucial difference between singular and plural quantification is that such a sentence might be true even if F is a

If extensions are pluralities, then identity between extensions must then be understood as an instance of plural identity. Two pluralities, the xx and the yy , are identical iff every one of the xx is identical to one of the yy and vice versa. Thus, if the xx are the argument-value pairs of some concept F and the yy are the argument-value pairs of some concept G , then the xx are identical to the yy iff every pair among the xx is identical to a pair among the yy and vice versa. However, whenever convenient, I will use singular expressions and speak of identity between extensions as a singular relation.

A plural analysis of extensions raises a concern similar to the one just discussed for the view of extensions as collections. This concern arises from consideration of the extensions of what we might call *plural concepts*: well-defined functions from arbitrary pluralities of objects to truth-values. Other than taking plural arguments, plural concepts are just like the concepts discussed so far. They are truth-functions defined on all pluralities of objects that serve as the extensional meanings of plural predicates in the same way that ordinary concepts do for singular predicates.

A predicate can be plural in a strong or a weak sense. A predicate is weakly plural if any sentence in which it occurs can be reinterpreted in terms of singular predication. Thus, the plural predicate ‘are green’ is weakly plural since its applications, e.g., in the sentence ‘All the leaves are green’, can be reinterpreted in terms of singular predication as the sentence ‘Each of the leaves is green’, in which ‘is green’ takes only singular arguments. A strongly plural predicate cannot be reinterpreted in this way. An example is the predicate ‘are heavy’ in the sentence ‘All the leaves are heavy’. This sentence cannot be reinterpreted as ‘Each of the leaves is heavy’. If we accept ordinary concepts and we allow for strongly plural predication, then it seems that we must also accept plural concepts.

property that no individual has. An intuitive example is the relational property of being married to each other. No individual instantiates this, although many couples do.

Let ‘*TWO*’ stand for the plural concept *at least two in number*. According to the plural analysis of extensions, the extension of *TWO* is the plurality of all the argument-value pairs determined by *TWO*. Moreover, since *TWO* is a total function, it is defined on every plurality of objects; in particular it is defined on the extension of *TWO*. This means that the extension of *TWO* is the plural argument of one of its own argument-value pairs. Since according to the plural analysis of extensions, the extension of *TWO* is the plurality of all the argument-value pairs determined by *TWO*, it follows that this plurality is a member of one of its own argument-value pairs. Admittedly, this sounds strange. But I suspect that this strangeness is due to the assumptions, which are illicit in this context, that the argument-value pairs of a concept are each “built-up” from some original arguments and values and that the extension of a concept is “built-up” from its argument-value pairs. If these assumptions were true, then the arguments of each pair would be prior to the pair, and each pair would be prior to the extension. And the present example would then involve the absurd consequence that the extension of *TWO* is prior to one of its argument-value pairs and therefore prior to itself. But these assumptions are false. According to the present view of extensions as concept mappings, each of the pairs in the extension of *TWO* as well as the extension of *TWO* itself are determined “from above” by applications of the plural concept *TWO*.

4.4 Inconsistency of naive set theory

The logical conception of *set* as the extension of a concept is relatively easy to understand and provides a simple, uniform account of the existence conditions for sets. As it was formulated in ignorance of the paradoxes, it is free from ad hoc devices designed to avoid them. Unfortunately, the set theories that most accurately reflect this conception—naive set theories—are inconsistent. In this section, I formulate a naive set theory based on the logical conception of *set*. I then prove its inconsistency

by deriving Russell's paradox. In the following sections, I consider two modifications to this set theory that seek to restore consistency while remaining as close as possible to the logical conception of *set*.

Naive set theories can be formulated using either a classical first or second-order logic. I prefer a second-order logic, as this allows for quantification over concepts, which makes it possible to express the relation between concepts and sets in the most straightforward way. I call the theory 'second-order naive set theory' (SN). The language of SN includes the first-order variables: ' x ', ' y ' and ' z ', which range over all objects (including extensions) and the second-order variables: ' F ', ' G ' and ' H ', which range over all concepts. In addition, I include the restricted first-order variable ' e ', which ranges exclusively over extensions. I attach subscripts to variables as needed. The membership relation \in is defined by D2.

SN has two axioms. The axiom of full comprehension axiom (FC) specifies the existence conditions for extensions:

$$\mathbf{FC} \quad (\forall F)(\exists e)(\forall x)(x \in e \leftrightarrow Fx).$$

The axiom of extensionality (Ext) specifies identity conditions:

$$\mathbf{Ext} \quad (\forall F)(\forall G)[\text{Ext}(F) = \text{Ext}(G) \leftrightarrow (\forall x)(x \in \text{Ext}(F) \leftrightarrow x \in \text{Ext}(G))].^{13}$$

¹³If we wanted, we could replace (FC) and (Ext) with the single axiom:

$$(1) \quad (\forall F)(\forall G)[\text{Ext}(F) = \text{Ext}(G) \leftrightarrow F = G],$$

from which they can be derived as theorems.

The derivation of FC takes place in two steps. First, we show that (1) implies that every concept has an extension. *Proof:* Instantiate the concept variables ' F ' and ' G ' in (1) to an arbitrary concept, say H , to get: $\text{Ext}(H) = \text{Ext}(H) \leftrightarrow H = H$. By the law of identity, $H = H$. Therefore, $\text{Ext}(H) = \text{Ext}(H)$. Since the logic is classical, there are no empty terms, so $\text{Ext}(H)$ exists. Moreover, H is arbitrary, so it follows that every concept has an extension, i.e., $(\forall F)(\exists e)(e = \text{Ext}(F))$. Next, we apply D2, which specifies the conditions under which something is a member of an extension, to e . Given $e = \text{Ext}(F)$, D2 tells us that for any object, x , $x \in e \leftrightarrow Fx$. This gives us FC.

The principle of extensionality (Ext) follows from an application of D2 directly to (1). First, since concepts are extensional, ' $F = G$ ' is equivalent to ' $(\forall x)(Fx \leftrightarrow Gx)$ '. Second, we know by D2 that ' Fx ' (' Gx ') is equivalent to ' $x \in \text{Ext}(F)$ ' (' $x \in \text{Ext}(G)$ '). Therefore, ' $F = G$ ' is equivalent to ' $(\forall x)(x \in \text{Ext}(F) \leftrightarrow x \in \text{Ext}(G))$ '. In effect, the derivation of Ext from (1) plus D2 amounts to replacing the relation of identity between concepts in (1) with the definitionally equivalent relation of coextensionality between the extensions of these concepts.

FC and Ext jointly entail that each of the concepts *object*, *extension*, *pure extension*, *cardinal number*, *ordinal number* and *non-self-membered extension* defines a unique extension differing from the others with respect to its members.¹⁴ Let U , V , V_P , K , O and R stand for the corresponding extensions, so that SN vindicates A1–A6. This is the first step in deriving the contradictions discussed in chapter two. The second step is to justify any additional principles used in the derivations. As we saw in chapter 2, Russell’s paradox requires no additional principles and is derived in SN as follows:

Instantiate FC to the concept *non-self-membered extension* to get:

$$(4.1) (\exists e)(\forall x)(x \in e \leftrightarrow x \text{ is an extension} \wedge x \notin x).$$

By (4.1) and Ext, ‘ e ’ has a unique value, R , whose members are all and only the non-self-membered extensions. This allows us to eliminate the existential quantifier, replacing the bound occurrence of ‘ e ’ with ‘ R ’:

$$(4.2) (\forall x)(x \in R \leftrightarrow x \text{ is an extension} \wedge x \notin x).$$

(4.2) entails:

$$(4.3) (\forall e)(e \in R \leftrightarrow e \notin e).$$

Instantiate ‘ e ’ to R in (4.3) to get:

$$(4.4) R \in R \leftrightarrow R \notin R.$$

The other paradoxes require additional principles, such as Cantor’s theorem, SBT, BF1 and BF2 (chapter 2) and consequently might be blocked by denying some of the assumptions needed to prove these principles. Any response of this sort, however, would be an instance of the No Function strategy (chapter 3) and, as we’ve seen, this strategy provides a diagnosis of the paradoxes that is less uniform than that provided by the No Set strategy. For this reason, a No Set strategy is preferable.

¹⁴Thus, the concept *object* defines $\text{Ext}(\text{object})$, the concept *ordinal number* defines $\text{Ext}(\text{ordinal number})$ and $\text{Ext}(\text{object}) \neq \text{Ext}(\text{ordinal number})$ since not every object is an ordinal number.

Below, I consider two No Set strategies that strive to preserve as much of the logical conception of set as possible. The first—limitation of size—replaces FC with an explicitly restricted comprehension principle, according to which only “small” concepts have extensions. The second—first-order set restrictivism—reinterprets the first-order quantifiers in FC as implicitly restricted by context so that while every concept defines an extension in every context, the extensions so defined are always small. Similar remarks apply to Ext. First-order set restrictivism amounts to replacing FC and Ext with implicitly restricted principles of comprehension and extensionality. The difference between these approaches can be illustrated by comparing what each says about the concepts that engender the contradiction-inducing sets under FC. Limitation of size tells us that these concepts do not have extensions and hence do not define sets at all. First-order set restrictivism tells us that while these concepts define extensions in every context; there is no context in which these extensions include *all* the objects that intuitively fall under the corresponding concepts.

4.5 Limitation of Size

The limitation of size doctrine has two parts. The first is the imposition of a limit on the size of sets. The second is the claim that this limit is intrinsic to the existence conditions for sets in such a way that the mere fact that some objects, the xx , exceed the limit explains why there is no set of the xx . Both parts can arguably be traced back to Cantor’s view of the “absolutely infinite” as the size of totalities too great to be measured or in any way bounded, though it is somewhat controversial whether Cantor would have endorsed the sort of explanatory claim expressed by the second part (Cantor, 1899, 114).

If the doctrine is going to be informative, we must have some independent grasp on the size of the absolutely infinite. It would seem that the absolutely infinite cannot be measured, or assigned any definite cardinality, since that would imply that it was

bounded or had a limit. On the other hand, in order to avoid a vicious sort of circularity, we cannot define the absolutely infinite simply as the size that is too big for set formation and then apply this standard in our explanation for why some things don't form sets. Fortunately, we are able to get a grip on the size of the absolutely infinite in another way, by discovering independently intelligible totalities which can be proved to be limitless. We can then use these totalities as relative measures in the sense that any totality as large as one of them is shown to be absolutely infinite. Cantor (1899) proceeds in just this way. First, he applies Burali-Forti's reasoning to prove by reductio that the totality of all ordinals is absolutely infinite. Second, he proves that other totalities, such as the totality of all cardinals and the totality of all thoughts, are absolutely infinite, by showing that they are as big as the totality of ordinals. There is nothing sacred about the totality of all ordinals. One might select another totality to start with, such as the totality of all sets (von Neumann, 1925), apply Cantor's reasoning to prove by reductio that this totality is absolutely infinite and then go on to show that other totalities are absolutely infinite, by showing that they are as big as the totality of all sets. I will discuss worries of circularity facing these strategies below.

The limitation of size doctrine seeks to fix SN by replacing FC with a restricted comprehension principle (LOS), according to which only "small" concepts define extensions:

LOS $(\forall F)[\text{Small}(F) \rightarrow (\exists e)(\forall x)(x \in e \leftrightarrow Fx)]$.

The size of a concept is defined in terms of the size of the absolutely infinite.

D3. A concept is *small* if it applies to fewer than absolutely many things.

D4. A concept is *big* if it is not small.

Because LOS, like FC, says that all sets (extensions) are defined by concepts, it can be seen as an attempt to preserve as much of the logical conception of *set* as

possible.¹⁵ The paradoxes are blocked as follows. Given some relative metric for the absolutely infinite, such as the concept *ordinal*, we show that each of the concepts *object*, *extension*, *pure extension*, *cardinal number*, and *non-self-membered extension* has at least as many instances as this metric and is therefore big. It follows by LOS that none of these concepts defines a set. Hence, none of the contradiction-inducing sets (U , V , V_P , K , O or R) exists. In this way, the limitation of size doctrine provides a uniform, No Set solution to the set-theoretic paradoxes.

Below, I consider two objections to limitation of size. The first claims that the limitation of size doctrine fails to respect the intuitive notion of an extension. The second claims that it leads to circular explanations.

4.5.1 Failure to respect the intuitive notion of an extension

Intuitively, a concept, F , has an extension if for every object, x , there is a determinate answer to the question of whether x falls under F . Intuitively, therefore, every determinate concept defines an extension. As we've seen, the limitation of size doctrine denies this. According to LOS, only small concepts have extensions. Big concepts do not. But why not? Can it be that big concepts fail to have extensions *because* they are true of too many things? I think not. If this sort of explanation appears at all plausible, I claim that this is because it forces an illicit shift from

¹⁵Although modern set theories, like ZF, are not generally thought of as representing the logical conception of *set*, some philosophers have argued that the axioms of these theories can also be seen as attempts to preserve as much of the logical conception of *set* as possible. Thus, Fraenkel et al. (1973, 32) write:

The axiom of [full] comprehension turned out to be inconsistent and therefore cannot be used as an axiom of set theory. However, since this axiom is so close to our intuitive concept of set we shall try to retain a considerable number of instances of this axiom schema . . . Our guiding principle, for the system ZF, will be to admit only those instances of the axiom schema of comprehension which assert the existence of sets which are not too “big” compared to sets which we already have. We shall call this principle *the limitation of size doctrine*.

the logical conception of *set* as an extension to the iterative conception of *set* as a collection. If we mix these conceptions, and think of extensions as collections of the objects falling under a concept, then we may find some plausibility in the claim that when the objects falling under a concept are too many, it is impossible for them to be collected together (by whatever means) so as to form an extension. But this plausibility is based on a conflation. Extensions as mappings exist, iff and because there is a determinate answer to the question of whether each object falls under the corresponding concept. The number of objects for which the answer to this question is yes is irrelevant.

But let us suppose that size does determine whether an extension exists and that this is an acceptable constraint of the logical conception. We might then wonder why size shouldn't also determine whether a concept exists. It seems arbitrary to accept concepts, no matter their size, but then to reject their extensions when the number of objects that would belong to them is too many. Furthermore, given our definition of extensions as mappings, there is a perfectly good sense in which even the extensions of small concepts have an absolutely infinite size; for every extension *involves* absolutely every object. Limitation of size only restricts the number of objects that can be *members* of an extension, not the number of objects that can be *involved* in an extension. What might justify such partiality?

A defender of the limitation of size doctrine might seek to meet this objection by modifying his view. According to the modified view, every concept has an extension, but there are two *kinds* of extension: *small extensions*—the extensions of small concepts—and *big extensions*—the extensions of big concepts. The difference between these two kinds has to do with the conditions under which they are members of other extensions. Small extensions behave in the way we expect extensions to behave: a small extension is a member of an extension, $\text{Ext}(F)$, just in case it falls

under the concept F . Big extensions behave differently: they are never members of anything. This amounts to replacing D2 and LOS with:

D2*. x is a *member* of $\text{Ext}(F) =_{df} Fx \wedge x$ is not a big extension

and

LOS* $(\forall F)(\exists e)(\forall x)[(x \text{ is not a big extension} \wedge x \in e \leftrightarrow Fx) \vee$
 $(x \text{ is a big extension} \wedge x \notin e)]$.

LOS* tells us that every concept has an extension. In particular, the concept *non-self-membered extension* has an extension. Ext tells us that this extension is unique (i.e., the concept *non-self-membered extension* has only one extension). We can therefore label it ' R '. However, the derivation of the contradictory $R \in R \leftrightarrow R \notin R$ goes through only if R is small. We refer by reductio that R is big.

To see this in a bit more detail, we might try running through the derivation (4.1)–(4.4) of Russell's paradox in SN, with D2* and LOS* replacing D2 and FC. (To save space, I use 'A' to abbreviate ' x is an extension $\wedge x \notin x$ '.) Begin by instantiating LOS* to the concept *non-self-membered extension* to get:

$$(5.1) (\exists e)(\forall x)[(x \text{ is not a big extension} \wedge x \in e \leftrightarrow Ax) \vee$$

$$(x \text{ is a big extension} \wedge x \notin e)].$$

As before, Ext then allows us to infer:

$$(5.2) (\forall x)[(x \text{ is not a big extension} \wedge x \in R \leftrightarrow Ax) \vee$$

$$(x \text{ is a big extension} \wedge x \notin R)].$$

But here the derivation comes to a halt. For (5.2) entails:

$$(5.3) (\forall e)(e \in R \leftrightarrow e \notin e)$$

only if every extension is small. In particular, it follows that:

$$(5.4) \quad R \in R \leftrightarrow R \notin R$$

only if R is small. We infer by reductio that R is big.

D2* and LOS* preserve the letter of the thesis that every concept has an extension, but not the spirit. Intuitively, $\text{Ext}(F)$ is a self-member iff $F(\text{Ext}(F))$. Indeed, this is the result of substituting ‘ $\text{Ext}(F)$ ’ for ‘ x ’ in D2. But we have replaced D2 and this is no longer true. To see how unintuitive the consequences are, consider the concept *self-identical*. Since $\text{Ext}(\textit{self-identical})$ is self-identical, $\text{Ext}(\textit{self-identical}) \in \text{Ext}(\textit{self-identical})$ according to D2. But *self-identical* is big (since everything is self-identical). So $\text{Ext}(\textit{self-identical}) \notin \text{Ext}(\textit{self-identical})$ according to LOS*. Because this version of the limitation of size doctrine denies that big extensions like $\text{Ext}(\textit{self-identical})$ are members of extensions whose intuitive membership criteria they satisfy, it fails to preserve the spirit of the thesis that every concept has an extension.

Perhaps it is for this reason that in his attempt to fix naive set theory, Boolos (1998, ch. 6) chooses to drop the name ‘extension’ entirely, replacing it with ‘subtension’. In Boolos’s adaptation of the limitation of size doctrine, every concept F has a subtension, $*F$, where $*F = *G$ iff either F and G are both big or F and G are coextensional. Membership in a subtension is defined by: $x \in y =_{df} (\exists F)(y = *F \wedge Fx)$. Subtensions of small concepts behave just as we would expect extensions to. If F is small, then $x \in *F \leftrightarrow Fx$. Subtensions of big concepts do not behave this way. If F is big, then $x \in *F$ does not imply Fx . Let F be the concept *self-identical* and let G be the concept of *ordinal number*. Since F and G are both big, $*F = *G$. Since Socrates is self-identical, $\text{Socrates} \in *F$. Therefore, $\text{Socrates} \in *G$. But Socrates is not an ordinal. Boolos’s theory preserves neither the letter nor the spirit of the thesis that every concept has an extension.

In chapter 2 (p. 55), I described a pragmatic solution to the paradoxes as one “which seeks to preserve a theory of sets without any underlying conception of *set*.” I argued that this sort of solution is unsatisfying because it leaves us unable to offer

explanations such as why there exist only those sets that our theory says exist and why none of the contradiction-inducing sets exists. The preceding objections show that the limitation of size doctrine provides only a pragmatic solution to the paradoxes. It enforces restrictions on logical set theory that are sufficient to restore consistency; but (as I have argued) these restrictions cannot be justified by the underlying logical conception of *set*. Despite his efforts to resuscitate logical set theory by means of the limitation of size doctrine, I think that Boolos would agree. After lamenting the inconsistency of naive set theory, which he describes as “simple to state, elegant, initially quite credible, and natural in that it articulates a view about sets that might occur to one quite naturally” (1998, 15), he characterizes the limitation of size doctrine quite differently (90).

Unlike the naive and the iterative conceptions, *limitation of size* ... is not a natural view, for one would come to entertain it only after one’s preconceptions had been sophisticated by knowledge of the set-theoretic antinomies, including not just Russell’s paradox, but those of Cantor and Burali-Forti as well.

(Note how the grounds for unnaturalness here cited call to mind Russell’s third condition on a satisfactory explanation (p. 54).) I conclude that because logical set theory can only be motivated by the logical conception of *set*, we can adopt the limitation of size doctrine as a fix to logical set theory only if we are willing to do so in a purely pragmatic spirit, with no reference to an underlying conception of *set*. This solution by fiat is philosophically unsatisfactory.

4.5.2 Circular explanations

As early as 1883, Cantor had determined that the totalities of cardinal and ordinal numbers were ‘absolutely infinite’ and could not be assigned any definite cardinal number, finite or transfinite. This position led to a conflict with Frege’s (1884) definition of ‘the cardinal number of F ’ as ‘the extension of the concept *equinumerous to*

F '.¹⁶ Under a replacement of ' F ' with 'the concept *ordinal number*', Frege's definition implies that the cardinal number of the concept *ordinal number* is the extension of the concept *equinumerous to the concept ordinal number*. Similarly, Frege's definition implies that the cardinal number of the concept *cardinal number* is the extension of the concept *equinumerous to the concept cardinal number*. Thus, Frege's definition implies that each of these concepts has a definite cardinal number. This contradicts Cantor.

In his critical review of Frege (1884), Cantor (1885) argued that Frege had overlooked the fact that some extensions have no definite cardinal number, but are "quantitatively indeterminate". The relevant passage is the following (Cantor here uses the terms 'number' and 'power' for 'finite cardinal' and 'transfinite cardinal'):

[Frege] entirely overlooks the fact that the 'extension of a concept' in general may be quantitatively completely indeterminate. Only in certain cases is the 'extension of a concept' quantitatively determinate: Then it has of course, if it is finite, a definite number or, in the case it is infinite, a definite power. For such a quantitative determination of an 'extension of a concept' the concepts 'number' and 'power' must already be given from another source, and it is a *reversal of direction* if one undertakes to found the latter concepts on that of 'extension of a concept'.¹⁷

The argument is not obvious; however, if we interpret 'quantitatively indeterminate' as 'absolutely infinite', the following plausible line of reasoning emerges.¹⁸

- (a) Cantor observes that in order to avoid false implications, such as that the concept *cardinal number* has a definite cardinal number, Frege's definition must be restricted so that it applies only to concepts that have definite cardinal numbers (or whose extensions do).

¹⁶In this discussion, I treat 'the cardinal number of F ' and 'the cardinal number of $\text{Ext}(F)$ ' as interchangeable.

¹⁷This translation is from Tait (1997).

¹⁸See Tait (1997) and Burgess (2008). For an argument that Cantor did not understand quantitative indeterminacy in this way, see Ebert and Rossberg (2009).

- (b) Although he does not call it such, the necessary restriction is a version of the limitation of size doctrine, according to which there is no extension of the form ‘Ext(*equinumerous to the concept F*)’, whenever F is quantitatively indeterminate.
- (c) Assuming that no concept is indeterminate in the sense of lacking a sharp boundary (see 4.1, pp. 89–91), quantitative indeterminacy is entirely a matter of lacking a definite cardinal number. Therefore, this limitation of size principle makes use of the very notion of ‘cardinal number’ that Frege is attempting to define. This is the “reversal of direction” that Cantor speaks of.

I believe that Cantor’s full objection may then be properly understood as a dilemma: Frege’s definition of ‘cardinal number’ is either inconsistent (if it is not amended by a limitation of size principle) or viciously circular (if it is).

Cantor’s argument is worrisome only insofar as one is sympathetic to Frege’s definition of the cardinal numbers as extensions of concepts, although it may be generalized so that it also applies to Russell’s definition of the cardinal numbers as classes of equivalent classes. However, recently, Øystein Linnebo (2010, 151–154) has articulated an argument that a limitation of size principle leads to circular explanations no matter how one defines the cardinal numbers. This variant is based on the notion of a minimal size—a “threshold cardinality”—for set formation: any things whose cardinality is at or above the threshold cardinality are too many to form a set. This leaves room for questions about what the threshold cardinality is and why it is what it is. Linnebo exploits these questions, arguing that attempts to answer them lead to circular explanations.

A standard answer to the first question is that the threshold cardinality is the cardinality of the plurality of all the ordinals ω . Linnebo assumes this identity and proceeds to argue that explanations for why it obtains are circular.

Consider the question why there are not more ordinals than oo . For instance, why cannot the plurality oo form a set, which would then be an additional ordinal, larger than any member of oo ? According to the view under discussion, the explanation is that oo are too many to form a set, where being too many is defined as being as many as oo . Thus, the proposed explanation moves in a tiny circle. The threshold cardinality is what it is because of the cardinality of the plurality of all ordinals, but the cardinality of this plurality is what it is because of the threshold.¹⁹

This circularity arises for any self-reproductive property ϕ , for the set of all ϕ s, if it existed, would entail the existence of an additional ϕ . In some cases, the set of all ϕ s would be an additional ϕ ; in other cases it would not be an additional ϕ , but would allow us to define it.²⁰ Thus, if we identify the threshold cardinality with the cardinality of the plurality of ϕ s and then say that there is no set of all ϕ s because of the number of ϕ s being what it is, our explanation will necessarily be circular. To illustrate, consider the self-reproductive property *set* and suppose we identify the threshold cardinality with the cardinality of the plurality of all sets ss (von Neumann, 1925). Finally, consider the question why cannot the plurality ss form a set—the universal set—which would be more inclusive than any member of ss ? If we answer that the plurality of ss cannot form a set because of the threshold cardinality being what it is, our explanation will be circular. For on the one hand, the threshold cardinality is what it is because of the cardinality of the plurality of all sets. But since the set of all sets, if it existed, would be an additional set, the cardinality of this plurality is what it is (at least in part) because of the threshold.

These circles may be broken if Cantor’s notion of the absolutely infinite is coherent. Equipped with this notion, we are in a position to define the threshold cardinality as

¹⁹In this passage, Linnebo identifies ordinals with sets of ordinals. But this is not necessary for his argument. It can be proved from the definition of ordinal numbers as measures of well-ordered sets that any set of ordinals starting from 0 (and containing no gaps) has an ordinal number greater than any of its members (see the proof of BF2 in the appendix to chapter 2). Consequently, the number of ordinals depends in part on whether O exists.

²⁰See my discussion of Russell’s schema, 3.2.

a size without any bound or limit, in a way that does not involve reference to any particular pluralities, such as the ordinals or the sets. While it may still be true that the ordinals or the sets are absolutely infinite and consequently that their size is that of the threshold cardinality, it will not be true that these pluralities *fix* or *determine* the threshold cardinality; at least not in the strong sense of making it what it is. If we are able to define the threshold cardinality independently of any particular plurality of objects, we can say that the threshold cardinality is what it is (and not more or less) because this is the greatest possible size, not because the cardinality of the ordinals or the cardinality of the sets is what it is (and not more or less). This breaks the explanatory circles above at the same step: we deny the proposition that the threshold cardinality is what it is because of the cardinality of the ordinals or the cardinality of the sets.

4.6 First-order set restrictivism

First-order set restrictivism (FSR) seeks to block the derivations of paradoxes in SN by placing contextual restrictions on the first-order quantifiers in FC and Ext. The effect is that big concepts always define small extensions, though the same concept may define different extensions in different contexts. There is a sense in which these restrictions allow the restrictivist to maintain the full comprehension axiom (FC). Interpreted within the context in which the restrictions are imposed, the statement ‘Any concept F determines the set of all the F s’ is true. The same statement will be false if interpreted within an absolutist context in which no quantificational restrictions are imposed. In the absolutist context, it will be true to say that FSR replaces FC with a new comprehension principle, in which the first-order quantifiers are always restricted. Indeed, it is this restriction which blocks the paradoxes. (It also raises issues for expressibility, which I discuss below.)

The contextual restrictions on the first-order quantifiers are expressed in FSR as restrictions on first-order domains of quantification as follows:

- ‘ $\forall_D x \phi(x)$ ’ abbreviates ‘everything ϕ s’, where ‘everything’ is interpreted as ranging over all objects in the domain D .

Expressions denoting extensions in FSR are indexed to the domain over which the extension is defined. Thus:

- ‘ extension_D ’ abbreviates ‘extension defined over D ’. (Similarly, ‘ $\text{Ext}_D(F)$ ’ abbreviates ‘ $\text{Ext}(F)$ defined over D ’, and ‘ e_D ’ abbreviates ‘ e defined over D ’.)

Contextual restrictions on the first-order quantifiers are always enforced. While domains vary between contexts, there is no context in which quantification over absolutely all sets is permitted. In particular, these contextual restrictions are applied to FC and Ext so as to yield the following interpretations:

$$\mathbf{FC}^* (\forall F)(\exists e_D)(\forall_D x)(x \in e \leftrightarrow Fx)$$

and

$$\mathbf{Ext}^* (\forall F)(\forall G)[\text{Ext}_D(F) = \text{Ext}_D(G) \leftrightarrow (\forall_D x)(x \in \text{Ext}(F) \leftrightarrow x \in \text{Ext}(G))].$$

Note that the outermost second-order quantifier is unrestricted. \mathbf{FC}^* tells us that every concept, F , has an extension_D , whose members are all and only those objects from D that fall under F . \mathbf{Ext}^* gives identity conditions for extension_D ; in particular, \mathbf{Ext}^* tells us that every concept defines a unique extension_D .

As I mentioned above, the same concept defines different extensions relative to different domains. This may be unsurprising in some cases. We might expect that the concept *horse* would have a different extension relative to the domain of yearlings—in which all and only yearlings are members—than it would relative to the domain of all horses—in which all and only horses are members. This may be more surprising

in other cases. It is somewhat surprising to think that the concept *horse* would have a different extension relative to the domain of all horses than it would relative to domain of all animals. After all, these extensions have all the same members! But given our definition of an extension as a mapping, this result seems to be the right one: the latter extension includes many argument-value pairs that the former does not. These extensions cannot really be identical after all.

In order to avoid Russell’s paradox, FSR must not recognize a universal domain, D_U . For if it did, we could derive FC and Ext in FSR as instances of FC* and Ext* by substituting D_U for ‘ D ’. (As we saw above, FC and Ext lead directly to Russell’s paradox.) In fact, Russell’s paradox can be derived even without a universal domain. All that is required is that there is *some* D that contains $\text{Ext}_D(\textit{non-self-membered extension})$, i.e., some D that contains the extension of *non-self-membered extension* that is defined over it. Fortunately, the FSRist is free to deny that this ever occurs. To see this in a bit more detail, we might try running through the derivation (4.1)–(4.4) of Russell’s paradox once more, but this time with FC* and Ext* replacing FC and Ext.

Letting ‘ D ’ stand for an arbitrary domain, the first three steps below—(6.1)–(6.3)—parallel (4.1)–(4.3) from the original derivation almost exactly. We begin by instantiating FC* to *non-self-membered extension* to get:

$$(6.1) \quad (\exists e_D)(\forall_D x)(x \in e_D \leftrightarrow x \text{ is an extension}_D \wedge x \notin x).$$

By (6.1) and Ext*, ‘ e_D ’ has a unique value, R_D , which is the extension _{D} whose members are all and only the non-self-membered extensions in D . This allows us to eliminate the existential quantifier, replacing the bound occurrence of ‘ e_D ’ with ‘ R_D ’:

$$(6.2) \quad (\forall_D x)(x \in R_D \leftrightarrow x \text{ is an extension}_D \wedge x \notin x).$$

(6.2) entails:

$$(6.3) \quad (\forall_D e_D)(e_D \in R_D \leftrightarrow e_D \notin e_D).$$

But here the derivation comes to a halt. If R_D was in the range of \forall_D we could instantiate ‘ e_D ’ to R_D in (6.3) to get the contradictory: $R_D \in R_D \leftrightarrow R_D \notin R_D$. However, since \forall_D is restricted, this assumption is open to question. We conclude by reductio that R_D lies outside the range of \forall_D ; consequently that:

$$(6.4) \quad R_D \notin R_D.$$

If we want, we can shift to a more inclusive context, in which quantification ranges over a more inclusive domain D^+ , which includes R_D . In this context, the result of instantiating FC^* to *non-self-membered extension* is:

$$(7.1) \quad (\exists e_{D^+})(\forall_{D^+}x)(x \in e_{D^+} \leftrightarrow x \text{ is an extension}_{D^+} \wedge x \notin x).$$

By the same reasoning as before, ‘ e_{D^+} ’ has a unique value, R_{D^+} , which is the extension_{D^+} whose members are all and only the non-self-membered extensions in D^+ . However, since \forall_{D^+} is restricted, we cannot instantiate to R_{D^+} to derive the contradictory $R_{D^+} \in R_{D^+} \leftrightarrow R_{D^+} \notin R_{D^+}$. Instead, we conclude by reductio that R_{D^+} lies outside the range of \forall_{D^+} and consequently that $R_{D^+} \notin R_{D^+}$.

Burali-Forti’s paradox is blocked in precisely the same way. FSR allows us to define sets of the form $\text{Ext}_D(\text{ordinal number})$. But the existence of each of these sets leads to a contradiction only if it is in the D over which it is defined. To see this, recall my discussion of Burali-Forti’s contradiction (2.3.1 and 3.2.3). In these chapters, I described the contradiction as based on the fact that the ordinal number of the initial segment defined by any set of ordinal numbers is greater than any member of the set. But the contradiction might equally well be described in terms of domains. Simply replace ‘set’ with ‘domain’ above and Burali-Forti’s contradiction is based on the fact that the ordinal number of the initial segment defined by any *domain* of ordinal numbers is greater than any member of the *domain*. This leads to a contradiction if this ordinal number must also be a member of the domain; but not if we are free to infer that it is not a member of the domain. Since according to FSR, D never

includes all the ordinals, we are always free to infer that the ordinal number of the initial segment defined by any domain of ordinal numbers lies outside the domain. Hence, there is no paradox.

Cantor's paradox is also blocked. FSR allows us to define sets of the form $\text{Ext}_D(\textit{cardinal number})$, $\text{Ext}_D(\textit{object})$, $\text{Ext}_D(\textit{extension})$ and $\text{Ext}_D(\textit{pure extension})$ but none of these are contradiction-inducing. To see this, recall the derivations of contradictions from the sets K , U , V and V_P (2.2.2 and 3.2.2). Consider first, the derivation of a contradiction from K . This is based on the fact that given any set of cardinals, we are able to define a cardinal number which is provably greater than any member of the set. But the contradiction might equally well be described in terms of domains. So understood, the contradiction is based on the fact that given any domain of cardinals, we are able to define a cardinal number which is provably greater than any member of the domain. This leads to a contradiction if this cardinal number must also be a member of the domain; but not if we are free to infer that it is not a member of the domain. Since according to FSR, D never includes all the cardinals, we are always free to infer that the cardinal number defined over any domain of cardinal numbers lies outside the domain.

Next, consider the derivations of contradictions from U , V and V_P . These are based on the fact that we are able to define bijections between these sets and their powersets, which contradicts Cantor's theorem.²¹ Of course, the definition of each bijection involves quantification over the members of the corresponding set. In the language of FSR, this means that it presupposes the existence of a domain, D , whose members include all members of the set. Moreover, since each of these sets necessarily includes every set of its members as a member (including the set of all its members)—for every set of objects is an object, every set of sets is a set and every set of pure

²¹See my discussion of Cantor I, pp. 36–38

sets is a pure set—it follows that any domain over which these sets are defined must include every set of its members as a member. For whenever we introduce a set by means of FC^* , we quantify over all its members. If such a D exists, we can derive the contradictions as before. To illustrate, let $D-$ be any restriction on the domain D and suppose that D contains both $\text{Ext}_D(\text{object})$ and every set of the form $\text{Ext}_D(\text{object})$. Such a D would include all its “subsets” as members. This would enable us to define a bijection between $\text{Ext}_D(\text{object})$ and its “powerset”, $\mathcal{P}(\text{Ext}_D(\text{object}))$, which would contradict Cantor’s theorem as applied to D . FSR blocks this derivation by denying that such a domain exists. Consequently, any attempt to define the necessary bijection between $\text{Ext}_D(\text{object})$ and $\mathcal{P}(\text{Ext}_D(\text{object}))$ fails. This amounts to a variant of the restrictivist solution to Russell’s paradox above. Just as in that case any attempt to derive the contradictory $R_D \in R_D \leftrightarrow R_D \notin R_D$ only succeeds in showing that R_D lies outside of D , so in this case any attempt to establish a bijection between $\text{Ext}_D(\text{object})$ and $\mathcal{P}(\text{Ext}_D(\text{object}))$, which would contradict Cantor’s theorem, only proves that some “subsets” of $\text{Ext}_D(\text{object})$ lie outside D . This means that there can be no domain of *all* objects, *all* sets or *all* pure sets. For any of these domains would—by definition—include all sets of their members as members. Similar remarks apply to $\text{Ext}_D(\text{extension})$ and $\text{Ext}_D(\text{pure extension})$. Again there is no paradox.

4.6.1 A modalized variant of FC

In the interests of providing a foundation for iterative set theory, Charles Parsons (1983, chaps. 10, 11) introduces a modalized variant of full comprehension, according to which any property of sets *can* determine a set, which exists at later stages in the

hierarchy.²² Using second-order variables for properties and first-order variables for sets, it is tempting to formulate (the natural necessitation of) this idea as:

$$\mathbf{FC}^\diamond \quad \Box(\forall F)\diamond(\exists y)(\forall x)(x \in y \leftrightarrow Fx).$$

However, \mathbf{FC}^\diamond is inconsistent. In terms of the usual possible worlds semantics, \mathbf{FC}^\diamond says that at any world, w , for every property F at w : $\diamond(\exists y)(\forall x)(x \in y \leftrightarrow Fx)$. This leads to Russell’s paradox when ‘ F ’ is instantiated to the property *non-self-membered set*. The inconsistent instance reads:

$$(8.1) \text{ At } w: \diamond(\exists y)(\forall x)(x \in y \leftrightarrow x \notin x).$$

Applying the standard possible worlds interpretation to the \diamond , (8.1) says that at w there is an accessible world v at which the Russell set $R_v = \{x|x \notin x\}$ of all the non-self-membered sets in v exists. Because R_v is one of these sets, it follows that:

$$(8.2) \text{ At } v: R_v \in R_v \leftrightarrow R_v \notin R_v.$$

The source of the trouble with \mathbf{FC}^\diamond appears to be quantification into the scope of \diamond . (The quantifier $\forall F$ occurs outside of the \diamond operator, while the bound occurrence of the variable ‘ F ’ falls within it.) In response, Parsons (pp. 295, 315–318) suggests “fully rigidifying” properties relative to worlds. F is *rigidified* relative to a world w if anything that is F at w is necessarily F and anything that is not F at w is necessarily not F . F is *fully rigidified* relative to a world w if in addition, there are no F s that do not exist at w , i.e., the F s at w are all the possible F s.²³ When F is fully rigidified relative to a world w , something counts as an F at another world v iff it is an F at w . \mathbf{FC}^\diamond is consistent if F is (fully) rigidified relative to the world w

²²Parsons is motivated by the thought, which he traces back to Cantor (1899), that any “multiplicity” of sets that “can exist together” can form a set. For discussion of another modal variant of full comprehension, according to which any property F determines the set of all things that are F in a special way, see Fritz et al. (nd).

²³See Parsons (pp. 288, f.n. 29; 301–302) for a definition of ‘the (full) rigidification of F ’.

to which \Box is instantiated. The result is that the occurrence of ‘ F ’ in ‘ Fx ’ is read as falling outside the scope of \Diamond so that instead of referring to the things that are F at the world v to which \Diamond is instantiated, it refers to the things that are F at w . Parsons’s rigidification strategy is striking since it blocks derivations of Russell’s paradox in a way that closely resembles FC*.

Let ‘ F_w ’ symbolize the (full) rigidification of F to w and let ‘ $\{x|F_w x\}$ ’ denote the possible set—one whose existence is asserted within the context of the \Diamond operator—that is defined by “instantiating” FC $^\Diamond$ to the world w and the (fully) rigidified property F_w . The result of replacing the occurrence of ‘ \notin ’ in (8.1) with its (full) rigidification to w is:

$$(9.1) \text{ At } w: \Diamond(\exists y)(\forall x)(x \in y \leftrightarrow x \notin_w x).$$

In terms of the usual possible worlds semantics, (9.1) says that at w there is an accessible world v at which the Russell set $R_w = \{x|x \notin_w x\}$ of all the non-self-membered sets in w exists. (9.1) does not say that R_w exists at w . So (9.1) does not entail the contradiction: At $w: R_w \in R_w \leftrightarrow R_w \notin R_w$. Of course, we might apply FC $^\Diamond$ a second time (“instantiating” to v instead of w) to define the Russell set R_v :

$$(9.2) \text{ At } v: \Diamond(\exists y)(\forall x)(x \in y \leftrightarrow x \notin_v x).$$

In terms of the usual possible worlds semantics, (9.2) says that at v there is an accessible world u at which the Russell set $R_v = \{x|x \notin_v x\}$ of all the non-self-membered sets in v exists. We cannot derive a contradiction from (9.2), however, for the same reason that we cannot derive a contradiction from (9.1). Just as R_w does not exist at w , R_v does not exist at v , but rather at some v -accessible world, u .

The role played by possible worlds here is very much like that played by quantificational contexts under FSR. While it is always possible to define a Russell set relative to any context (or world), there is no context (or world) that contains *all* the non-self-membered sets (that exist at all possible worlds). This allows us to infer that every

Russell set is outside of the context (or world) in terms of which it is defined, which is how paradox is avoided. In light of these similarities, it seems that a modal set theory based on FC^\diamond might constitute another type of set restrictivism (in addition to FSR). Indeed, Parsons (1983, 288–293) endorses a variant of restrictivism, according to which claims about *all* sets are “systematically ambiguous” in the sense that ‘all’ is interpreted as ranging over the sets at some stage in the iterative hierarchy, though it is ambiguous which stage this is. But even if a modal theory based on FC^\diamond counts as a type of restrictivism, it is unclear whether it can be understood as a development of the logical conception of *set*, or rather whether it can only be understood, as Parsons intends, as a development of the iterative conception of *set*.

4.6.2 The All-in-One principle

In chapter 3, I argued that the doctrine of self-reproductive properties (SRP) provided a plausible interpretation of Russell’s 1906 attribution of the paradoxes to “self-reproductive processes or classes.”

SRP If ϕ is self-reproductive, then there is no set of all ϕ s.

Interpreted in light of Russell’s remarks two years later—“When I say that a collection has no total, I mean that statements about *all* its members are nonsense” (Russell, 1908, 63)—SRP tells us that self-reproductive properties imply restrictivism; in particular, that self-reproductive properties such as *non-self-membered extension* imply set restrictivism.

In chapter 1, I argued that any restrictivist argument of this sort is committed to the All-in-One principle, which I argued is most likely false. We might wonder whether FSR is also committed to the All-in-One principle. The suspicion that it may be so committed is fueled, in part, by the fact that every occurrence of the first-order quantifier \forall in FSR is associated with a domain of quantification. This establishes half of the All-in-One principle, viz., that quantification is always associated with

a domain of quantification (P1 from chapter 1). The other half of the All-in-One principle (P2 from chapter 1) is the claim that domains are sets (or set-like). Are they? It would be highly problematic if they were, and for reasons having to do with explanation that are independent of the All-in-One principle. FSR is a theory of sets which employs domains as interpretive parameters on its principles governing set existence and set identity. If domains were sets, then a prior understanding of sets would be needed to understand the principles governing set existence and set identity. Such circularity in the order of explanation appears vicious.

Fortunately, there is no need to treat domains as sets, or as singular objects of any sort. We are free to interpret ‘ D ’ as a plural term, referring plurally to all the objects over which \forall ranges. In other words, we are free to interpret reference to domains in FSR as reference to the objects “in” these domains and quantification over domains as plural quantification over objects “in” domains. Thus, we might speak of an object over which \forall_D ranges as being among the dd instead of being in D . Accordingly, we might read ‘ $\forall_{dd}x\phi(x)$ ’ as “everything ϕ s,” where ‘everything’ is restricted by context to quantification over objects among dd . And we might rewrite the FSR axioms accordingly:

$$\mathbf{FC}^* (\forall F)(\exists e_{dd})(\forall_{dd}x)(x \in e_{dd} \leftrightarrow Fx)$$

and

$$\mathbf{Ext}^* (\forall F)(\forall G)[\text{Ext}_{dd}(F) = \text{Ext}_{dd}(G) \leftrightarrow (\forall_D x)(x \in \text{Ext}(F) \leftrightarrow x \in \text{Ext}(G))].$$

The reason why we are free to interpret FSR in this way has to do with the role domains play in FSR: they restrict quantification. In order to play this role, they must allow us to refer to the objects over which a quantifier ranges. In order to do this, it is not necessary for these objects to be the members of a singular objects such as a set or class.

Still, it might be thought that if not FSR itself, then certainly FSR’s solutions to the paradoxes require identifying sets with domains. For example, in my remarks on the Burali-Forti paradox, I claimed that if ‘set’ were uniformly replaced with ‘domain’, the derivation of Burali-Forti’s contradiction would remain intact. Doesn’t this show that ‘domain’ and ‘set’ are simply two words for the same thing? I think that it does not. The reason is that the derivation of Burali-Forti’s contradiction also remains intact under a plural reconstruction. (In chapter 7, I argue that plural reference and plural quantification over ordinals is enough to generate the paradox.) This means that we are free to interpret domain-talk plurally. Under a plural reconstruction, Burali-Forti’s contradiction is based on the fact that the ordinal number of the initial segment (which is understood as a plurality) defined by any plurality of ordinal numbers is greater than any member of the plurality. This version of the paradox cannot be blocked by placing restrictions on sets, for it is not based on the existence of a set of all ordinals, but on the existence of a plurality of all ordinals. However, it is blocked by FSR, given a plural interpretation of domains. The paradox is blocked because there is no domain—which in this context must be thought of as a plurality—that contains all ordinals. Similar remarks apply to my discussion of the other paradoxes. I conclude that FSR has no commitment to domains as singular objects and therefore has no commitment to the All-in-One principle.

4.6.3 Explanation and Expressibility

FSR blocks the paradoxes by restricting quantification so that there is no D containing all instances of a big concept. Insofar as this restriction on quantification amounts to a restriction on set-existence—there can be no set of all F s whenever it is impossible to quantify over all F s—the FSRist appears committed to a variant of the limitation of size doctrine. According to this variant, it is quantification rather than FC, that is the immediate target. This is problematic; for it seems that FSR

should then be subject to the same criticism, levied against the limitation of size doctrine above, that the restriction imposed on quantification is unmotivated by the logical conception of *set* and can provide only a pragmatic solution to the paradoxes. I think, in fact, that there is an important difference between quantification and existence and that limitations on quantification do not call out for explanations in the way that limitations on existence do. However, even granting this, there is a more fundamental problem. Not only does the FSRist appear unable to *explain* her restriction, she appears unable to *state* it. And this is surely something that we should require her to do.

To illustrate, let ‘F’ stand for an arbitrary big concept. A statement of the quantificational restriction imposed by FSR is:

$$(10) \neg(\exists D)(\forall x)(Fx \rightarrow x \text{ is in } D),$$

or, if domains are pluralities:

$$(10^*) \neg(\exists dd)(\forall x)(Fx \rightarrow x \text{ is among } dd).$$

However, (10) (or (10*)) conveys the intended thought only if \forall ranges over absolutely all *F*s, precisely what is forbidden in FSR.

Of course, the fact that (10) cannot be expressed *in* FSR does not mean that (10) cannot be expressed at all. Indeed, it might be pointed out that the inexpressibility of (10) in FSR should be expected. After all, (10) quantifies directly over FSR’s domains of quantification. This indicates that (10) belongs to FSR’s semantic theory. And, it is widely accepted that the semantic theory for a formal theory is only expressible *outside* of the theory. Perhaps, then, (10) is expressible in a metalanguage for FSR. Quantification in this metalanguage must be unrestricted; but that is no problem, for acceptance of FSR doesn’t commit one to the much stronger claim that unrestricted quantification over sets occurs in no language.

Unfortunately, things are not so easy. For consider whether FC, expressed in FSR's metalanguage, is true. If quantification in the metalanguage is unrestricted, we can interpret the occurrences of D in FC as D_U , the universal domain. Since this allows us to define the contradiction-inducing sets and leads to contradictions, we must say that FC is false when expressed in FSR's metalanguage. But how can this be if anything like the logical conception of set is true? If we want to defend FSR, we must find another way.

CHAPTER 5

EXPRESSING RESTRICTIVISM

Inexpressibility is not unique to set restrictivism. It is taken by many authors to be *the* challenge facing restrictivists in general. As David Lewis (1991, 68) puts it, the restrictivist, “violates his own stricture in the very act of proclaiming it!” Lewis’s challenge may be put as follows: if restrictivism is true, any attempt to state it must fail.

Timothy Williamson (2003, 428) presents the problem in this way: the restrictivist is committed to quantified claims with contradictory truth conditions, for example:

(1) I am not quantifying over everything,

or equivalently:

(2) Something is not being quantified over by me.

Given two plausible semantic principles governing quantification and assertion, (2) can be shown to have a contradictory truth condition. Let S be any speaker. The first semantic principle states that S truly utters a sentence of the form ‘something F s’ iff something over which S quantifies satisfies the predicate expression ‘ F s’. In the present case, ‘ F s’ is the predicate expression ‘is not being quantified over by me’ and so the first principle tells us that S truly utters (2) iff something over which S quantifies satisfies ‘is not being quantified over by me’. The second semantic principle states that something satisfies the predicate expression ‘ F s’ as uttered by S iff it F s. In the present case, this second principle tells us that something satisfies ‘is not being quantified over by me’ as uttered by S iff it is not being quantified over by S .

Taken together, these two principles tell us that S truly utters (2) iff something over which S quantifies is not being quantified over by S .¹ Since this truth condition is contradictory, (2) cannot be true.

5.1 Semantic ascent

The fact that (1) and (2) have contradictory truth conditions is simply a consequence of the fact that they involve the very sort of quantification that they deny. In the case of set restrictivism, I argued that this problem cannot be solved by shifting to a metalanguage. My argument there appealed to special considerations involving the logical conception of set which don't apply in the case of restrictivism simpliciter. However, there are more general reasons that seem to show that such a strategy of semantic ascent cannot succeed in expressing restrictivism of any sort.

Suppose the restrictivist attempts to express her view by ascending from her object language L to a suitable metalanguage, L^* . In L^* , she asserts:

(3) Something is not being quantified over by me in L .

The sort of contradiction noted above is avoided, since the occurrences of 'something' and 'not being quantified over by me in L ' are evaluated relative to different languages with different domains. The occurrence of 'something' is evaluated relative to L^* , while the occurrence of 'not being quantified over by me in L ' is evaluated relative to L .

The problem with this strategy is that it is too weak to say what needs to be said. For the restrictivist is committed to the claim that unrestricted first-order quantification occurs in *no* language. (Likewise, set restrictivism implies that unrestricted

¹For ease of exposition, I have left out time indices. With time indices, the first principle reads: S truly utters a sentence of the form 'there is something that Fs ' at t iff something over which S quantifies at t satisfies ' Fs ' at t . And the second principle reads: something satisfies the predicate expression ' Fs ' as uttered by S at t iff it Fs at t .

first-order quantification over sets occurs in *no* language.) (3) leaves open the possibility that quantification in L^* is unrestricted. Of course, the restrictivist might plug this hole by invoking a third language, L^{**} , and say in L^{**} :

- (4) Something is not being quantified over by me in L^* .

But this only pushes the problem back, for it leaves open the possibility that quantification in L^{**} is unrestricted.

The obvious way of overcoming this problem is to employ a statement that quantifies over all languages. We might try:

- (5) For any language L_0 , there is a language L_1 such that L_0 is restricted relative to L_1 .

But statements such as (5) are problematic for another reason: they are too strong. Timothy Williamson (2003) and Kit Fine (2006b) give different arguments that purport to establish this. Williamson (2003, 429–430) argues that the strength of statements like (5) can be shown to entail statements with contradictory truth conditions. Fine (2006b, 27–28) argues that the strength of such statements enables us to introduce unrestricted quantification and are therefore self-defeating.

5.1.1 Williamson

Using variation between contexts in a natural language instead of variation between formal languages, Williamson considers expressing restrictivism as:

- (6) For any context C_0 , there is a context C_1 such that not everything that is quantified over in C_1 is quantified over in C_0 .

Unfortunately, as he shows, (6) cannot be true. Without loss of generality, we can assume that the restrictivist asserts (6) in the context C . In C , (6) entails:

- (6C) There is a context C_1 such that not everything that is quantified over in C_1 is quantified over in C ,

which in turn entails:

(7) Not everything is quantified over in C .

Assuming that commitment in a context is closed under entailment, it follows that by asserting (6) in C , the restrictivist is committed to assert (7) in C . (7) may be true if uttered in another context, but it cannot be true in C . This can be proved by adapting the two semantic principles to contexts in the obvious way, and then using them to show that (7) has a contradictory truth condition like (2). In the context C , the first principle reads: An assertion of the form ‘not everything F s’ is true as uttered in C iff something that is quantified over in C does not satisfy the predicate expression ‘ F s’. The second principle reads: something that is quantified over in C does not satisfy the predicate expression ‘ F s’ iff it does not F . Taken together, these two principles tell us that (7) is true as uttered in the context C iff something that is quantified over in C is not quantified over in C . Since this truth condition is contradictory, (7) cannot be true. And so (6) cannot be true either.

In response, Williamson considers combining quantification over contexts with semantic ascent:

(8) For every context C_0 , there is a context C_1 such that ‘Not everything is quantified over in C_0 ’ is true as uttered in C_1 .

This blocks the argument that (8) has a contradictory truth condition. To see this, try applying the argument to (8). Without loss of generality, we can assume that the restrictivist asserts (8) in C . In C , (8) entails

(8C) There is a context C_1 such that ‘Not everything is quantified over in C ’ is true as uttered in C_1 ,

which in turn entails

(9) ‘Not everything is quantified over in C ’ is true as uttered in C_1 .

(9) has a consistent truth condition, for, intuitively, (9) is true iff something that is quantified over in C_1 is not quantified over in C . The reason (9)’s truth condition is consistent is that the occurrence of ‘everything’ is not interpreted according to the context of utterance, C , but according to C_1 , which we might call the context of interpretation. Notice that no matter what context of utterance we might instantiate (8) to, we are free to hold that the context of interpretation is different. By contrast, in (6) the occurrence of ‘everything’ is interpreted according to the context of utterance. This is what underwrites the first semantic principle according to which something that *is quantified over in C* does not F iff the utterance in C of ‘not everything F s’ is true. This principle does not hold for (8), since the occurrence of ‘everything’ in (8) is tagged to a context outside the context of utterance.

Unfortunately, this solution raises its own problems concerning the ability of the speaker to understand what the truth of the quoted sentence—‘Not everything is quantified over in C_0 ’—amounts to in contexts other than his own. As Williamson (2003, 430) points out: “If the speaker does not know what the sentence ‘Not everything is quantified over in C_0 ’ expresses in another context, then the claim that it is true as uttered in that context is not very helpful.”

5.1.2 Fine

Fine (2006b, 26) considers expressing restrictivism as the thesis that necessarily, any interpretation, I , of the first-order quantifier \forall has an expansion, J , from which it follows that there can be no maximal interpretation. He argues that this attempt to express restrictivism fails since it is possible to join any number of interpretations I, J, K, \dots to form a new *sum interpretation* whose domain includes everything in any of the domains of I, K, J, \dots . Since there is no limit to the number of interpretations that can be joined, unrestricted quantification over interpretations leads directly to the existence of a maximal interpretation, whose domain includes

absolutely everything. Before discussing the argument in more detail, I will say a bit about interpretations.

Intuitively, interpretations of the first-order quantifier (henceforth: simply interpretations) are functions, whose value, $I(\forall)$, is the domain of \forall under I . (Keep in mind that domains needn't be understood as singular objects and so the interpretation function can be viewed as a multi-valued function from \forall to the objects over which \forall ranges under I —the *dd* of \forall under I . Nevertheless, I will help myself to singular domain talk wherever grammatically convenient.) However, the identification of interpretations with functions raises an immediate worry. On the standard analysis, functions are sets of ordered pairs: each pair consisting of an argument and its value. Thus, under the standard analysis, any interpretation, I , would be identified with the ordered pair $\langle \forall, I(\forall) \rangle$. Again, on the standard analysis, ordered pairs are defined as sets; the standard definition being Kuratowski's: $\langle x, y \rangle =_{df} \{\{x\}, \{x, y\}\}$. Applying the Kuratowski definition to I yields: $I = \langle \forall, I(\forall) \rangle = \{\{\forall\}, \{\forall, I(\forall)\}\}$.

This analysis is inappropriate in the present context since it implies that absolutism is committed to the existence of a universal set. This misconstrues the quantificational question of whether there is a maximal interpretation with the set-theoretic question of whether there is a universal set. In particular, if I is a maximal interpretation, then $I(\forall)$ is the plurality of absolutely all things, which means that $\{I(\forall)\}$ is the universal set. Moreover, since \forall is one of the things included in the plurality of absolutely all things, the plurality of \forall and $I(\forall)$ is identical to $I(\forall)$. Consequently, according to the Kuratowski analysis, I is a set containing the universal set as a member.

To avoid this problem, one might replace the standard analysis of functions as sets of ordered pairs with an “ontologically innocent” plural analysis. Under this analysis,

functions are paraphrased away using plural quantification (and possibly mereology).² It follows that the only ingredients needed to define an absolute interpretation are (i) the argument, \forall , (ii) its plural value, the plurality of absolutely all things, and (iii) plural quantification. While absolutism is committed to the existence of an all-encompassing plurality, it is not committed to the existence of a universal set. Alternatively, functions may be taken as primitive as in Oliver and Smiley (2013). In what follows, I will assume that interpretations are viewed in one of these two ways.

Interpretation variables play the role of the domain variables in FSR. They are indexed to the first-order quantifiers, $\forall_I/\exists_I, \forall_J/\exists_J, \forall_K/\exists_K$, and the expression ‘ $\exists_I x\phi(x)$ ’ is read as saying that there is some x under the interpretation I for which $\phi(x)$. Compare to FSR, in which domain variables are indexed to the first-order quantifiers, \forall_D/\exists_D , and ‘ $\forall_D x\phi(x)$ ’ abbreviates ‘everything ϕ s’, where ‘everything’ is interpreted as ranging over all objects in the domain D .

For any interpretations I and J , J is an *expansion* of I if the things in the domain of \forall under I are properly among the things in the domain of \forall under J .³

D1. J is an expansion of $I =_{df} I(\forall) \prec J(\forall)$.

To illustrate: since all the dogs are properly among all the animals, the interpretation that assigns all the animals to \forall is an expansion of the interpretation that assigns all the dogs to \forall . Assuming an infinite universe, restrictivism can then be formulated as the thesis that every interpretation has an expansion:

(10) $(\forall I)(\exists J)(J \text{ is an expansion of } I)$.

²In the Appendix to Lewis (1991), John Burgess, Alan Hazen and David Lewis show how to replace the standard analysis with first-order plural quantification and mereology. Hazen (1997) replaces the mereological component from Lewis (1991) with higher-order plural quantification. For more on the resources and limitations of higher-order plural quantification, see Rayo (2006).

³The xx are properly among the yy iff every one of the xx is one of the yy , but not visa versa.

(10) conveys the intended thought only if it quantifies over all interpretations. (If it does not, (10) cannot rule out the existence of an absolute interpretation and is therefore too weak to express restrictivism.) But, Fine (2006b, 27) contends that any interpretations I_1, I_2, \dots can be joined to form a *sum interpretation*, J , where $\exists_J x\phi(x)$ iff $\exists_I x\phi(x)$ for some I in I_1, I_2, \dots . Since quantification over interpretations is unrestricted, it follows that we can join all interpretations to form an absolute interpretation J , whose domain includes absolutely everything. (10) is therefore self-defeating: by quantifying over all interpretations—which is necessary in order to convey the intended thought—(10) smuggles in absolutely unrestricted quantification. Fine is not alone in this observation. Michael Glanzberg (2004, 560) considers formulating restrictivism along the lines of (10) as the thesis that “for any domain of quantification that can be specified in interpretation, there can be specified a wider one.” He then goes on to observe that:

If we could ‘read through’ talk about domains to talk about the objects in them, then we might see talk about all domains or all interpretations as talk about whatever falls under them.

To make the argument a bit more precise, I introduce the notion of a *summation principle* for interpretations. Let ‘ α ’ be a schematic variable for ordinal numbers. A summation principle says that given a fixed number of interpretations, I_1, \dots, I_α , there exists a sum interpretation, J , where $J(\forall) = I_1(\forall) \cup \dots \cup I_\alpha(\forall)$. There is a family of summation principles, of various strengths. Weaker principles allow for the summation of fewer interpretations; stronger principles allow for the summation of more interpretations. The argument requires a maximally strong summation principle, which allows for the summation of *any* number of interpretations. I will refer to this as “the absolute summation principle” (ASP).

ASP. $(\forall I_0)(\forall I_1) \dots (\exists J)(J(\forall) = I_0(\forall) \cup I_1(\forall) \cup \dots)$ ⁴

⁴I assume that the ellipses can be filled in by an infinite number of interpretations.

The argument that (10) is self-defeating is derived from ASP as follows. Instantiate ASP to all interpretations to derive their sum interpretation, J_{ALL} . Since absolutely every object is included in the domain of some interpretation, absolutely every object is among the $dd = J_{\text{ALL}}(\forall)$. Consequently, J_{ALL} is an absolute interpretation.

To block this argument, the restrictivist might advocate replacing ASP with a weaker summation principle that limits the number of interpretations that can be summed. The limitation may be imposed by introducing a cut-off cardinality for interpretation summation, κ , defined as the least number of interpretations that are too many to sum. This determines not just one, but a family of weaker summation principles. The weaker summation principles can then be expressed by the following “weak summation principle schema” (WSPS):

$$\text{WSPS. } |\alpha| < \kappa \rightarrow (\forall I_0) \dots (\forall I_\alpha)(\exists J)(J(\forall) = I_0(\forall) \cup \dots \cup I_\alpha(\forall))$$

No instance of WSPS can be instantiated to all interpretations provided the number of interpretations is equal to or greater than κ .

This response is not very satisfying. What grounds can there be for placing limits on the number of interpretations that can be joined if no limits are to be placed on the number of interpretations that can be quantified over? One possible answer is that summation is a sort of mental activity: one must join interpretations one-by-one, and there is a limit to how far this can be carried out. But such a constructivist approach to summation is highly implausible. That there exists an interpretation whose domain includes all dogs-and-cats, given that there is an interpretation whose domain includes all dogs and an interpretation whose domain includes all cats seems completely independent of any mental activity.

Fine seeks to replace (10) with a modal variant of restrictivism he calls “expansionism”. Expansionism involves a special *postulational* modality modeled on the self-reproductive properties discussed in chapter 3. Russell’s paradox may be taken to show that the property *set* is self-reproductive since it is always possible to expand

a plurality of sets by the addition of a new set—the set of all non-self-membered sets among the plurality—that is provably not among them. The addition of this new set is a postulational possibility relative to the original plurality of sets. Expansionism is the view that any interpretation I of the quantifier \forall can be extended by the postulation of new objects which do not exist according to I but do exist according to I 's expansion. Using \Box and \Diamond for postulational necessity and postulational possibility, expansionism can be expressed as:

$$(11) \Box\forall I\Diamond\exists J(J \text{ is an expansion of } I).$$

There are two problems with (11). First, it is unclear whether the postulational modality involved is an intelligible notion. (I take up this challenge in next chapter, when I discuss modal interpretations of set theory.) Second, assuming postulational modality is an intelligible notion, it is unclear whether the view that (11) expresses should count as restrictivist. In chapter 1 (p. 15) I defined restrictivism as the view that “absolutely unrestricted quantification is impossible, or incoherent, but not because semantic indeterminism, or conceptual relativism, or mysticism is true.” Later, (p. 27) I argued that the ontological thesis that there is no plurality of all things is not properly restrictivist: the restrictivist “does not deny that there are some things that are all the things. What he denies is that it is possible to quantify over all the things there are.” Yet, some of the things Fine says when discussing postulational possibility suggest either a conceptual relativism along the lines of Carnap and Putnam, or an ontological conception of indefinite extensibility, such as that expressed by Spencer (2012) and Yablo (2006).

Fine (2006b, 40) suggests that the possibility of extending any given interpretation of \forall is not due to any failure on our part to talk about all that there is, but rather to a failure on the world's part. This suggests that it is the indefinite extensibility of the universe, not the limitations of language or logic that prevent us from quantifying unrestrictedly. He goes on to describe his view as follows:

We should bear in mind that, on the present view, there is no such thing as *the* ontology, one that is privileged as genuinely being the sum-total of what there is. There are merely many different ontologies, all of which have the same right (or perhaps we should say no right) to be regarded as the sum-total of what there is.

This sounds strikingly like the sort of conceptual relativism espoused by Carnap (1950) and Putnam (1977, 1981, 1987). Consider once more Putnam’s thought experiment. How many objects are there in a world containing three mereological atoms? Relative to a conceptual framework that counts arbitrary mereological sums, the answer is seven (or eight, if we include the null-object). Relative to a conceptual framework that counts only atoms, the answer is three. But (we are invited to infer) neither of these frameworks can claim to be the true framework.

Given these rather sobering difficulties, the restrictivist might feel forced to abandon the project of expressing her view and choose instead to describe herself as raising a challenge to her absolutist opponent. Tim Button (2010, 391–392) recommends this sort of approach, which he calls “militant quietism”:

Restrictivism should not be thought of as a positive doctrine, but rather as a form of *militant quietism* . . . The restrictivist should simply pick a fight: *Give me a sentence which you think quantifies over absolutely everything; engage with me in conversation for a while; and by the time you leave, I will have convinced you that you failed—by your own lights—to quantify over absolutely everything.*

Alternatively, one might seek to express restrictivism in some other way. I believe that the problems brought forward by Williamson and Fine show that restrictivism cannot be expressed using quantification at all. However, the moral to be drawn is not that restrictivism is inexpressible, at least not wholly inexpressible, but rather that it must be expressed in a non-quantificational way.

5.2 Ambiguous assertion: any and every

The restrictivist finds herself in the following predicament: in order to formulate her view as a quantified claim about interpretations she must quantify over all interpretations, but doing so commits her to quantifying over absolutely everything. This predicament is fatal, however, only if quantification provides the only means for expressing generality. If there is some non-quantificational means of doing this, the restrictivist might hope to replace (10) with a general claim about interpretations that is quantification-free. Since this claim would not involve quantification at all, it would not carry commitment to quantification over absolutely everything (at least not for the reasons given above).

Famously, Russell (1908, 64–65) argues that general claims can be expressed in two ways: (a) by quantifying over all the values of a variable and (b) by “ambiguous assertion.” In an ambiguous assertion, the speaker asserts something about a particular object, but leaves it indeterminate what object this is. Generality arises from the resulting indeterminacy. Russell claims that ambiguous assertions are common in mathematical proofs. We make ambiguous assertions whenever we prove that some property holds of all objects of some type by selecting an ambiguous particular of the type and showing that the property holds of this particular.

If we say: ‘Let ABC be a triangle, then the sides AB, AC are together greater than the side BC’, we are saying something about *one* triangle, not about *all* triangles; but the one triangle concerned is absolutely ambiguous, and our statement consequently is also absolutely ambiguous. We do not affirm any one definite proposition, but an undetermined one of all the propositions resulting from supposing ABC to be this or that triangle.

I have two remarks about ambiguous assertion in light of this passage. First, I take it that what Russell means when he says that the triangle ABC is ambiguous is not that ‘ABC’ names an ambiguous triangle—a vague object that is not definitely identical to any one particular triangle. What he means is, rather, that it’s ambiguous what triangle ‘ABC’ names. In other words, the ambiguity here is linguistic, not

ontological. This linguistic ambiguity, in turn, is due to the fact that the selection of ABC is arbitrary. Second, I take it that Russell intends ambiguous assertion to be non-quantificational. The relevant quantified claim would be a claim about all objects in the domain—in this case, all triangles—whereas Russell says that in this case, “we are saying something about *one* triangle, not about *all* triangles.”

Immediately prior to the passage above, Russell notes that ‘any’ is often used to make ambiguous assertions. In this respect, ‘any’ functions differently from related words such as ‘all’ and ‘every’. Compare ‘let ABC be any triangle’ (which I take to be equivalent to Russell’s statement ‘let ABC be a triangle’) to the absurd ‘let ABC be every triangle’. In the former, ‘ABC’ refers to an arbitrarily selected triangle. In the latter, ‘ABC’ refers non-arbitrarily to every particular triangle. Dropping the term ‘ABC’, a similar distinction might be drawn between the expressions ‘any triangle’ and ‘every triangle’. The former is used to refer to an arbitrary (and now, unnamed) triangle. The latter is not used to refer, but to quantify over every particular triangle. At a level of greater abstraction, we might say that expressions of the form ‘any F’ are used to refer to an arbitrary F whereas expressions of the form ‘every F’ are used to quantify over all Fs. Finally, statements of the form ‘any F Gs’—any-statements—involve reference to an arbitrary particular and are used to make ambiguous assertions. Statements of the form ‘every F Gs’—every-statements—involve quantification over all Fs and are not used in this way.

If any-statements provide a viable alternative to quantificational generality, the restrictivist might formulate her view by means of an any-statement referring to an arbitrarily selected (but unnamed) interpretation as follows:

(12) Any interpretation has an expansion.

She might also choose a name for an arbitrary interpretation—say, I_a , and formulate her view as:

(12a) Any interpretation, I_a , has an expansion.

As far as I can tell, the difference between (12) and (12a) is unimportant. In what follows, I will confine myself to (12), but what I say should hold for (12a) as well. The first thing to notice about (12) is that it seems to have both of the features we want: (i) it is general, due to the arbitrariness of the interpretation that is selected; (ii) it is also non-quantificational since it is not a claim about all interpretations, but a claim about one arbitrarily selected interpretation.

Unfortunately for the restrictivist, most philosophers doubt whether Russell's distinction between 'any' and 'every' is of any real philosophical importance. On the orthodox view, while 'any' and 'every' can be used to generate generalities in different ways, the generalities generated are the same. This is confirmed by examples. Compare: 'any triangle has three sides' to 'every triangle has three sides'. Each is commonly taken to express the same universal truth about all triangles. The orthodox position is further confirmed by the practice of regimenting both claims as the quantified claim: $(\forall x)(T(x) \rightarrow S(x))$, where ' $T(x)$ ' is ' x is a triangle' and ' $S(x)$ ' is ' x has three sides'.

According to the orthodox position, (12) expresses the same generality as:

(13) Every interpretation has an expansion.

Furthermore, the generality expressed by (12) and (13) is the same as the generality expressed by (10), which is a standard regimentation of both.

5.3 The generality of any-statements

I believe that the orthodox view is mistaken. In this section, I review an argument by Patrick Dieveney that any-statements and every-statements express distinct types of generality. I will then show how his argument can be applied to (12) to support the claim that (12) expresses a distinct sort of generality from (13) and (10). In the next section, I will argue that (12) is properly regimented as a schema, not as a quantified statement.

Dieveney (2013, 123) argues that any-statements differ from every-statements in the following two ways.

First, any-statements employ arbitrary generality whereas every-statements employ universal generality. Second, any-statements can be ‘open’ whereas every-statements are not.

Dieveney’s account of the difference between arbitrary and universal generality is very similar to Russell’s account of the difference between ambiguous assertion and quantificational assertion. Arbitrary generality and universal generality are two different mechanisms for making general claims. The mechanism of arbitrary generality consists in selecting an arbitrary object (of the relevant sort) and showing that some property holds of it. One then infers that the same property holds of any object of the relevant sort. Thus, to show that any F Gs, one selects an arbitrary F and shows that G holds of it. This is the procedure at work in Russell’s triangle example above. The mechanism of universal generality involves considering all objects of the relevant sort and showing that some property holds of each of these objects. To show that every F Gs, one considers all Fs and shows that G holds of each of them.

Unfortunately, such differences between arbitrary generality and universal generality (and between ambiguous assertion and quantificational assertion) are not enough, on their own, to settle the question of whether any-statements express a different type of generality from every-statements. It may be that arbitrary generality and universal generality are merely two mechanisms for generating the very same type of generality.⁵

To show that these statements generate different types of generalities, Dieveney turns to openness. Any-statements not only employ arbitrary generality, they can be open. What this means is that a speaker who makes an any-statement may leave the number of its true instances undetermined. This means that he does not limit

⁵As Dieveney (2013, 126) points out.

himself to one instance, or two instances, or any determinate number of instances. However, he does not go so far as to definitely commit himself to ‘all’ instances. This indeterminacy is critical to the notion of generality that the speaker conveys. By contrast, every-statements are closed. A speaker who makes an every-statement definitely commits himself to all instances in the domain of quantification.

Dieveney (123–124) illustrates this distinction with the following example:

Imagine a scientist on a beach shoveling sand into a small bucket, weighing it, and then dumping it out again. After watching the scientist repeating this process, a bystander approaches him and asks the following questions:

Can any handful of sand be shoveled?

Can every handful of sand be shoveled?

Dieveney claims that the scientist might reasonably answer ‘yes’ to only the first question. That is, he might reasonably assert:

(14a) Any handful of sand can be shoveled

while refusing to assert:

(14e) Every handful of sand can be shoveled.

If (14a) and (14e) express different types of generality—open and closed—then there are two considerations that justify these answers. First, the scientist might worry about the amount of sand: while he believes that any one handful can be shoveled, and any two handfuls can be shoveled, and so on; he might doubt whether he could ever possibly shovel all the handfuls of sand on the beach. For this reason, he might be willing to assert (14a), which (131) “leaves open exactly how many handfuls of sand can be shoveled,” while being unwilling to assert (14e), which makes the stronger claim that all of them can be shoveled. Dieveney calls this the ‘Domain is Too Large’ reason for asserting (14a) but not (14e). Second, the scientist might worry that the intended domain of quantification, the domain containing every handful of sand, is not well-defined. This worry is based on the thought that the boundaries

between individual handfuls of sand as well as the boundaries of the beach are likely to be vague. As a result, while it may be possible to make sense of the notion of an arbitrarily selected handful of sand, it may not be possible to make sense of the domain of all handfuls of sand. This is the ‘Domain is not Well-Defined’ reason for asserting (14a) but not (14e).

In his discussion of the ‘Domain is Too large’ reason, Dieveney writes (127):

This result, that one could always shovel one more, but yet never (practically speaking) run out of new handfuls to shovel is captured by understanding ‘any’ as expressing open arbitrary generality.

It is important to note the parenthetical phrase “practically speaking.” This qualification seems necessary in light of the fact that if the scientist were to continue shoveling sand long enough, he would eventually exhaust all the sand on the beach. In recognition of this, we might say that while, in theory, he would eventually run out of new handfuls to shovel, practically speaking, he would not. But this raises another worry. Practically speaking, it seems that the scientist would reach a limit: at some point, he would be unable (out of boredom or exhaustion) to shovel even one more bucketful of sand, or at least he would reach a point at which it was unclear whether he was able to shovel one more bucketful. Wouldn’t these practical limitations on the extent of the beach and the scientist’s abilities undermine the claim that he could always shovel one more bucketful of sand and consequently the claim that (14a) expresses open generality by conveying the thought that the scientist “could always shovel one more”?

In response, we might modify the case so as to abstract away from such “practical” limitations. In the modified case, the scientist never reaches a point at which he is too tired or too bored to continue shoveling. No matter how long he has been shoveling, he can always shovel one more bucketful of sand. But the beach has become much larger. It now contains infinitely many bucketfuls of sand. While the scientist is tireless, the beach is inexhaustible. So long as the scientist has not been shoveling for

an infinitely long time, he will never exhaust all the sand on the beach. Therefore, no matter how long he shovels, he could always shovel one more bucketful of sand.

However, I believe that this is unnecessary. The openness in the assertion of (14a) may be attributed to the scientist's ignorance of any (definite) limitations on his ability to shovel more, or perhaps to the nonexistence of any precise limitations on this ability. Note that these explanations are compatible with the existence of a fixed limit (whether or not this is a precise limit is irrelevant) to the number of bucketfuls of sand that can be shoveled.⁶

5.3.1 An alternative account

There is another possible explanation for why the scientist might reasonably assert (14a) but not (14e), which does not require that these statements express different types of generality. Quine (1960, 138–140) and Geach (1962, ch. 4) treat both every-statements and any-statements as quantified statements. They explain the felt differences between them as arising from differences in the scope of quantification. The rough idea is that 'any' signifies an occurrence of universal quantification with wide scope; whereas 'every' signifies an occurrence of universal quantification with narrow scope. One of Quine's examples is the pair of English statements:

(15a) I do not know any poem

(15e) I do not know every poem

Using ' Px ' and ' Kx ' as abbreviations for ' x is a poem' and 'I know x ', (15a) and (15e) may be properly regimented as

(16a) $(\forall x)(Px \rightarrow \neg Kx)$

⁶This limit is fixed either by the size of the beach (if the scientist never tires) or the strength of the scientist (if he does tire). If the scientist never tires, then the number of true instances of (14a) is limited by the total number of bucketfuls of sand on the beach and so (14a) is equivalent to (14e). If the scientist does tire, then the number of true instances of (14a) is limited to something less than the total number of bucketfuls of sand on the beach.

$$(16e) \neg(\forall x)(Px \rightarrow Kx)$$

In (16a), \forall has wider scope than it does in (16e), in which it occurs within the scope of the negation operator \neg . Extending this account to the present case, we might regiment (14a) as ‘For every handful of sand, it is possible that it is shoveled’—in which the quantifier has wide scope—and (14e) as ‘It is possible that every handful of sand is shoveled’—in which the quantifier has narrow scope, falling within the scope of the modal operator ‘it is possible that’. Using ‘ $H(x)$ ’ and ‘ $S(x)$ ’ as abbreviations for ‘ x is a handful of sand’ and ‘ x is shoveled’ these can be expressed as:

$$(14a^\diamond) (\forall x)(Hx \rightarrow \diamond Sx)$$

$$(14e^\diamond) \diamond(\forall x)(Hx \rightarrow Sx)$$

Dieveney (129) argues that the scoped account of the difference between any-statements and every-statements has limited explanatory force. First, while it may accommodate the ‘domain is too large’ reason for asserting only (14a)—the scoped translation (14e $^\diamond$) expresses the claim that the scientist might exhaust the sand on the beach, which is stronger than the claim expressed by (14a $^\diamond$)—it cannot accommodate the ‘domain is not well-defined’ reason for asserting only (14a). For the occurrence of \forall in (14a) is meaningful only if the domain of handfuls of sand is well-defined. Therefore, according to this account, it cannot be that one of the reasons for asserting only (14a) is that the domain of handfuls of sand is not well-defined. Second, there are cases in which one may be willing to assert an any-statement but not the corresponding every-statement that are not amenable to scope distinctions. Thus, a finitist may be willing to assert:

(17a) Any number has a successor

but not

(17e) Every number has a successor

on the grounds that (17e) presupposes an infinite domain of numbers, but (17a) does not. Third, there is no satisfactory explanation for why ‘any’ generally takes wider scope than ‘every’. (As evidence of this, Dieveney quotes Quine’s (1960, 139) claim that the correlation of these words with scope distinctions is “a simple irreducible trait of English usage.”)

It is unclear whether these arguments succeed. In reply to the first, one might appeal to the method of supervaluations (Lewis, 1993) and claim that under any eligible precisification of the domain of handfuls of sand, (14e) is false. In this way, one can preserve the thought that (14e) is false without requiring that there be one particular domain of handfuls of sand that stands out as the unique domain of quantification intended by (14e). In reply to the second, one might point out that a strict finitism is untenable and that mathematics requires, if not an actual infinity, at the very least, the notion of a potential infinity, and that this latter notion is highly objectionable in its own right. Furthermore, any philosophical defense of the potential infinite would seem to be unavailable, given my own criticisms of primitive modality in the context of set theory presented in the following chapter. In reply to the third, one might claim that a similar complaint might be brought against Dieveney’s account: there is no satisfactory explanation for why ‘any’ expresses arbitrary open generality and ‘every’ does not. In what follows, I will simply assume that Dieveney’s account of open arbitrary generality is correct. The success of my proposed method for expressing restrictivism should be understood as conditional on this assumption.

5.4 Expressing restrictivism

Recall the attempt to derive Russell’s paradox in the restrictivist set theory FSR (ch. 4). Replacing ‘extension’ and the variable ‘ e ’ with ‘set’ and the variable ‘ x ’, this can be described as follows: We begin with an arbitrary domain D and apply the comprehension principle FC* to define the set R_D by the condition: $(\forall_D x_D)(x_D \in$

$R_D \leftrightarrow x_D \notin x_D$).⁷ Since \forall_D is restricted, we are led to conclude that R_D lies outside the range of \forall_D . We can then shift to a more inclusive context, in which quantification ranges over a more inclusive domain D^+ , which includes R_D . Call this ‘the Russell expansion of D ’. Applying the comprehension principle a second time, we can define the set R_{D^+} , which is then shown to lie outside the range of \forall_{D^+} . We can continue in this way indefinitely to derive the existence of indefinitely many expansions of the original domain (each domain being an expansion of its predecessor). Each of these domains corresponds to an interpretation of the quantifier \forall . Let I be the original interpretation—corresponding to the domain D . Denote the Russell expansion of I by ‘ $r(I)$ ’ and the Russell expansion of $r(I)$ by ‘ $r(r(I))$ ’, and so on.

With this in mind, consider the attempt to express restrictivism by means of the any-statement (12), which (we are assuming) expresses arbitrary open generality. (12) tells us that any arbitrarily selected interpretation I has an expansion. Because it is open, (12) tells us that $r(I)$ has an expansion; and $r(r(I))$ has an expansion, and so on; however, (12) does not tell us that *every* interpretation has an expansion. According to the restrictivist, this claim cannot be expressed on account of the size of the purported domain of quantification (which contains absolutely all interpretations). Moreover, we’ve seen that any attempt to make this quantificational claim is self-defeating. But now, we might worry that insofar as it falls short of making such a strong claim, (12) cannot rule out the existence of a maximal interpretation and is therefore too weak to express restrictivism. I agree that there is a sense in which (12) cannot rule out the existence of a maximal interpretation. On the other hand, if the absolutist were to claim that there was such an interpretation, then we could run through the reasoning of Russell’s paradox to derive its expansion.⁸

⁷I replace ‘extension’ with ‘set’ throughout.

⁸This way of speaking may suggest that (12) is to be understood in terms of what the restrictivist is able to prove. This is objectionable, for restrictivism is not a modal view about what can be proved (as I claimed in chapter 1). Fortunately, (12) needn’t be understood in this way; rather, it should

I contend that (12) should be viewed as making a claim that is intermediate between the strong quantificational claim that (11) purports to make and the quietism advocated by Button, according to which any attempt by the absolutist to quantify over absolutely everything can be shown to fail. While (12) is too weak to say directly that *no* interpretation is maximal; it is sufficiently strong to tell us that, for *any* particular interpretation that the absolutist may select as absolute, is not; that it has an expansion. Indeed, it would be mistaken to expect more than this. As I wrote in chapter 1, “the heart of restrictivism is a picture of the world on which ontology outstrips the resources of language and logic.” According to restrictivism, our ability to quantify, like the scientist’s ability to shovel, is limited; however, in this case, the limitations are not “practical” but deeply rooted in the nature of logic and language. It is a direct consequence of this that we cannot speak about all interpretations; nevertheless, we can say that any arbitrary interpretation has an expansion and consequently that any interpretation that may be specified by the absolutist cannot be absolute.

A similar strategy for expressing restrictivism is pursued by Shaughan Lavine (1994), (2006). In place of any-statements, Lavine appeals to “full sentential schemata”. A sentential schema (henceforth simply schema) is a sentential template, containing free schematic variables, that is associated with a substitution rule, which specifies what expressions from the language can be substituted for the variables. The instances of a schema are the sentences that result from appropriate substitutions. A schema is not itself a truth-bearer; however it carries commitment to its instances and may be called ‘true’ if its instances are true. To accept a schema is to accept all the instances that result from the associated substitution rule.

be understood as saying that any purportedly absolute interpretation has an expansion. This may be validated by running through the procedure described above; but the validation is one thing; the expansion of an interpretation exists independently of its being proved to exist.

It is generally agreed that schemas differ from quantified statements in one respect: whereas a quantified statement carries commitment to its instances by making a single, general claim about them, a schema carries commitment to its instances, without making a single, general claim about them. But this difference alone does not show that schematic generality can replace quantificational generality. There are two reasons why most philosophers have held that it cannot.

The first stems from the fact that whereas quantified variables stand for objects, schematic variables stand for the linguistic expressions from a suitable substitution class. This means that schemas are limited by the expressive resources of the language to which they belong. As a result, we cannot express the generality of ‘every number has a successor’ with a schema such as:

(18) t is a number $\rightarrow t$ has a successor

where ‘ t ’ is a schematic variable that can be replaced by any numeral in English. There are not nearly enough numerals in English—or even in formalized variants of English—to name all the numbers.

The second has to do with the fact that a schema is committed to its instances. To illustrate, suppose we were able to solve the first problem by expanding the substitution class of numerals so that it included numerals for every number. Since (18) is committed to its instances and there is one instance for every numeral, (18) is committed to quantification (in the metalanguage) over all numerals. Given the correspondence between numerals and numbers, this seems just as problematic as commitment to numbers.

Lavine (1994, 230–232) (2006, 117–120) introduces the notion of a “full schema” to overcome these difficulties. First, a full schema is not associated with a particular substitution class. To accept a full schema is to (i) to accept any of its instances that is obtained by replacing the schematic variables with suitable expressions and (ii) to be willing to accept more instances that may be obtained in any suitable expansion

of one's language. This overcomes the first problem since one is not limited by the vocabulary of any particular language. Second, the acceptance conditions for a full schema are not expressed quantificationally, in terms of acceptance of all its instances, but schematically. This overcomes the second problem since (18) is not committed to quantification over linguistic expressions (names) for every possible interpretation, but rather to a corresponding schematic claim in the metalanguage.

If Lavine is right, restrictivism may be formulated as the full schema:

(19) $r(\mathcal{L})$ is an expansion of \mathcal{L}

To accept (19) is (i) to accept any of its instances that is obtained by replacing the variable expressions ' \mathcal{L} ' and ' $r(\mathcal{L})$ ' with names for an interpretation and its Russell expansion and (ii) to be willing to accept more instances that may be obtained from suitable expansions of the language. In accepting (19), one is not committed to quantification over all the possible names that may be substituted for these variable expression; these are specified schematically.

5.5 Concluding Remark

It might be thought that the restrictivist owes us some sort of explanation for *why* logic and language are limited as she claims. The only explanation I can think of is that quantification is limited by size, so that some domains are simply too large to quantify over. However, there are two problems with this account.

First, in 4.5, I objected to the limitation of size doctrine on the grounds that there is no suitable explanation for why size should be relevant to the question of set existence. Why think that a limitation of size doctrine for quantification fares any better in this respect? Second, the limitation of size doctrine faces problems of expressibility similar to those encountered by restrictivism itself. (If some things are too many to quantify over, it would seem that we cannot say so without quantifying over them.)

In response to the first question, I would claim that the demand for explanation in logic is not as pressing as the demand for explanation in set theory. In a set theory based on the limitation of size doctrine, we are dealing directly with the existence of sets; and it seems fair to demand some explanation for why all and only the sets that are said to exist by this theory do indeed exist. In the theory of logic endorsed by the set restrictivist, the case is different. We are not directly making claims about what exists, but about what it is possible to say. According to the restrictivist, the error of FC is precisely that it purports to say something that cannot be said. In fact, if properly understood, FC is true *as it stands*. For it is the case that any sets (where quantification is properly understood as restricted by context to a domain D) form a set. The solution is to recognize that it is impossible to quantify over absolutely any sets at once. However, the question of how to understand these logical truths and the degree to which they can be explained in terms of other facts is a difficult question, which I will not discuss further.

In response to the second question, the restrictivist might give up on the project of expressing the limitation of size conception of quantification and contend that the burden of doing so is unfair. After all, the absolutist has no trouble with understanding it. And isn't the restrictivist appealing to limitation of size in order to satisfy the absolutist's allegation that her grounds for rejecting absolutely unrestricted quantification are ad hoc? Can't limitation of size fulfill this role even if it is inexpressible by the restrictivist's own lights? In direct contrast to Wittgenstein's claim that the limits of language are the limits of reality⁹, the heart of restrictivism is a picture of the world on which ontology outstrips the resources of language and logic. For the restrictivist, any assumption that language and the world go hand in hand is unjustifiably parochial. She seeks to make her view expressible in order to address the

⁹Wittgenstein (2001, §5.6): "The limits of my language mean the limits of my world."

charge of incoherence. But expressibility can only go so far. She hopes to express her view. Perhaps she should not expect to be able to express what motivates and explains it as well.

CHAPTER 6

THE ITERATIVE CONCEPTION OF *SET*

The most popular response to the set-theoretic paradoxes has been to adopt the iterative conception of *set*, according to which sets are not the extensions of concepts (or predicates) but are collections that are “formed out of” or “constituted by” their members. This idea goes back to Cantor’s (1895, 481) definition of ‘set’ as “a collection into a whole of definite distinct objects of our intuition or of our thought.”¹ The notion of a collection is intuitive and may not be fully analyzable; however, it can be characterized as the notion of a singular object related to a unique plurality of objects—its members—in a special, non-mereological way.²

The label ‘iterative’ refers to the process—originally encoded in the axioms of Zermelo (1908, 1930), Fraenkel (1922) and von Neumann (1925)—by which sets are successively formed out of (or built up from) their members.³ Very roughly, this process goes as follows. One begins by forming collections—sets—out of some original things (urelements). One then proceeds to form *new* collections out of these collections and urelements. Next, one forms more new collections out of the most recently formed

¹This translation is from Fraenkel et al. (1973, 15). Set theorists and philosophers who follow Cantor in identifying sets with collections include Shoenfield (1967, ch. 9), Wang (1974, 530), Enderton (1977, 1), Potter (2004, 36), Linnebo (2010) and Oliver and Smiley (2013, ch. 14).

²Mereological sums are also singular objects related to pluralities of objects—their parts—in a special way. However, collections differ logically from sums in at least two important respects. First, ‘the members (of a collection)’ refers to a unique plurality whereas ‘the parts (of a mereological sum)’ does not. Second, parthood is a transitive relation; membership in a collection is not.

³Zermelo’s original 1908 set theory did not include the axioms of Replacement or Foundation. The necessity of Replacement to derive large sets was discovered (independently) by Fraenkel (1922) and Skolem (1922). Foundation was subsequently added by von Neumann (1925). Zermelo incorporated both axioms in his revised 1930 set theory.

collections, the earlier collections and the urelements. One goes on in this way, forming new collections out of previously formed collections and the urelements, forever.

Building upon this foundational work in axiomatic set theory, Joseph Shoenfield (1967), George Boolos (1998, chaps. 1 and 6) and Michael Potter (2004) have articulated alternative axiomatizations of set theory, which represent the iterative process of set formation as occurring in a series of primitive “stages” (Shoenfield, Boolos) or as comprising a series of successively generated “levels” (Potter). Shoenfield and Boolos proceed to derive the traditional axioms of set theory from the set formation rule that at each stage all pluralities of sets formed at previous stages are formed into sets. (I discuss this rule and Boolos’s theory in greater detail below.) Potter adopts an explicit, recursive definition of ‘level’ (originally formulated by Dana Scott (1974)), according to which (i) the first level, V_0 , is the empty set and (ii) each successive level, V_α , is the set containing all members and sub-collections of all preceding levels V_β for $\beta < \alpha$.⁴ He then proceeds to derive the traditional axioms from a version of the Separation axiom, according to which any members of a level form a set. All three accounts qualify as iterative: Shoenfield and Boolos’s because each consists of repeated applications of the rule of set formation (once at every stage); Potter’s because it involves a successive generation of levels (as he himself remarks (2004, 41)). These accounts are also iterative in the original sense according to which sets are formed out of (or constituted by) their members: Shoenfield and Boolos’s insofar as each application of the rule of set formation generates new sets whose members are previously formed sets; Potter’s insofar as each level contains all members and sub-collections of members from previous levels (and sets are defined via Separation as sub-collections of levels).

⁴In standard set theories, levels are recursively defined in terms of the operations of powerset and union as follows: (i) $V_0 = \emptyset$; (ii) for any successor ordinal, $\alpha + 1$, $V_{\alpha+1} = \mathcal{P}(V_\alpha)$; (iii) for any limit ordinal, λ , $V_\lambda = \cup_{\gamma < \lambda} V_\gamma$. The novelty of Scott’s definition is that it does not presuppose the operations of Powerset and Union.

The iterative process of set formation suggests the existence of a real relation of priority between members and the sets they form, which in turn provides a principled explanation for why the contradiction-inducing sets do not exist. This explanatory power gives the iterative conception of *set* an advantage over its competitor: the logical conception of *set*. As we saw in chapter 4, the logical conception is inherently vulnerable to the set-theoretic paradoxes and requires some addendum to save it from contradiction; however, each of the solutions we considered—the limitation of size doctrine and set restrictivism—appears either to be unmotivated by the logical conception and consequently ad hoc, or it raises additional problems, such as the problem of inexpressibility. The iterative conception, on the other hand, appears to be naturally immune to the set-theoretic paradoxes and requires no special addendum to prevent them. Gödel (1947, 518–519) praises the iterative conception for these reasons when he writes:

This concept of set, however, according to which a set is something obtainable from the integers (or some other well-defined objects) by iterated application of the operation “set of”, not something obtained by dividing the totality of all existing things into two categories, has never led to any antinomy whatsoever; that is, the perfectly “naive” and uncritical working with this concept of set has so far proved completely self-consistent.

The iterative conception of *set* can be divided into two theses: (a) the identification of sets with collections; (b) the iterative process of set formation. There are both idealist and a realist variants of (b). According to the idealist version, sets are constructed, built up from, or formed out of their members. A proponent of this view is Shoenfield (1967), who describes a set as a collection that is “formed at some stage” in “the construction of sets.” According to the realist version, sets are constituted by (or ontologically depend upon) their members. A proponent of this view is

Øystein Linnebo (2010, 149), who describes “the concept of a set” as “the concept of a collection ‘constituted by’ its elements.”⁵

One might wonder about the connection between (a) and (b) and, in particular, whether (b) is independent of (a). Two observations support the claim that this is so. The first is that, in a more or less formal sense, we can grasp the nature of the iterative process without forming any beliefs about the intrinsic nature of sets. The second is that philosophers sometimes describe the iterative conception of *set* in terms of (b) alone. The excerpt from Gödel above is one example of this. Donald Martin (2001, 6) provides a second:

When I talk of the concept of set, I mean the *iterative concept*, according to which sets are viewed as being formed in a transfinite “process,” starting with a perhaps empty domain of non-sets (*urelements*).

On the other hand, it might be argued that merely grasping the formal process of set formation is not enough. In order to have a satisfactory conception of *set*, we must understand not only *what* sets are formed according to this process, but *why* it is precisely these sets (no more and no less) that are so formed. Arguably, an answer to the second question requires identifying sets with collections.

In this chapter, I show how the restrictions the iterative conception of *set* imposes on the existence axioms of iterative set theories prevent the formation of the contradiction-inducing sets, with special attention being paid to the two existence axioms—Class Comprehension and Separation—that are closest to the inconsistent axiom of full comprehension (FC) from SN. I then raise two objections to the iterative conception: the first concerns its ability to motivate the empty set axiom of iterative

⁵He goes on to describe three aspects of this:

One aspect of this idea is that part of the nature of a set is what elements it has . . . Another aspect is the converse, namely, that the nature of a set is exhausted by what elements it has. . . A third aspect is that the elements of a set are “prior to” the set itself.

set theory; the second concerns the nature of the priority relation. My plan is the following. In 6.1, I present the iterative process of set formation as a stage theory (ala Boolos) and I show how this process can be used to motivate a No Set solution to the paradoxes. In 6.2, I consider two standard iterative set theories that express the iterative conception: Zermelo-Fraenkel set theory and Von-Neuman-Bernays-Gödel set theory; and I show how restrictions incorporated into the existence axioms of these theories prevent derivations of the paradoxes. In 6.3, I argue that (a liberalized version of) the iterative conception motivates all the existence axioms that are common to these two theories with the exception of the empty set axiom (and two other axioms that depend on this). In 6.4, I discuss problems concerning the nature of the priority relation. I argue that understanding this in a straightforwardly constructivist sense threatens mathematical realism; but the leading realist interpretations—modal and ontological interpretations of priority—are problematic as well.

6.1 The process of set formation

Sets are formed in a series of *stages* according to the rule that at each stage, every plurality, xx , consisting of sets formed at earlier stages and/or urelements is formed into a set $X = \{x|x \prec xx\}$.⁶ If we ignore urelements, this rule can be expressed as the following principle of set formation at a stage (SFS):

SFS: $(\forall s)(\forall xx)[\{x|x \prec xx\}$ is formed at $s \leftrightarrow (\forall x)(x \prec xx \rightarrow x$ is formed at some stage earlier than $s)]$.

The ordering of stages is partially described by the following three axioms:

Ax1. There is a *first stage*.

Ax2. Every stage is immediately earlier than some stage (its successor stage).

⁶This means that each set gets formed repeatedly: at each stage after the stage at which it is first formed.

Ax3. The stages are well-ordered under the relation *earlier than*.⁷

Ax1 tells us that there is a first stage (stage 0). Ax2 then tells us that stage 0 has a successor (stage 1) and that stage 1 has a successor (stage 2) and that stage 2 has a successor (stage 3) and so on. Finally, Ax3 tells us that each of these stages is distinct from the others. In this way, Ax1–Ax3 entail the existence of the infinite sequence: stage 0, stage 1, stage 2, stage 3, ... which is bounded by the first infinite ordinal ω .

SFS tells us what sets are formed at each stage. At stage 0, no sets have been formed at earlier stages, so the only plurality to be formed into a set is the empty plurality, which forms the set \emptyset . At stage 1, \emptyset has been formed at an earlier stage, so two sets are formed: \emptyset and $\{\emptyset\}$. At stage 2, \emptyset and $\{\emptyset\}$ have been formed at earlier stages, so four sets are formed: \emptyset , $\{\emptyset\}$, $\{\{\emptyset\}\}$ and $\{\emptyset, \{\emptyset\}\}$. This process continues for as long as the stages go.⁸

The process of set formation captured by SFS is *exhaustive*: absolutely every set is formed according to this process at some stage. This establishes an iterative existence principle for sets (ISE), according to which $\{x|x \prec xx\}$ exists iff there is some stage s later than any stage at which any $x \prec xx$ is first formed. ISE can be formulated as:

⁷Well-orderings are typically defined on sets; but the standard definition can easily be extended to pluralities as follows:

Any xx are well-ordered under the relation R iff:

- R is transitive on the xx
- The xx satisfy trichotomy (i.e., for any $x_1, x_2 \prec xx$, exactly one of the cases (a) x_1 is earlier than x_2 , (b) x_2 is earlier than x_1 , (c) $x_1 = x_2$ holds)
- Any $yy \prec xx$ have a least element under R .

⁸If there are urelements, things get a bit more complex. Suppose we begin with two urelements, a and b . Then at stage 0, four sets are formed: \emptyset , $\{a\}$, $\{b\}$, and $\{a, b\}$. At stage 1, six elements are available to form new sets, the two urelements a and b , and the four sets from stage 0. Therefore, at stage 1, 2^{2+4} sets are formed. In general, if k is the number of urelements and m sets are formed at stage n , then 2^{k+m} sets are formed at stage $n + 1$. When urelements are included, it is common to delay the formation of sets by one stage so that at stage 0 only urelements exist and no sets are formed. See for example, Studd (2013, 698) and Linnebo (2010, 144).

ISE: $(\forall xx)[(\exists y)(y = \{x|x \prec xx\}) \leftrightarrow (\exists s)(\forall x)(x \prec xx \rightarrow s \text{ is later than any stage at which } x \text{ is first formed})]$.

The right-to-left direction of ISE is a plural version of Boolos’s (1998, 21, 91–93) axiom schema of specification (*Spec*), which uses the schematic predicate ϕ in place of the quantified plural variable ‘ xx ’.

Spec: $(\exists s)(\forall x)(\phi(x) \rightarrow s \text{ is later than any stage at which } x \text{ is first formed}) \rightarrow (\exists y)(y = \{x|\phi(x)\})$.⁹

A direct consequence of ISE is that questions of what sets there are, e.g., whether there is an infinite set, whether every set has a powerset, whether there is a set of cardinality \aleph_ω , and so on, are determined by the length of the series of stages. It is generally agreed that there are infinitely many stages; however, their exact number and even whether this is something that is determined by the iterative conception is open to question.

In order to secure the existence of infinite stages (stages that are infinitely far removed from the first stage) an axiom of infinity is required. Boolos (1998, 21) proposes the following:

Ax4. There is a limit stage, ω , later than all successors of the first stage but not immediately later than any stage.

Repeated applications of Ax2–Ax3 then deliver the infinite sequence of infinite stages: stage ω , stage $\omega + 1$, stage $\omega + 2$, stage $\omega + 3$, . . . Nevertheless, without additional axioms, the stages will still be bounded by the second order of infinity $\omega + \omega$ ($\omega \cdot 2$).

⁹Boolos (1998, 91) introduces the dyadic predicate $B(x, s)$ —read: “ x is formed at a stage prior to s ”—to express the condition that s is later than the stage at which x is first formed. This allows him to formulate *Spec* as: $(\exists s)(\forall x)(\phi(x) \rightarrow B(x, s)) \rightarrow (\exists y)(y = \{x|\phi(x)\})$ Although his official formulation involves only singular quantification, he writes (p.92) that the thought behind *Spec* “can be put better” using plural quantification as: “for any stage s and any sets (notice the plural) that have all been formed before s , there is a set to which exactly those belong.” This in turn can be expressed by a plural variant of *Spec* in which occurrences of ϕ are replaced with plural variables: $(\forall xx)[(\exists s)(\forall x)(x \prec xx \rightarrow B(x, s)) \rightarrow (\exists y)(y = \{x|x \prec xx\})]$.

Boolos, who views every axiom of infinity stronger than Ax4 as exceeding the proper content of the iterative conception, is satisfied with this boundary; however, most set theorists and philosophers of set theory would argue that the series of stages goes much further. This requires a stronger axiom of infinity. We might try:

Ax4*. For any successor stage, s , there is a limit stage, $l(s)$, later than all successors of s but not immediately later than any stage.

Ax4* entails the existence of an infinite sequence of limit stages: stage ω , stage $\omega \cdot 2$, stage $\omega \cdot 3$, ... Each has the form $\omega \cdot n$ and so repeated applications of Ax2–Ax3 deliver infinitely many infinite sequences of the form: stage $\omega \cdot n$, stage $(\omega \cdot n) + 1$, stage $(\omega \cdot n) + 2$, stage $(\omega \cdot n) + 3$, ... Note that for every one of these sequences in which $n > 1$, every stage in the sequence either reaches or exceeds the previous boundary $\omega \cdot 2$. Not only that, each of these sequences has its own boundary, and each of the members of the next sequence (in which the value of n is one greater) either reaches or exceeds this boundary. Nevertheless, the addition of Ax4* is still not strong enough to break through all boundaries. For instance, Ax1–Ax4* do not entail the existence of the limit stage ω^2 and so, if we restrict ourselves to these axioms, we must conclude that the stages are bounded by the order of infinity ω^2 . We might decide to adopt progressively stronger axioms, according to which the stages are bounded by progressively higher orders of infinity, e.g., ω^ω , ϵ_0 and ω_1).¹⁰ How strong should our axiom of infinity be?

The iterative conception is frequently described as one according to which the process of set formation (or the hierarchy of levels, or the height of the set-theoretic universe) is *absolutely infinite* and *surpasses all boundaries*. This opinion is bolstered by an argument that only a maximally strong axiom of infinity is defensible.

¹⁰ ϵ_0 is the limit of the sequence $\omega, \omega^\omega, \omega^{\omega^\omega}, \dots$; ω_1 is the least uncountable ordinal.

P1. Anything less than a maximally strong axiom of infinity imposes a fixed boundary on the universe of sets.

P2. Any fixed boundary is arbitrary.

C. Therefore, only a maximally strong axiom of infinity is defensible.

Fraenkel et al. (1973, 118) appeal to this argument when they object to an “axiom of restriction” which would impose a fixed boundary on the universe of sets:

The axiom of restriction points to the existence of some fixed natural universe of sets, but if the collection of all sets in this universe is again a Platonistic entity, then why should it not be admitted as a new set by allowing a wider universe than that allowed by the axiom of restriction?

The crux of the argument is P2, which can be motivated as follows. Suppose a restricted axiom of infinity, R, is adopted. Since R is restricted, it imposes a boundary on the stages. Whatever the boundary is, ISE tells us that the sets formed in all the stages leading up to it do not form a set. (Such a set could only be formed at the boundary itself, which by assumption is never reached.) Now it seems that we could always have selected a more liberal axiom of infinity, in which case the aforementioned sets *would* have formed a set. And why shouldn't we? Unless there is an answer to this, the choice of R appears arbitrary. (It will not do to reply that there can be no set of *all* sets. For in arguing that it is arbitrary to deny that all the sets formed prior to some boundary form a set, we are not arguing that it is arbitrary to deny that all the sets form a set, but rather that it is arbitrary to assert that all the sets formed prior to the boundary are all the sets.)

The same idea is sometimes used as a *reductio* against an actualist conception of the universe of sets. Thus Linnebo (2013, 206) writes:

According to the actualist conception, the set-theoretic quantifiers range over a definite totality of sets. Why should the objects that make up this totality not themselves form a set? ... To disallow such a set would be to truncate the iterative hierarchy at an arbitrary level.

As I understand him, Linnebo is pointing out a very general difficulty that is inherent in the notion of an actual sequence that “surpasses all boundaries.”¹¹ Just as Russell’s yacht determines a length—say 20 meters—that it does not exceed; so it might be argued, the sequence of stages itself determines a length—whatever this might be—that it does not exceed. It might be replied that Burali-Forti’s paradox has taught us that there is no ordinal number of all the ordinal numbers and so also that there is no length determined by the sequence of all the stages (for what else could this be if not an ordinal number?). However, this reply is a bit perplexing; for if we cannot assign a definite ordinal to the sequence of stages, then how can we be justified in referring to its length? And if we cannot refer to the length of this sequence, then how can we be justified in saying that it surpasses all boundaries? Perhaps the best we can do is to say that the sequence of stages surpasses every independently specifiable boundary (even if, strictly speaking, it does not surpass *every* boundary). This implies that the sequence of stages surpasses every independently given ordinal and therefore that the sequence of stages is at least “as long as” the ordinals. It would seem that this is informative only insofar as the ordinals can be grasped independently of set theory.

Is it possible to express the absolute infinity of the stages without presupposing the ordinals? Shoenfield (1967, 239) proposes a strong “cofinality principle” which is equivalent to the following “boundary schema” for stages which Boolos (1998, 26–27) considers (but rejects):

If each set is correlated with at least one stage (no matter how), then for any set z there is a stage s such that for each member w of z , s is later than some stage with which w is correlated.¹²

¹¹I discuss the potentialist conception that Linnebo favors later in 6.4.

¹²Boolos (1998, 97) also discusses a second principle of infinity, according to which any sets injectable into a plurality of sets formed before some stage s form a set. He rejects both principles on the grounds that they express thoughts about the length of stages that go beyond whatever is plausibly taken to be implied by the iterative conception.

This schema might be described as stating that there is always a stage later than any stage among a given set-sized plurality of stages; however, this is a heuristic and not an analysis, for under the iterative conception, it is the length of the stages that fixes the limits on set-size (and not the other way round). If quantification over functions is permitted,¹³ Boolos’s schema can be expressed in axiom form as:

(B) For any set x and any function f from sets into stages: $(\exists t)(\forall z \in x)(t \text{ is later than any stage } s = f(z))$.

A simple implication of (B) is that there is a stage later than any stage among any given countable sequence of stages s_1, s_2, \dots . *Proof:* Let a_1, a_2, \dots be sets first formed at the stages 1, 2, \dots . By Ax4, there is a stage ω at which the set, A , containing every a_i is formed. Let f_a be the function that maps each a_i to the stage s_i . (B) then tells us that there is a stage t later than every s_i .

Another idea, recommended by Tait (2005, ch.6), Paseau (2007, 33), Welch and Horsten (2016) and (arguably) Burgess (2008, 117–124) is to use a principle of *reflection*, according to which any statement of set theory, in particular, any statement involving quantification over all sets, is “reflected down” onto initial segments of the set-theoretic hierarchy. This is done by means of truth-value preserving reinterpretations under which quantification is restricted to those sets formed prior to some definite stage t . Intuitively, reflection conveys the thought that the stages are absolutely infinite: they go on for so long that initial segments of the domain of all sets become indiscernible from the domain of all sets.¹⁴ To formalize reflection, let ‘ ϕ^t ’

¹³Presumably, this requires a plural analysis of functions, according to which their domains and ranges are identified with pluralities of arguments and values (instead of sets of these).

¹⁴The discovery and development of Reflection principles is largely due to the work of Richard Montague (1957, 1961) and Paul Bernays (1961). For a comparison of various Reflection principles and an overview of some more recent developments, see Incurvati (nd, sec. 2). Reflection is commonly formulated as a schema of ZF set theory, in which particular sets (levels) play the role of stages (see p. 157, f.n. 4). Unlike Tait and Paseau, Burgess presents Reflection as expressing a version of the limitation of size doctrine according to which some xx are too many to form a set if any true

stand for the result of restricting the quantifiers in ϕ to sets formed prior to t and let $\forall x^t/\exists x^t$ symbolize quantification over all sets formed prior to t . Reflection can then be formulated as the following schema:

(Ref) For any stage s , there is a later stage t such that: $(\forall^t x_1, \dots, x_n)(\phi(x_1, \dots, x_n) \rightarrow \phi^t(x_1, \dots, x_n))$.

To illustrate, replace ϕ with the formula $\neg(\exists y)(\forall x)(x \in y)$ (read: there is no set that contains all the sets). (Ref) tells us that for any stage s there is a later stage t such that: ϕ implies the restricted formula $\phi^t = \neg(\exists y^t)(\forall x^t)(x \in y)$ (read: there is no set formed prior to t that contains all sets formed prior to t).

6.1.1 Priority and Explanation

The iterative conception of *set* provides a satisfactory response to the paradoxes only if it provides an explanation for why the contradiction-inducing sets (U , V , V_P , O and K) do not exist (2.5, pp. 52–56). In this section, I will show how SFS blocks the formation of each of these sets (from which it follows, by ISE, that none of them exist). I will then argue that this provides an explanation for why the contradiction-inducing sets do not exist only if either SFS is literally true and sets really are formed from their members in stages or SFS is not literally true, but is a metaphor that aptly models the real facts about set existence.

The non-existence of the contradiction-inducing sets follows from two premises: (i) each of these sets (if it exists) is self-membered and (ii) there can be no self-membered sets according to SFS and ISE. To establish (ii), it suffices to note that according to SFS, if x is a self-membered set that is formed at a stage s , then x is also formed at a stage prior to s . Therefore, no stage can be the stage at which x is first-formed. But since the stages are well-ordered, if there were any stages at which x was formed,

statement ϕ that involves quantification over all of them has a reflection that involves quantification over only those xx belonging to some level V .

there would be a first stage at which x was formed. Therefore, there are no stages at which x is formed. So, by ISE, x does not exist. Below, I will give some arguments in support of (i).

- *The universal set:* Since U is defined as the set of all things, and U , if it exists, is a thing, U is self-membered (if it exists). Similar remarks apply to the set of all sets, V , and the set of all pure sets, V_P . (V , if it exists, is a set and hence a self-member. V_P , if it exists, is a pure set, and hence a self-member.)
- *The Russell set:* Since self-membership is impossible, R , if it exists, is identical to the set of all sets, V . Therefore, R , if it exists, is self-membered.
- *The set of all ordinals:* In iterative set theories, ordinals are defined as transitive sets (all their members are subsets) that are well-ordered by \in . We can prove that O , if it exists, meets these conditions, so O is an ordinal.¹⁵ It follows that O is a self-member (since O is the set of *all* ordinals).
- *The set of all cardinals:* The explanation for the nonexistence of K follows suit. In iterative set theories, each cardinal κ is identified with the least ordinal of

¹⁵Proof sketch: \in is transitive on O since its members are all the ordinals and, by definition, all ordinals are transitive. To show that O is well-ordered by \in , we need to show two more things:

1. \in satisfies trichotomy on O
2. Any nonempty subset $X \subseteq O$ has a least element under \in .

To prove 1, we use the definition of ‘ordinal’ to show that for any ordinals α, β : (i) if $\alpha \subset \beta$ then $\alpha \in \beta$ and (ii) either $\alpha \subseteq \beta$ or $\beta \subseteq \alpha$. (See Jech (2003, 19) for proofs.) From (i) and (ii), it follows that *at least* one of the cases: (a) $\alpha = \beta$, (b) $\alpha \in \beta$ or (c) $\beta \in \alpha$ holds. The well-ordering of α and β under \in ensures that *at most* one of these cases holds. (If more than one of (a)–(c) holds, then $\alpha \in \alpha$ and $\beta \in \beta$, which contradicts trichotomy on α and β .)

To prove 2, take an arbitrary $\alpha \in X$. If $\alpha \cap X = \emptyset$, then α is least, since for any distinct $\beta \in X$: $\beta \notin \alpha$ and therefore $\alpha \in \beta$ by (i) and (ii) above. Next suppose $\alpha \cap X \neq \emptyset$. Since $\alpha \cap X \subset \alpha$, it contains a least element γ . We now show that γ is the least member of X . For any $\beta \neq \alpha \in X$, either $\beta \in \alpha$ or $\beta \notin \alpha$. If $\beta \in \alpha$, then $\beta \in \alpha \cap X$ and so $\gamma \in \beta$. If $\beta \notin \alpha$, then $\alpha \in \beta$ by (i) and (ii) above, and therefore $\gamma \in \beta$. Either way, γ is least.

cardinality κ . Given the isomorphism of the cardinals with the ordinals, the nonexistence of O entails the nonexistence of K .

These remarks explain why none of the contradiction-inducing sets exists *according to the iterative conception*. Our goal, however, is to explain why none of these sets exists *simpliciter*. To achieve this, we must show that the iterative conception is true; in particular that SFS and ISE are true. There are two possibilities to consider:

- (a) SFS and ISE are literally true: this would mean that all sets really are formed from their members in stages, just as SFS and ISE say they are
- (b) SFS and ISE are literally false but metaphorically true: sets are not really formed from their members in stages; however, talk of set formation at a stage is a true metaphor that aptly models the real facts about set existence.

Note that the process of set formation imposes a priority relation between any elements and the set they form. I claim that it is this core idea of priority that provides the explanatory backbone in our account of why certain sets exist or do not exist (and which must be preserved by any metaphorical renditions of these principles).

Priority: The set $A = \{x|x \prec xx\}$ exists iff and *because* the xx are prior to A .

In order to provide a satisfactory response to the paradoxes, we must give an account of this priority relation. Doing so is simple if (a) obtains; for then the priority of the xx to A *just is* the fact that there is a stage later than any stage at which any $x \prec xx$ is first formed. It is more difficult if (b) obtains. In that case, priority must be understood in terms of the real facts about set existence that the metaphor of set formation at a stage aptly models. One idea is that talk of set formation at a stage models real facts about *mathematical modality*. Priority emerges from the fact that each set is merely possible relative to its members. Another idea is that talk of set formation at a stage models real facts about *ontological dependence*. Priority emerges

from the fact that each set ontologically depends upon its members. I will evaluate each of these ideas in 6.4. At present, my goal is to show that some account must be given; that we cannot just leave the metaphor of set formation at a stage as a useful heuristic, informing us of what sets there are according to the iterative conception.

The contrary position is defended by Luca Incurvati (2012, 82), who describes her own view as a type of “minimalism,” according to which the iterative conception, “is exhausted by saying that it is the conception of set according to which sets are the objects that occur at one level or another of the cumulative hierarchy.” I take it that the idea is that principles such as SFS and ISE are purely heuristic, and do not stand for any genuinely explanatory relation. Of course, one might still say that the xx are prior to A , but this is not to be taken in an explanatory sense as providing an account of why the set A exists; rather, it is simply a means of conveying the fact that A exists.

A similar perspective is suggested by George Boolos’s (1998, 91) remarks (which Incurvati cites) that it is possible to explain the iterative conception without invoking a priority relation.

In any case, for the purpose of explaining the conception, the metaphor is thoroughly unnecessary, for we can say instead: there are the null set and the set containing just the null set, sets of all those, sets of all *those*, sets of all *Those*, . . . There are also sets of all *THOSE* Let us now refer to these sets as “those”. Then there are sets of those, sets of *those*, . . .

Let’s grant that Boolos’s recursive list provides a feasible means of conveying what sets there are. In other words, let’s grant that Boolos has provided a recursive list of the sets that exist according to standard, iterative set theory.

I say that this is not enough. Boolos’s recursive list includes a great many sets. Why think that all sets on this list exist? It also leaves out a great many contenders to set-hood, for example, the self-membered sets and the non-well-founded sets. Why think that none of these contenders to set-hood is a set? Stripped of the metaphor of set formation at a stage, we have no assurance that the list is even extensionally

correct. Among the excluded contenders to set-hood are the contradiction-inducing sets. That's certainly a good thing. But stripped of any real priority relation, we have no satisfactory explanation for why these are left out. It's not much help to stress that they must be ruled out to avoid inconsistencies. Excluding them for this reason is, in Michael Dummett's words, "to wield the big stick, not to offer an explanation" (1991, 316).

The shortcomings of minimalism might be further illustrated by comparing it to an imaginary "fix" to naive set theory (SN) which consists of restricting the comprehension principle (FC) to all and only the "safe properties." A property is safe if it leads to no contradiction-inducing sets. Such a theory would be consistent and, it seems, at least as powerful as standard set theory. The problem is that we would be left with no satisfactory explanation for why the "unsafe properties" are excluded. Again, it will not do to observe that they must be excluded to avoid inconsistencies. Minimalism cannot guarantee that the sets that exist according to standard set theory are all and only the sets that there are. Nor, supposing standard set theory gets the right results, can it explain why this is the case. I conclude that minimalism is unacceptable.

6.2 Iterative set theories

There are two standard set theories that are said to express the iterative conception: Zermelo-Fraenkel set theory (ZF) and Von-Neuman-Bernays-Gödel set theory (NBG). Up to this point, I have used 'set' as a generic label for objects that include extensions and classes, as well as the sets of standard set theory. I have done this in order to frame questions about competing conceptions of these objects (the logical and iterative conceptions). In this section, it will be convenient to use 'set' in a more specific way to refer to the objects of ZF. Following standard practice, I use 'class' to refer to the objects of NBG. Thus, ZF is (by definition) a theory of sets; whereas NBG

is (by definition) a theory of classes. Every set is a class and so the NBG universe includes all the sets in the ZF universe. It is therefore unsurprising that NBG incorporates all the ZF axioms with little or no change. However, the NBG universe also includes certain large classes—*proper classes*—that are not sets and consequently do not exist in the ZF universe. The existence of classes is handled by a class comprehension schema unique to NBG. I will first present the ZF axioms in full and then discuss their NBG variants and the NBG class comprehension schema. At the end of this section, I will argue that this schema is philosophically problematic.

ZF is a first-order theory whose quantifiers range over all sets. The ZF axioms are:

ZF1. Extensionality $(\forall x)(\forall y)[(\forall z)(z \in x \leftrightarrow z \in y) \rightarrow x = y]$

Read: coextensional sets are identical.

ZF2. Empty Set $(\exists y)(\forall x)(x \notin y)$

Read: there exists the empty set \emptyset .¹⁶

ZF3. Pairing $(\forall x)(\forall z)(\exists y)(y = \{x, z\})$

Read: any two sets x and z define the set $y = \{x, z\}$.

ZF4. Union $(\forall x)(\exists y)(\forall z)(z \in y \leftrightarrow \exists w(w \in x \wedge z \in w))$

Read: for any set x , there exists the set y of all members of members of x ($\bigcup x$).

ZF5. Power Set $(\forall x)(\exists y)(\forall z)(z \in y \leftrightarrow z \subseteq x)$

Read: for any set x , there exists the set y of all subsets of x ($\mathcal{P}x$).

ZF6. Infinity $(\exists y)(\emptyset \in y \wedge (\forall x)(x \in y \rightarrow x \cup \{x\} \in y))$

Read: there exists the infinite set $\omega = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots\}$.

¹⁶Uniqueness (*the* empty set) follows from ZF1. Without ZF1, ZF2 only asserts that there is *an* empty set. Similar remarks apply to ZF3–ZF8.

ZF7. Separation $(\forall x)(\exists y)(\forall z(z \in y \leftrightarrow z \in x \wedge \phi z))$

Read: for any set x , there exists the (sub)set y of all $z \in x$ such that ϕz .

ZF8. Replacement $(\forall x)[(\forall z \in x)(\exists! w)(\phi(z) = w) \rightarrow$

$$(\exists y)(\forall w)(w \in y \leftrightarrow (\exists z \in x)(\phi(z) = w))]$$

Read: if the condition ϕ is functional on x , then there exists the set y of all $\phi(z)$ such that $z \in x$.¹⁷

ZF9. Foundation $(\forall x)(x \neq \emptyset \rightarrow (\exists y)(y \in x \wedge y \cap x = \emptyset))$.

Read: any nonempty set has a member from which it is disjoint.

While most of these axioms have positive existential import, the first and last do not. ZF1 provides the standard extensional identity conditions for sets (and says nothing about whether any sets exist). ZF9 too says nothing about whether any sets exist; rather, it says that certain sets do not exist. These include self-membered sets as well as sets with infinitely descending chains of membership, i.e., any set x_0 for which there is an infinite series of the form: $\dots \in x_3 \in x_2 \in x_1 \in x_0$. Both types of sets violate SFS (the latter because SFS asserts that every x_j in the series, where $j > i$, must be formed at an earlier stage than x_i ; and this is impossible given the ordering of stages). The remaining seven axioms, which I call “existence axioms”, have positive existential import. These are of two kinds. Some, like ZF2 and ZF6, directly posit the existence of a particular set. Others are conditional: they lay out a procedure for defining “new” sets given the existence of other sets. ZF3, ZF4, ZF5, ZF7 and ZF8 are existence axioms of this kind.

NBG is a two-sorted theory, with distinct variables for sets ($'x'$, $'y'$, $'z'$) and classes ($'X'$, $'Y'$, $'Z'$). The NBG axioms consist of the nine ZF axioms, either entirely un-

¹⁷For a definition of ‘ ϕ is functional on x ’, see 2.1, p. 32.

changed or slightly modified, plus a new class comprehension axiom schema.¹⁸ The NBG set existence axioms (NBG2–NBG8) are exact duplicates of their ZF counterparts (ZF2–ZF8). The NBG axioms of Extensionality and Foundation (NBG1 and NBG9) are slightly modified so that they apply to classes.

NBG1. Extensionality $(\forall X)(\forall Y)((\forall z)(z \in X \leftrightarrow z \in Y) \rightarrow X = Y)$

Read: coextensional classes are identical.

NBG9. Foundation $\forall X(X \neq \emptyset \rightarrow \exists x(x \in X \wedge x \cap X = \emptyset))$.

Read: any nonempty class has a member from which it is disjoint.

The NBG class existence axiom is:

NBG10. Class Comprehension $(\exists X)(\forall x)(x \in X \leftrightarrow \phi(x))$, where $\phi(x)$ is any predicative wff.

Read: any predicative wff ϕ defines the class $X = \{x | \phi(x)\}$.

NBG10 has no ZF counterpart; instead it resembles the full comprehension schema FC of the naive set theory SN (see 4.4 pp. 105–106).

6.2.1 How Class Comprehension blocks the paradoxes

What saves NBG10 from the paradoxes that FC engenders is the restriction to sets imposed by the restricted quantifier $\forall x$. This means that substitutions for ϕ in NBG10 determine classes of *sets*, not classes of *classes*. As a result, while the existence conditions for classes of sets correspond to something very much like the logical conception—every predicative wff ϕ defines the class of all *sets* such that ϕ —the existence conditions for classes of classes do not—it's not the case that every predicative wff ϕ defines the class of all *classes* such that ϕ —for proper classes are

¹⁸This formulation of NBG follows Mendelson (1997, chap. 4). In his 1925 axiomatization, von Neumann leaves out Separation and Replacement. Instead, he employs a limitation of size axiom, which says that any class X defines a set $x = X$ iff there is no surjection from X onto the class of all (pure) sets V_P . Separation and Replacement are derived as theorems from this axiom.

never members of other classes. This prevents the existence of any contradiction-inducing classes.

To illustrate, consider the NBG formulas ‘ x is a non-self-membered class’, ‘ x is a class’, ‘ x is a cardinal’ and ‘ x is an ordinal’. NBG10 tells us that each defines a class. These are: the class of all non-self-membered sets, the class of all sets, the class of all cardinals (which are defined as particular ordinals) and the class of all ordinals (which are defined as particular transitive, well-ordered sets). Contradictions are avoided because these are proper classes and therefore not self-members *even if* they satisfy the relevant defining formulas (for proper classes are outside the range of the restricted quantifier $\forall x$). It is on account of this restriction that we are not forced to infer from the fact that the Russell class defined by replacing ϕ with ‘ x is a non-self-membered class’ in NBG10 is not a self-member that it also *is* a self-member (even though it does satisfy the formula ‘ x is a non-self-membered class’). Similarly, we are not forced to infer that the class defined by replacing ϕ with ‘ x is a class’ is a self-member (even though it satisfies the formula ‘ x is a class’). Instead, we infer that because it is a proper class and therefore outside of the range of the restricted quantifier $\forall x$ that not all of its subclasses are members (even though each of them satisfies the formula ‘ x is a class’). Since not all its subclasses are members, we cannot use identity to define an injection (from the powerclass of the class of all sets into the class of all sets) that contradicts Cantor’s theorem (see 2.2, pp. 36–37). Nor can we derive Cantor’s paradox from the existence of the class of all cardinals. In iterative set theories, only sets have cardinal numbers. Since the class of all cardinals is not a set, it has no cardinal number. Similarly, because the class of all ordinals is not a set, it has no ordinal number.

6.2.2 Is NBG set theory restrictivist?

If we think of NBG classes as extensions of concepts (and not as collections of their members), then NBG10 may be thought of as the result of restricting quantification in FC. Earlier (in ch. 4), I described the comprehension principle FC^* of the restrictivist set theory FSR in the same way: as the result of restricting quantification in FC. Since both these theories (NBG and FSR) are the result of restricting quantification in FC, we might be led to believe that both theories should count as restrictivist set theories. However, this would only invite confusion; for there are two senses in which a theory might be called ‘restrictivist’: an ontological sense and a quantificational sense (up to this point, I have used ‘restrictivism’ in only the quantificational sense). I will now attempt to draw this distinction in the present case by comparing the pictures of set-theoretic reality which underlie the principles NBG10 and FC^* . I will argue that NBG is ontologically restrictivist but not quantificationally restrictivist, whereas FSR is quantificationally restrictivist but not ontologically restrictivist.

The picture underlying NBG10 is one on which there are two kinds of class-like objects: sets (extensions of small concepts) and proper classes (extensions of big concepts). (In what follows, I use these pairs of terms interchangeably.) NBG10 is restricted because it applies only to sets, specifying the conditions under which any *sets* form a class (they do iff they share a property expressible by a predicative formula). The picture underlying FSR is quite different. On this picture, there is only one kind of class-like object: extensions (classes). FC^* specifies the conditions under which any extension $\text{Ext}_D(F)$ exists (it does iff the property F exists). The restriction to D does not reflect the ontological fact that there are special extensions (the proper classes) to which FC^* does not apply, but rather the quantificational fact that there is no single context in which FC^* quantifies over all extensions *at once*.

These two pictures provide two different accounts of where FC goes wrong. On the NBG picture, FC fails to recognize the distinction between sets and proper classes.

FC says (falsely) that any classes (which may include proper classes) form a class (provided they share a property F). In so doing, it mistakenly attempts to apply the principle of comprehension to pluralities of which it does not hold. The solution is to recognize that the principle of comprehension only applies to *sets*. This is the sense in which FC is to be restricted. On the FSR picture, the error of FC is that it purports to say something that cannot be said. FC purports to say that absolutely any extensions (classes), which may be all the extensions, form an extension (provided they share a property F). But this cannot be said. In fact, if properly understood, FC is true *as it stands*. For it is the case that any extensions (where ‘any’ is properly understood as restricted by context to a domain D) that fall under a single property form an extension. The solution is to recognize that it is impossible to quantify over absolutely all extensions at once. There is no all-inclusive D .

The fact that NBG10 employs a restricted quantifier where FC employs an unrestricted quantifier does not reflect any important disagreement between these principles regarding *quantification* (whether or not it is possible to quantify over all classes/extensions); rather, it reflects an important disagreement between these principles regarding *ontology*—whether every class (or extension in SN) is governed by NBG10 or FC. The fact that this ontological restriction is expressed by means of a quantificational restriction is irrelevant. This is merely a means for expressing a restriction that is not itself quantificational, but ontological in nature. Concerning the quantificational question, only FSR is restrictivist: both the inconsistent SN based on FC and NBG based on NBG10 agree that it is possible to quantify over absolutely all objects treated by these theories (extensions for SN; classes for NBG). Concerning the ontological question, only NBG is restrictivist: SN and FSR (but not NBG) agree that comprehension determines the existence conditions for every object treated by these theories.

6.2.3 How Separation blocks the paradoxes

Another axiom that bears similarities to the problematic FC is the axiom schema of Separation, which is common to both ZF and NBG (ZF7/NBG7). Here, the critical difference is the restriction to the members of a given set. This restriction means that substitutions for ϕ in Separation determine *subsets* of sets. In effect, while the existence conditions for subsets of sets correspond to something very much like the logical conception—relative to any set z , every wff ϕ defines the *subset* of all members of z such that ϕ —the existence conditions for sets do not—it’s not the case that every wff ϕ defines the set of all sets such that ϕ . This prevents the derivation of contradiction-inducing sets.

To illustrate, consider once more the formulas ‘ x is a non-self-membered set’, ‘ x is a set’, ‘ x is a cardinal’ and ‘ x is an ordinal’. Separation tells us that each defines a set, by “separating” out all the members from a given set that satisfy it. But unless we can prove that there are *maximal* sets whose members include *all* the sets satisfying these formulas, this procedure will never yield the contradiction-inducing sets. We know, at least in an informal way, that there can be no such maximal sets, since the process of set formation is always churning out new sets satisfying each of these formulas.

Separation’s restriction to the members of a given set bears striking similarities to FC*’s restriction to the elements from a given domain of quantification (see 4.6, pp. 117–122). To further appreciate this similarity between iterative set theories and FSR, consider the following attempt to derive Russell’s paradox from Separation. We start with a set, a . By Separation, we can define the Russell set on a , by replacing ϕ with $x \notin x$.

$$(1.1) \quad (\exists y)(\forall x)(x \in y \leftrightarrow x \in a \wedge x \notin x).$$

By (1.1) and ZF1/NBG1, ‘ y ’ has a unique value, R_a , which is the set of all non-self-membered sets in a .¹⁹ This allows us to eliminate the existential quantifier, replacing the bound occurrence of ‘ y ’ with ‘ R_a ’ (the Russell set on a):

$$(1.2) \quad (\forall x)(x \in R_a \leftrightarrow x \in a \wedge x \notin x).$$

Since $\forall x$ ranges over all sets, we can instantiate ‘ x ’ to R_a :

$$(1.3) \quad R_a \in R_a \leftrightarrow R_a \in a \wedge R_a \notin R_a.$$

If $R_a \in a$, then (1.3) would entail the contradictory:

$$(1.4) \quad R_a \in R_a \leftrightarrow R_a \notin R_a.$$

However, we know, at least in an informal way, that a cannot be maximal with respect to the non-self-membered sets because the process of set formation is always churning out new non-self-membered sets. Consequently, we are free to infer that $R_a \notin a$.

Compare this conclusion and the role played by Separation in the attempt to derive Russell’s paradox in ZF/NBG to the conclusion and the role played by FC* in the attempt to derive Russell’s paradox in FSR. In *that* argument, it was the restriction on quantification in FC* that licensed the conclusion that R_D —the Russell extension defined on D —lies outside the range of \forall_D (pp. 119–120). In the present case, it is the restriction to subsets in Separation that licenses the conclusion that $R_a \notin a$. Indeed, throughout both arguments, a and D play the same role: they limit a principle of set existence (FC or Separation) by restricting the application of the items that define sets (concepts in FC; predicates in Separation) to those sets in a or in D . It is the fact that these entities are not all-inclusive, or maximal, that blocks the paradox. Notice, however, that the restrictions imposed by Separation, like the restrictions imposed by NBG10, are ontological, not quantificational in nature. Separation places restrictions

¹⁹I retain use of capital ‘ R ’ for the Russell set for stylistic reasons.

on which sets can be defined by predicates (only subsets); it does not place restrictions on quantification over sets. As the preceding discussion of NBG10 makes clear, even if Separation were expressed using restricted quantifiers (by restricting quantification to the members of a given set), this restriction would only be the means of expressing an ontological restriction.

6.2.4 Expressing the Iterative Conception: ZF vs. NBG

At the beginning of this section, I wrote that both ZF and NBG are said to express the iterative conception. But NBG does more than this. Its class comprehension schema (NBG10) reflects (a restriction of) the logical conception and the proper classes whose existence is entailed by NBG10 cannot be attributed to any iterative process of class formation. Thus, NBG is a sort of mixed theory, in which both conceptions are at work: the existence axioms NBG2–NBG8 reflecting the iterative conception (applying to all sets), while NBG10 reflects a limited version of the logical conception (applying to all classes).

The addition of proper classes has some advantages. For one, it provides a theory that can be finitely axiomatized.²⁰ For another, it provides a more inclusive theory, one whose domain includes the entire ZF universe as a single object (viz., the class of all sets). Nevertheless, it is philosophically objectionable for two reasons. First, the NBG axioms reflect two very different conceptions—NBG2–NBG10 reflect an iterative conception of *set*, while NBG10 reflects a logical conception of *class*. This causes the theory to be both overdetermined and disunified. The overlap between ‘set’ and ‘class’ means that the existence conditions for sets are overdetermined: once by each conception. The non-overlap between ‘set’ and ‘proper class’ makes for a

²⁰It can be proved that NBG10 is equivalent to a small number of its instances, in which the schematic ϕ is replaced by quantification over classes. Mendelson (1997) proves that NBG10 is equivalent to seven of its instances, which he calls “axioms of class existence.” If the occurrence of the schematic ϕ in Separation and Replacement is also replaced with quantification over classes, we get a finite axiomatization of NBG (the nine ZF axioms plus the seven axioms of class existence).

disjunctive, disunified account of ‘class’ as ‘either a set or a proper class’. John Burgess (2008, 112) goes so far as to call the distinction between sets and proper classes “mystifying”.

Second, the version of the logical conception of *class* that NBG10 reflects is a restricted one according to which all predicative wffs determine classes of *sets* but not classes of *classes*. Like the limitation of size doctrine (see 4.5.1), this restriction appears to be arbitrary and ad hoc. If classes are really the extensions of (predicative) predicates, then any proper classes that satisfy such a predicate should be among the members of the class the predicate defines. (If, on the other hand, classes are not the extensions of predicates, then what reason do we have to think that every predicate determines the class of all the sets that satisfy it? In other words, what reason do we have to think that NBG10 is true?) If NBG10 is true and the wff ‘ x is a set’ determines the proper class whose members are all the sets, then shouldn’t the wff ‘ x is a class’ determine the proper class whose members are all the classes? How is it that these formulas determine the same class? To answer that classes are the extensions of predicates *defined in terms of quantification over sets*, but not *over classes* is to impose an arbitrary and ad hoc restriction on the logical conception of *class*.

For these reasons, I conclude that ZF is philosophically on better footing than NBG. But we’ve seen that NBG is really just ZF set theory plus NBG10; in other words that the theory of sets expressed by the ZF axioms is common to ZF and NBG. Henceforth, I will refer to this theory as ‘iterative set theory’. It is with the set existence axioms of iterative set theory that I will be concerned for the remainder of this chapter.

6.3 Motivating the set existence axioms

The claim that iterative set theory expresses the iterative conception is an interpretive one. The theory itself does not mention the process of set formation: it

includes expressions for sets and the membership relation, not for stages or the operation of set formation. Consequently, the explanation offered on pp. 167–169 above for why there can be no contradiction-inducing sets—viz., because (i) each of these sets (if it exists) is self-membered and (ii) there can be no self-membered sets according to SFS and ISE—cannot be applied to the theory directly. The immediate reason why there can be no contradiction-inducing sets in iterative set theory is because none of these sets can be derived from the existence axioms of iterative set theory.²¹ Alternative explanations that make explicit appeal to the process of set formation (and principles such as SFS and ISE) are appropriate only if, and only insofar as, the process of set formation motivates the existence axioms. Boolos (1998, chaps. 1 and 6) argues that this process does motivate (nearly) every existence axiom of iterative set theory. The notable exception is Replacement.²² In singling out replacement, he is not alone. For example, Putnam (2000) writes:

Quite frankly, I see no intuitive basis at all for ... the axiom of replacement. Better put, I do not see that a *notion* of set on which that axiom is clearly true has ever been *explained*.

Boolos’s own position on Replacement follows from his conservatism on the axioms of infinity for stages. In this section, I take issue with Boolos on two points: first, I

²¹Self-membered sets are inconsistent with the axiom of Foundation (ZF9/NBG9), and therefore it might be thought that it is the prohibition of this single axiom (ZF9/NBG9) and not the limited reach of the seven existence axioms (ZF2/NBG2–ZF8/NBG8) that provides the *immediate* reason why there can be no contradiction-inducing sets in iterative set theory. There is certainly a reading of ‘immediate’ under which this is correct; however, note that Foundation’s prohibition is effective only if it is not contravened by the existence axioms. Considered in isolation, Foundation can only secure a disjunction: either there are no self-membered sets or iterative set theory is inconsistent. The limited reach of the existence axioms is needed to rule out the second disjunct.

²²Boolos also argues that the axioms of Extensionality and Choice are not properly motivated by the iterative conception. Extensionality is not properly motivated by the iterative conception because it is analytic (given the meaning of ‘set’), or, if not analytic, then an obvious truth that is largely independent of the iterative conception (1998, 27–28; 93–94). Choice cannot be motivated by the iterative conception because it cannot be derived from the axioms of stage theory (1998, 28–29; 96–97). These are not exceptions to the claim that every *existence* axiom (but Replacement) can be justified. Extensionality is not an existence axiom (6.2 p. 173). Choice is an addition to iterative set theory (ZF + Choice yields ZFC; NBG + Choice yields NBGC).

argue for a “liberalized”²³ version of the iterative conception, which licenses axioms of infinity for stages that are strong enough to justify Replacement; second, I argue that the iterative conception has difficulty making sense of the notion of a set with no members and therefore has trouble justifying the axioms of Empty Set and Infinity, which directly assert that \emptyset exists, and Separation, which conditionally asserts that \emptyset exists (it does if at least one set exists).

Before discussing these cases of disagreement, however, I will run through the cases of agreement (Pairing, Union and Powerset). My strategy will be to attempt to informally derive each axiom from the principles governing the iterative process of set formation: Ax1–Ax3, together with the appropriate axiom(s) of infinity and the principles SFS and ISE.

Pairing: For any sets a and b , Pairing asserts the existence of the set $\{a, b\}$. Because the process of set formation is exhaustive and well-ordered, there are some stages s and t at which a and b are first formed. Without loss of generality, let t be no earlier than s . Instantiate ISE to a and b so that $\{a, b\}$ exists iff there is a stage later than t . By Ax2, t has a successor stage, $t + 1$, which is later than t . So the set $\{a, b\}$ exists.

Union: For any set a , Union asserts the existence of the set of the members of members of a . Because the process of set formation is exhaustive, there is a stage t at which a is formed. Instantiate ISE to the members of members of a so that $\{x | (\exists y)(x \in y \in a)\}$ exists iff there is a stage later than any stage at which any $x \in y \in a$ is first formed. Instantiate SFS to the stage t and the yy such that $y \in a$ (i.e., the yy such that $(\forall y)(y \prec yy \leftrightarrow y \in a)$). Since a is formed at t , it follows that each $y \in a$ is formed at a stage earlier than t . Given an arbitrary $y_0 \prec yy$, instantiate

²³The label comes from Paseau (2007, 33).

SFS to the stage s_0 at which y_0 is first formed and the $x_0 \in y_0$ such that $x_0 \in y_0$. Since y_0 is formed at s_0 , it follows that each $x_0 \in y_0$ is formed at a stage earlier than s_0 and therefore at a stage earlier than t . Generalizing, t is a stage later than any stage at which any $x \in y \in a$ is first formed. So the set $\{x | (\exists y)(x \in y \in a)\} = \bigcup a$ exists.

Powerset: For any set a , Powerset asserts the existence of the set of all subsets of a . Instantiate ISE to the members of a . Since a exists, there is a stage s later than any stage at which any $x \in a$ is first formed. Next, instantiate ISE to the subsets of a so that $\{x | x \subseteq a\}$ exists iff there is a stage later than any stage at which any $x \subseteq a$ is first formed. Let x_0 be an arbitrary subset of a . Since x_0 's members are members of a , they are first formed at stages prior to s . So, by SFS, it follows that x_0 is formed at s . Since x_0 is an arbitrary subset, every $x \subseteq a$ is formed at s . By Ax2, s has a successor stage, $s + 1$, which is later than s and therefore later than any stage at which any $x \subseteq a$ is first formed. So the set $\{x | x \subseteq a\} = \mathcal{P}a$ exists.

Replacement: For any functional condition ϕ defined on a set a , Replacement asserts the existence of the set $y = \{\phi(x) | x \in a\}$. Alternatively, if ' ϕ 's domain' and ' ϕ 's range' are understood as referring directly to the objects in a and their values (not sets of these), Replacement can be understood as saying that ϕ 's range forms a set if ϕ 's domain does.²⁴ The limitation of size doctrine provides a simple and straightforward motivation for Replacement: the range of a functional condition can be no larger than its domain and so the former forms a set if the latter does. However, this motivation is objectionable in the present context, since considerations of size seem to be independent of the iterative conception of *set*.

²⁴The domain and range of a function are typically defined as sets (see ch. 1), in which case ' ϕ 's range forms a set if ϕ 's domain does' is trivially true.

Replacement entails the existence of sets, which, according to the iterative conception can only be formed at various infinite stages beyond ω . So any successful motivation of Replacement based on the iterative conception must show that all these stages exist. One such set is $\{\omega, \mathcal{P}(\omega), \mathcal{P}(\mathcal{P}(\omega)), \dots\}$, which can be derived by applying Replacement to the condition F , defined on the set \mathbb{N} of natural numbers by the sequence:

- $F(0) = \omega$
- $F(1) = \mathcal{P}(\omega)$
- $F(2) = \mathcal{P}(\mathcal{P}(\omega))$
- \vdots

F is clearly functional and so Replacement asserts the existence of the set $\{F(n) | n \in \mathbb{N}\} = \{\omega, \mathcal{P}(\omega), \mathcal{P}(\mathcal{P}(\omega)), \dots\}$. Let ' ω^* ' name this set. Because each $F(n) \in \omega^*$ is first formed at the stage $\omega + n$, SFS tells us that ω^* is not formed before the stage $\omega \cdot 2$. ISE then tells us that ω^* exists only if the stage $\omega \cdot 2$ exists. It follows that Replacement is true only if the stage $\omega \cdot 2$ exists. If the axioms of infinity for stages to which the iterative conception is committed are too weak to entail the existence of this stage, they are too weak to motivate Replacement.

We've seen that Boolos accepts only the weakest axiom of infinity for stages (Ax4). Since Ax4 is too weak to entail the existence of stage $\omega \cdot 2$, it is too weak to support (this instance of) Replacement and so Boolos concludes that Replacement cannot be motivated by the iterative conception. We've also seen (p.163) that each of the stages $\omega, \omega \cdot 2, \omega \cdot 3, \dots$ can be derived from the stronger axiom Ax4*. Under such a strengthening of stage theory, the set $\{\omega, \mathcal{P}(\omega), \mathcal{P}(\mathcal{P}(\omega)), \dots\}$ exists and the instance of Replacement currently under discussion is true. But it is easy to come up with other instances of Replacement that entail the existence of sets that can only be formed at infinite stages beyond these.

Let ω^{**} name the set $\{\omega^*, \mathcal{P}(\omega^*), \mathcal{P}(\mathcal{P}(\omega^*)), \dots\}$. For the same reason that ω^* is not formed until stage $\omega \cdot 2$, ω^{**} is not formed until stage $\omega \cdot 3$. We can continue in this way. Let ω^{***} name the set $\{\omega^{**}, \mathcal{P}(\omega^{**}), \mathcal{P}(\mathcal{P}(\omega^{**})), \dots\}$. ω^{***} is not formed until stage $\omega \cdot 4$. And so on. Now, consider the sets $\omega, \omega^*, \omega^{**}, \dots$ and the condition F^* , defined on \mathbb{N} by:

- $F^*(0) = \omega$
- $F^*(1) = \omega^*$
- $F^*(2) = \omega^{**}$
- \vdots

F^* is clearly functional. So the corresponding instance of Replacement asserts the existence of the set $\{F^*(n) | n \in \mathbb{N}\} = \{\omega, \omega^*, \omega^{**}, \dots\}$. By SFS, this set cannot be formed before the stage ω^2 and so, by ISE, it exists only if the stage ω^2 exists. But $\text{Ax}4^*$ does not entail the existence of this stage. So, unless we further strengthen the stage theory by adding additional axioms of infinity, this instance of Replacement is false.

The choice of which axioms of infinity to adopt is governed by how one thinks about the iterative conception of *set*. It's possible to distinguish two general approaches. According to the first, which we might call *the sociological view*, the iterative conception reflects actual beliefs about the set-theoretic hierarchy held by practicing set theorists, philosophers of mathematics and their students. According to the second, which we might call *the rational idealization view*, the iterative conception is a rational idealization based (to some extent) on these beliefs.²⁵ For someone who thinks of the iterative conception in the first way, there may be no compelling reason to adopt a maximal axiom of infinity. Although set theorists, philosophers and

²⁵Paseau (2007, 36–40) makes a similar distinction between two types of justification (internal and external) for the axioms of set theory. *Very* roughly put, an internalist justification for the axioms is responsive to the actual beliefs of practitioners in a way that an externalist justification is not.

students of set theory generally agree that the iterative hierarchy is infinite (consisting of infinitely many stages, or levels); there is no consensus on its precise height. This is reflected in the fact that proofs requiring strong axioms of infinity explicitly include the axioms as premises. Thus, a proof of Δ from a strong axiom of infinity γ would be written $\gamma \vdash \Delta$. On the other hand, if we think of the iterative conception in an idealized sense, the lack of actual consensus on the height of the iterative hierarchy is irrelevant and it is reasonable to adopt a maximally strong axiom of infinity in order to avoid the problems of arbitrariness discussed earlier (pp. 163–165). This “liberalization” of the iterative conception entails principles such as (B) or (Ref) from which Replacement can easily be derived.

To derive Replacement from (B), assume for conditional proof that ϕ is functional on an arbitrary set A :

$$(2.1) \quad (\forall z \in A)(\exists! w)(\phi(z) = w).$$

Define f_A on A as the function that maps every $z \in A$ to the stage s at which the set $\phi(z)$ is first formed. Instantiate (B) to the set A and the function f_A to get:

$$(2.2) \quad (\exists t)(\forall z \in A)(t \text{ is later than any stage } s = f_A(z)).$$

(2.2) tells us that for every $z \in A$, the set $\phi(z)$ is formed prior to t . It therefore follows by SFS that the set $\{\phi(z) | z \in A\}$ is formed at t . So $\{\phi(z) | z \in A\}$ exists.

Discharging our assumption,

$$(2.3) \quad ((\forall z \in A)(\exists! w)(\phi(z) = w)) \rightarrow (\exists y)(y = \{\phi(z) | z \in A\}).$$

Since A is arbitrary, Replacement follows by generalization.

To derive Replacement from (Ref), we introduce the notation ‘ A^t ’ to denote an arbitrary set A that is formed prior to some stage t . (‘ A^t ’ can stand for an arbitrary set just as well as ‘ A ’ since every set is formed prior to some stage.) We then proceed as before: assume for conditional proof that ϕ is functional on A^t :

$$(3.1) (\forall z \in A^t)(\exists!w)(\phi(z) = w)$$

Next, apply (Ref) to the monadic predicate ‘ $(\forall z)(z \in x \rightarrow (\exists!w)(\phi(z) = w))$ ’ to get:

$$(3.2) (\forall x^t)[(\forall z)(z \in x \rightarrow (\exists!w)(\phi(z) = w)) \rightarrow (\forall z^t)(z \in x \rightarrow (\exists!w^t)(\phi(z) = w))].$$

Instantiate (3.2) to A^t to get:

$$(3.3) (\forall z)(z \in A^t \rightarrow (\exists!w)(\phi(z) = w)) \rightarrow (\forall z^t)(z \in A^t \rightarrow (\exists!w^t)(\phi(z) = w)).$$

By modus ponens:

$$(3.4) (\forall z^t)(z \in A^t \rightarrow (\exists!w^t)(\phi(z) = w)).$$

Since A^t is formed prior to t , it follows from (3.4) that for *every* $z \in A^t$, the set $\phi(z)$ is formed prior to t . (The possibility that there is some $z \in A^t$ which is not formed prior to t is ruled out by SFS.) It then follows by SFS that the set $\{\phi(z)|z \in A^t\}$ is formed at s . So $\{\phi(z)|z \in A^t\}$ exists. Discharging our assumption,

$$(3.5) ((\forall z \in A^t)(\exists!w)(\phi(z) = w)) \rightarrow (\exists y)(y = \{\phi(z)|z \in A^t\}).$$

Since A^t is arbitrary, Replacement follows by generalization.²⁶

Empty Set: Empty Set asserts the existence of the set \emptyset . In 6.1, I described \emptyset as being formed at stage 0 from “the empty plurality.” This is the consequence of two premises:

- (4) At stage 0, *all* pluralities xx of previously formed sets are formed into sets of the form $\{x|x \prec xx\}$. (This is simply an instance of SFS.)
- (5) The *empty* plurality, consisting of all and only the sets formed prior to stage 0, is among the “all pluralities of previously formed sets” quantified over in (4).

²⁶This proof is suggested Paseau (2007, 33; f.n. 6).

Together (4) and (5) entail that the empty plurality is formed into a (memberless) set at stage 0. This set is \emptyset .

A very similar argument is endorsed by Penelope Maddy (1997, 42–43), who describes the process of set formation as one according to which “at every stage, all possible collections of things [sets] at previous stages are formed.” Since (presumably) (a) “forming a collection” is equivalent to “forming a set” and (b) at any stage and the maximal collection of *all* sets at previous stages is among “all possible collections of sets at previous stages,” she infers that “at any stage, the set of all previously-formed sets is formed.” Finally, since at stage 0, ‘all previously-formed sets’ refers to the empty plurality (the plurality of no sets), Maddy concludes that at stage 0, the set of all previously-formed sets is the set \emptyset :

At the first stage, the set of all previously-formed sets is the empty set.

It follows that \emptyset is formed at stage 0.

Notice, however, that this conclusion follows from Maddy’s description of the process of set formation only if the empty plurality is a “possible collection” of things. This calls for some explanation. The expression ‘possible collection’ involves a modal term—‘possible’—along with a (grammatically) singular noun—‘collection’. The latter seems to be used (in a rather peculiar sense) to refer to pluralities of sets as proto-sets: a mere collection is a proto-set which “becomes” a set when it is “formed.” One possible advantage of speaking this way is that it allows Maddy to avoid irreducible (collective) plural quantification: in place of ‘the *F*s are formed into a set’, which cannot be analyzed in terms of singular quantification as ‘every *F* is formed into a set’, she can say ‘the collection of *F*s is formed (into a set)’, which can be analyzed in terms of singular quantification as ‘the collection consisting of every *F* is formed (into a set)’. Still, a question remains: if collections are pluralities of sets, is *every* plurality of sets a collection? One might think that ‘collection’ is more demanding than ‘plurality’ so that in order to qualify as ‘a collection’, a plurality must meet some

additional condition, e.g., being definable by a distributive predicate, which might be formulated as the condition: $(\forall x)(x \prec xx \leftrightarrow Fx)$. Given widely accepted limitations of language, not all pluralities of objects from an infinite domain can be ‘collections’ in this restricted sense. This explains the need for the additional term ‘possible’, which is used to make it clear that no such restriction is intended: intuitively, *every* plurality of (previously-formed) sets is to count as a *possible* collection of sets.²⁷

Under this interpretation, ‘possible collection’ is a singular replacement for plural quantification and so Maddy’s statement that “at every stage, all possible collections of things at previous stages are formed (into sets),” is equivalent to: “at every stage all things at previous stages are formed into sets.” This is just the (right-to-left direction of) SFS:

SFS[←]: $(\forall s)(\forall xx)[(\{x|x \prec xx\} \text{ is formed at } s) \leftarrow ((\forall x)(x \prec xx \rightarrow x \text{ is formed at some stage earlier than } s))]$.

Maddy’s argument for the empty set then depends entirely on the assumption that the empty plurality of all things formed prior to stage 0 is among the values of the plural quantifier in SFS[←], which is to say that her argument depends entirely on (5).

(5) is logically controversial since it requires a reading of SFS under which the plural variables take empty values; and it is generally held that plural variables cannot take empty values (Linnebo, 2014, 1.2). This standard view is typically enforced by including an existential antecedent in the plural comprehension schema:

²⁷This interpretation of ‘possible collection’ is endorsed by Boolos (1998, 92), who uses the expression in his 1971 discussion of set formation. After describing *Spec* as an attempt to capture the thought that “the sets formed at any stage are ‘all possible collections’ of sets formed at the stages earlier than that one,” he asks: “What is the modal term “possible” doing?” He answers that it is used to say what might be better expressed using plural quantification as (the right-to-left direction of) ISE: “for any stage s and any sets (notice the plural) that have all been formed before s , there is a set to which exactly those sets belong.”

Plural Comprehension $(\exists x)(\phi(x)) \rightarrow (\exists xx)(\forall y)(y \prec xx \leftrightarrow \phi(y))$

Read: there are some things that are all and only the ϕ s if there is at least one ϕ .

and possibly by adding an axiom stating that every plurality is non-empty

$$(6) (\forall xx)(\exists y)(y \prec xx).$$

However, this does not settle the question, as it has been shown that plural logic can also be formulated in a way that allows for empty pluralities (Boolos, 1998, 67–68). The fundamental difference between this and the standard view comes down to a difference between how we read the plural existential quantifier $\exists xx$. According to the standard view, $\exists xx$ is read as ‘there are one or more objects xx ’. According to the alternative view, $\exists xx$ is read as ‘there are zero or more objects xx ’. These two readings are interdefinable.

If one wishes to keep (5) but also to avoid tinkering with plural logic, one might formulate the iterative conception in a way that avoids plural quantification altogether. The basic idea here is to use schematic predicate variables in place of quantified plural variables. We can easily derive the existence of \emptyset from the following schematic variant of ISE:

$$\mathbf{ISE}^\phi (\exists y)(y = \{x|\phi(x)\}) \leftrightarrow (\exists s)(\forall x)(\phi(x) \rightarrow s \text{ is later than any stage at which } x \text{ is first formed}).^{28}$$

Substitute ‘ $x \neq x$ ’ for ϕ to get:

$$(7.1) (\exists y)(y = \{x|x \neq x\}) \leftrightarrow (\exists s)(\forall x)(x \neq x \rightarrow s \text{ is later than any stage at which } x \text{ is first formed}).$$

²⁸ \mathbf{ISE}^ϕ is the result of replacing the quantified plural variable ‘ xx ’ in ISE with the schematic predicate variable ϕ . The right-to-left direction is equivalent to *Spec*. Boolos (1998, 22) derives the Empty Set axiom by substituting ‘ $x = x$ ’ for ϕ in an equivalent reformulation of *Spec*.

Every stage trivially satisfies the condition on the right side, so the left side

$$(7.2) (\exists y)(y = \{x|x \neq x\})$$

follows by modus ponens.

However, even if we can so derive the existence of the empty set from the iterative conception, we face deeper, conceptual problems. According to the iterative conception, sets are collections, and so the empty set is an empty collection. Yet, it is doubtful whether the notion of an empty collection is coherent. For it seems evident that

(8) Part of what it is to be a collection is to have some members.

Since, in this context at least, ‘some members’ cannot be empty, it follows that a “collection of no things” is not a special type of collection—an “empty collection”—rather, it is no collection at all. Consequently, under the iterative conception, there cannot be an empty set; for according to this conception, the notion of an empty set is simply the (incoherent) notion of an empty collection.

Nevertheless, the empty set is almost universally accepted today by set theorists and those philosophers of mathematics who endorse the iterative conception of *set*. Yet, I find it hard to believe that their acceptance of the empty set reflects a considered rejection of (8). To deny (8) is to countenance collections of what there is not. Seventy years ago, Quine taught us the error of countenancing things that are not; surely, having taken this lesson to heart, we should also recognize the error of countenancing collections of such things. But if (8) has not been so widely denied, then why has the argument above been so ineffective?

One answer is historical. Many arguments against the empty set from the late nineteenth and early twentieth centuries were based on claims such as, “sets are mere aggregates,” or “sets consist entirely of their members,” which appear to confuse sets

with either mereological sums or pluralities. If either sort of confusion played an essential role in the historical arguments, it would not be unreasonable for philosophers to ignore them today. A famous example is from Frege (1895, 212):

A class, in the sense in which we have so far used the word, consists of objects; it is an aggregate, a collective unity, of them; if so, it must vanish when these objects vanish. If we burn down all the trees of a wood, we thereby burn down the wood. Thus there can be no empty class.

Context here is critical. This passage is taken from Frege's own critical review of Schröder's logical treatise (*Vorlesungen über die Algebra der Logik*). In that work, Schröder had endorsed a pseudo-mereological view of classes as aggregates and had also accepted the empty class. Frege is pointing out that these positions are inconsistent: there can be no aggregate of what is not. This argument is quite straightforward; however, it is not so clear what is needed to provide a satisfactory response. Donald Gillies (2011, 55) provides a typical gloss when he identifies the underlying problem as a failure by Schröder to adequately distinguish the membership relation \in from the subset relation \subseteq (a set can be thought of as the mereological sum of its subsets, but not its members). This evaluation leads Gillies to sum up his discussion of Frege's argument as follows:

Once the distinction between ' \in ' and ' \subseteq ' had been clearly made (and this was no easy task historically), then the difficulties concerned with the empty set disappeared.

If Gillies is right, then Frege's objection poses no threat to the iterative conception, for surely the iterative conception invites no confusion between \in and \subseteq . However, I think that Gillies's evaluation fails to register the full force of Frege's point. To see this, consider Frege's example once more. According to the iterative conception, classes are collections and so the class of trees that survive the fire is the collection of trees that survive the fire. It is certainly true that we must be careful to distinguish this collection from the mereological sum of surviving trees (as well as the plurality of surviving trees); but merely making this distinction isn't enough to ensure that

the collection of surviving trees exists if no trees survive. Indeed, (8) provides us with a good reason for thinking it does not. It is at least plausible, that collections, while being distinct from sums (and pluralities), nevertheless share with these entities a relation of dependence on their members (or parts) and that as a result, Frege's argument succeeds not only in establishing that there can be no empty sum or empty plurality; but also that there can be no empty collection.

Russell (1903, ch. 6) presents a similar argument against the empty set; though, unlike Frege, he seems sympathetic to the view of sets he targets. He begins by observing that finite classes can be defined by simply enumerating their members. Infinite classes, of course, are an exception, but Russell places little importance on this, since he attributes our inability to define infinite classes by enumeration to theoretically unimportant physical limitations, such as not having enough time. He concludes (p. 67) that sets simply are their members:

Thus Brown and Jones are a class, and Brown singly is a class.

If classes are to be identified with their members, it follows that when there are no members there is no class (as Frege observed). Thus, Russell (1903, 68) notes that a consequence of his view is:

there is no such thing as the null-class, though there are null class-concepts.

(A null class-concept is a property with no instances, e.g., non-self-identity.)

Insofar as this argument seems to depend entirely on the identification of sets with their members, it is bound to strike the contemporary reader as a non-starter. But, as with Frege's argument above, it would be overly hasty to simply dismiss the argument at this point. The crux of Russell's argument is the inference from no members to no class; and while this inference certainly holds if classes are identical to their members, it is not obvious that it holds *only if* classes are identical to their members. It may be enough if classes are dependent on their members. Later, Russell (1903, 74) expresses the objection in slightly different words:

A class which has no terms fails to be anything at all: what is merely and solely a collection of terms cannot subsist when all the terms have been removed.

It's unlikely that Russell meant to distinguish the claim that a class is 'merely and solely a collection' of its members from his earlier identity claim; but it's noteworthy that his argument against the empty class seems to hold good if the former claim is interpreted as the weaker dependence claim that classes are distinct from but wholly dependent on their members. In other words, it's possible to disentangle Russell's argument against the empty set from the identity claim that classes are their members. The critical inference—from no members to no class—holds even if a class (collection) is not identical to its members, provided that it is entirely dependent on them.

A second reason why the argument from (8) has been ineffective is that many have viewed the identity of sets with collections as a mere metaphor, not a serious claim about the nature of sets. Metaphors are only rough guides to truth, and are bound to falter in at least some cases and in at least some respects. In this case, the 'sets are collections' metaphor falters when it comes to the limiting case of the empty set. But this is not evidence that the empty set is not really a set; it is merely an unsurprising consequence of the fact that sets are not really collections. At the beginning of this chapter, I distinguished two components of the iterative conception: (a) the identification of sets with collections; (b) the iterative process of set formation. The view that collection-talk is metaphorical suggests an alternate interpretation of the iterative conception which maintains (b) but denies (a). According to this view, sets are not really collections, but the iterative process of set formation either literally describes or at least aptly models the real facts about set existence. An immediate benefit is immunity from the preceding criticisms of the empty set based on (8). The denial of (a), however, leaves a crucial question unanswered. If sets are not collections, then what shall we say about the nature of sets? There are three possibilities:

- (9) Sets have a primitive nature and the notion of what it is to be a set cannot be analyzed in terms of any other type of entity;
- (10) Sets have a non-primitive nature and the notion of what it is to be a set can be analyzed in terms of some type of entity (other than ‘collection’);
- (11) Sets have no intrinsic nature.

(9), if true, would mark an end to certain explanations. In particular, if ‘set’ is a primitive notion then it is hard to see what (non-pragmatic, non-external) reason we could give for believing that (b) is either the literal truth or an apt model for the real facts about set existence. Similarly, it is difficult to see how we could explain why the sets that exist according to the iterative conception are all and only the sets that there are. These things, it would seem, would just be primitive facts.

Sets have been likened to a number of entities other than collections, which may lend some support to (10). To take one fairly well-known example, Dedekind (1897) described sets as closed sacks, containing their elements. Arguably, it was in the interest of making sense of the empty set that he did so.²⁹ Sacks, unlike collections, are not essentially non-empty, and so there is no argument against empty sacks that parallels the argument from (8) against empty collections. However, Dedekind’s proposal invites new problems. One concerns the uniqueness of the empty set (Oliver and Smiley, 2006, 137). Is there any reason why there can be only one empty sack? A second problem is that relations between sacks and their elements are ill-suited to the relations modeled by the process of set formation. A sack’s contents are in no plausible sense “prior” to the sack (see 5.1.1). Such observations, of course, are symptoms of a more fundamental problem: Dedekind’s comparison is clearly intended as a metaphor. Sacks are physical objects, after all! Other proposals: lassos, corrals, etc.,

²⁹Dedekind expressed this view in an 1897 conversation with Felix Bernstein, later published his recollections of the conversation. See Ewald (2005, 836) for an English translation.

are alike in this regard. Each may have its merits as a metaphor, but none provides the slightest plausibility as a genuine proposal about the nature of sets. Of course, it might be possible to abstract away from the physical characteristics of “sacks” so as to formulate a more eligible candidate for a serious identity. The discussion of “containerized” sets in Oliver and Smiley (2006, 138–140) is a good example.

(11) is the view favored by mathematical structuralists, who are much impressed by the observation that *for purposes of mathematics* all that matters is structure. This observation suggests that the singular terms in mathematics are properly treated as general terms, referring indifferently to any ordinary (non-mathematical) objects that exemplify the requisite mathematical structure. When we speak of “mathematical objects” we refer to ordinary objects without attending to their intrinsic natures.³⁰ Thus, “the sets” refers to any sufficiently numerous plurality of objects that exhibits the right structural properties. In particular, whatever object (from this plurality) is selected for the empty set must not bear the membership relation to any other object (from the same plurality). This structuralist perspective is championed by David Lewis (1991, 13):

We needn’t be ontologically serious about the null set. It is useful to have a name that is guaranteed to denote some individual [=urelement], but it needn’t be a special individual with a whiff of nothingness about it. Ordinary individuals will suffice. In fact, *any* individual will do for the null set—even Possum. Like any individual, he has the main qualification for the job—memberlessness.

Since (11) is motivated by an independently popular view in the philosophy of mathematics, it is likely to be viewed more favorably than (9) and (10). One might then wonder whether a new structuralist version of the iterative conception can be formed

³⁰By ‘structuralism’ I mean ‘eliminative structuralism’. According to the sort of non-eliminative structuralism advocated by Michael Resnik (1997) and Stewart Shapiro (1997), mathematical objects constitute a special ontological category of structural objects that serve as the unique referents of mathematical expressions. I do not consider non-eliminative structuralism here, since it is highly controversial whether it implies, or is even compatible with (11). See MacBride (2005), Shapiro (2006) and Keränen (2006) for discussion.

by joining (11) to (b). However, there is an immediate tension between structuralism and (b). According to structuralism, an object is only considered to be a set relative to the requisite set-theoretical structure. Consequently, it would seem to follow that the existence of each set depends on the entire set-theoretic hierarchy (the entire structure) and therefore on every set in the hierarchy. On the other hand, if sets are formed in stages, as they are according to the iterative conception, then it would seem that the only sets on which a particular set depends are its members. Perhaps this tension between structuralism and the iterative conception can be resolved. If not, then we may have to choose between abandoning the iterative conception in order to save the empty set or abandoning the empty set in the interests of saving the iterative conception.

The latter strategy is pursued by Oliver and Smiley (2006, 2013). In lieu of an empty set, they propose a domain of urelements and an axiom asserting the existence of these urelements to replace the empty set axiom. Sets are then formed from this base according to the iterative process. But there are two problems with this. The first is that the existence of urelements appears to be contingent in a way in which mathematics is not.³¹ A second problem arises from the possibility that there are so many urelements that the assumption that they form a set leads to contradictions in the iterative set theory that is built upon them. This is particularly problematic since there does not appear to be anything in the iterative conception of *set* that would prevent the formation of such a set or urelements (Lewis, 1986, 104).³²

Separation: Like the axioms of Pairing, Union and Powerset, Separation can be straightforwardly derived from SFS and ISE; however restrictions must be imposed in

³¹Oliver and Smiley (2006, 151) bite the bullet here and conclude that set theory cannot provide a foundation for mathematics.

³²Christopher Menzel (2014) considers several modifications to the axioms that would alleviate this tension.

order to prevent instances of Separation from entailing the existence of the empty set. I will run through the derivation of Separation first and then will discuss restrictions.

For any set a , Separation asserts the existence of the (sub)set of all $x \in a$ such that ϕx . Because the process of set formation is exhaustive, a is formed at some stage s . Instantiate ISE to the xx such that $x \in a \wedge \phi(x)$. It follows that $\{x|x \in a \wedge \phi(x)\}$ exists iff there is a stage later than any stage at which any x such that $x \in a \wedge \phi(x)$ is first formed. So $\{x|x \in a \wedge \phi(x)\}$ exists if there is a stage later than any stage at which any $x \in a$ is first formed. Instantiate SFS to the stage s and the members of a . Since a is formed at s , s is later than any stage at which any $x \in a$ is first formed. So the set $\{x|x \in a \wedge \phi(x)\}$ exists.

If there is no empty set then Separation must be restricted. Otherwise, given the existence of at least one set, a , and a predicate such as $x \neq x$ that applies to no sets, Separation asserts the existence of the (sub)set of all $x \in a$ such that $a \neq a$, which is the empty set. There are two ways to prevent this. The first imposes the usual restriction of plural variables to pluralities of one or more objects from the domain. Note that in the derivation of Separation from ISE above, the problematic instances require instantiating ISE to empty pluralities (e.g., the xx such that $x \in a \wedge x \neq x$). Banning empty values for the plural variables is therefore one way to block the unwanted instances of Separation. Alternatively, one might restrict permissible substitutions for ϕ to predicates that apply to at least one object in the domain.

Infinity: Infinity depends directly on the existence of \emptyset . Consequently, without the Empty Set axiom, Infinity is false. However, it is possible to give a conditional justification. And this is what I will do here. I will show that Infinity can be derived from the iterative conception *if* \emptyset exists.

Infinity asserts the existence of the set $\omega = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots\}$. Given \emptyset , it can be shown that ω exists by showing (i) how to derive its members by repeated

applications of Pairing and Union and then showing (ii) that there is a stage later than all stages at which any of the sets derived in this way is first formed. (i) By Pairing, every set x has a singleton, $\{x\}$, and any sets $x, \{x\}$ form a pair set $\{x, \{x\}\}$. By Union, the set: $\bigcup\{x, \{x\}\} = x \cup \{x\}$ exists. If we replace x with \emptyset , repeated applications of these axioms generate the countably infinite series:

- \emptyset
- $\{\emptyset\} = \emptyset \cup \{\emptyset\}$
- $\{\emptyset, \{\emptyset\}\} = \{\emptyset\} \cup \{\{\emptyset\}\}$
- $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} = \{\emptyset, \{\emptyset\}\} \cup \{\{\emptyset, \{\emptyset\}\}\}$
- \vdots

(ii) Since Each set in this series is formed at a finite stage. We know that if x is formed at a finite stage s , then $x \cup \{x\}$ is formed at the next finite stage $s + 1$. If \emptyset is formed at stage 0, then each of these sets will be formed at a finite stage. Instantiate ISE to all sets in this series so that ω exists iff there is a stage later than any stage at which any of these sets is first formed. By Ax4, there is such a stage, viz., the infinite stage ω . So the set $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots\} = \omega$ exists.

6.4 Accounts of Priority

In this section, I discuss three accounts of the priority relation between sets and their members (6.1.1): (i) a constructivist account, according to which SFS and ISE are (more or less) literally true; (ii) a modal account, according to which SFS and ISE are metaphorical for the fact that sets are merely potential relative to their members; (iii) a dependence account, according to which SFS and ISE are metaphorical for the fact that sets ontologically depend on their members (and on nothing else).

6.4.1 The constructivist account of priority

The constructivist account of priority may be motivated by Cantor's (1895, 481) definition of set (quoted at the beginning of this chapter) as "a collection into a whole of definite distinct objects of our intuition or of our thought." The identification of sets with *collections* whose members are objects "of our *intuition* or of our *thought*," suggests a mental process according to which new sets are formed by acts of mentally conceiving, or collecting, previously formed sets. This, in turn, suggests a temporal interpretation of the stages of set formation, according to which the priority of members to their set is to be understood as the temporal priority of mental construction: the members must be temporally *available* as objects of intuition or thought before it is possible to mentally collect them. If in addition, it is possible to mentally collect *any* elements that are temporally available, it follows that *any* available elements form a set. Availability may then be understood as equivalent to existence at a stage in the process of set formation. There are a number of objections that have been raised against this view.

The first concerns the empty set. It is based on the thought that if sets are formed by acts of collecting their members, then the act of collecting nothing ought to result in the failure to form any set at all, not in the formation of the empty set. Consequently, if we wish to maintain the coherence of an empty set, we must reject constructivism (Black, 1971, 618–622). In reply, it might be noted that all versions of the iterative conception identify sets with collections and are therefore subject to the original objection against the empty set based on (8) above. As a result, the relative force of this objection as an argument against constructivism in particular (as compared to other more realist interpretations of the iterative conception) is unclear.

The second objection concerns time. Since (presumably) mental acts take place in time, constructivism has the implication that the existence of infinite sets (at least, sets, such as ω , of infinite rank) depends on the existence of infinite sequences of tem-

porally ordered mental acts. This is problematic for two reasons. First, it is doubtful whether such sequences are possible for finite creatures, and so constructivism may be committed to immortal minds. Second, even granting the existence of immortal minds, time itself may not be long enough for the construction of sets larger than the continuum.

Instead of an actual sequence of mental acts, the constructivist might try understanding set formation in terms of a more abstract, idealized possibility of mentally collecting (Wang, 1974). She might then reason as follows. Assume that \emptyset is formed. Then it's possible to mentally collect \emptyset . So $\{\emptyset\}$ is formed. Then it's possible to mentally collect \emptyset and $\{\emptyset\}$ (individually and jointly). So $\{\{\emptyset\}\}$ and $\{\emptyset, \{\emptyset\}\}$ are formed. And so on. Since none of these *possible* acts need to be understood as actually occurring in time, there is no longer any reason to doubt the existence of sets that can only be formed after continuum-many steps.

However, this notion is quite obscure. Given the availability of a countable infinity of sets, e.g., the countable sequence $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots$, is it possible to mentally collect them in the relevant idealized sense, thereby forming an infinite set? What if we are given an uncountable infinity of sets? Is it possible to mentally collect these, thereby forming an uncountably infinite set? Wang (1974) answers that it is possible in both cases; but it seems more likely that his answer is guided by prior commitments to standard set theory than by an independent grasp of what it is possible to mentally collect. It is of little help when he equates the ability to collect the sets falling under a concept with the ability to “look through” or “run through” them “in an *idealized* sense, ... in such a way that there are no surprises as to the objects which fall under the concept” (1974, 531).

Individual axioms of set theory raise additional questions. The axiom of Separation states that given a set, x , any set-theoretic predicate ϕ defines the subset $y = \{z | z \in x \wedge \phi(z)\}$. Moreover, Separation is true, according to the constructivist,

only if it is possible to mentally collect the members of each of these subsets. Is it? What if ϕ involves quantification over *all* sets? Must it then be possible to mentally “look through” or “run through” every set in order to determine which sets belong to y ? In order to answer this, the constructivist must clarify her notion of an idealized (possible) act of mentally collecting available elements. The prospects of doing so are dim.

6.4.2 The modal account of priority

Modal accounts of the priority relation are presented by Parsons (1983, chaps. 10, 11), Linnebo (2010, 2013) and Studd (2013). The goal of these accounts is, in Parsons’s words, to maintain the idea that “any available objects can be formed into a set” but “to replace the language of time and activity by the more bloodless language of potentiality and actuality” (1983, 293). More specifically, talk of an infinite process of set formation in stages is to be treated as a metaphor for the underlying modal facts: a potential hierarchy of sets in which each set “is an immediate possibility given its elements” (1983, 294). Linnebo describes the modal account similarly, as one “based on the idea that the hierarchy of sets is a potential one, not a completed or actual one” in which “the existence of a set is potential relative to its elements” (2010, 155). And Studd writes that the tensed language used to characterize the iterative conception of *set* is to be taken “seriously,” but not “literally” (as describing an actual, temporal process of set construction); rather, the tense is to be “replaced with suitable modal operators governed by a tense-like logic” in which the potential nature of sets is expressed by the thesis that “any sets *can* form a set” (2013, 699–700). (This thesis may appear stronger than Parsons’s qualified claim that any *available* objects can be formed into a set. However, a similar qualification is built into the modal context in which Studd’s statement applies so that the occurrence of

‘any sets’ in this context has the force of ‘any *actual* sets’ outside of it.) In what follows, I shall refer to this as “the core thesis” of the modal account.³³

In order to properly characterize the potential nature of sets, and to properly formulate the core thesis, we need a suitable modal logic. Recall that according to the iterative process of set formation—which will serve as our intuitive guide for the underlying modal facts—(a) new sets are formed at every stage and (b) sets, once formed at a stage s , are formed at every later stage t . (a) and (b) can be represented in a modal logic by an ordering on possible worlds (of sets) under the relation $>$, where $w_j > w_i$ iff w_j ’s domain is a proper expansion of w_i ’s domain. Intuitively, the sets that exist at w_j are (b) all the sets that have been formed at previous stages and which therefore exist at every preceding world w_i , as well as (a) any (new) sets that *can* be formed out of these. This ordering is captured by a directed, S4 logic, in which the accessibility relation is defined as \geq (a relation that is transitive, reflexive, and anti-symmetric).³⁴

The greatest challenge for the modal account is explicating the relevant notions of *possibility* and *necessity*. Linnebo represents these by means of the standard modal operators \Box and \Diamond , which are semantically modeled by quantification over possible worlds in the usual way:

- $\Box\phi$ is true at w iff ϕ is true at every w -accessible world.
- $\Diamond\phi$ is true at w iff ϕ is true at some w -accessible world.

³³Studd (2013, 699) calls it “the maximality thesis.” Linnebo (2010, 157) refers to it as “Collapse \Diamond ”: a modal interpretation of the inconsistent “Collapse” principle, according to which any things *do* form a set. Parsons (1983, 280–297) traces this back to the claim, which he attributes to Cantor (1899), that any “consistent multiplicity” can form a set. That any sets *can* form a set is also a consequence of Kit Fine’s (2006a; 2006b) modal formulation of restrictivism, according to which any mathematical domain can be “expanded” by appropriate procedural postulates.

³⁴In the interests of deriving a modal set theory, Linnebo strengthens this logic to S4.2, in which accessibility is also *convergent*. (Accessibility between worlds is *convergent* if whenever a world w_1 accesses worlds w_2 and w_3 , there is a world, w_4 , that w_2 and w_3 both access.) Studd strengthens the logic to S4.3, in which accessibility is also *connected*. (Accessibility between worlds is *connected* if for any distinct worlds w_1 and w_2 , either w_1 accesses w_2 or w_2 accesses w_1 .)

In the interests of deriving the well-foundedness of sets, Studd complicates things slightly by employing a “bi-modal” logic, in which \Box is replaced by a pair of operators: the (strictly) forwards looking operator $\Box_{>}$ and the (strictly) backwards looking operator $\Box_{<}$. Similarly, \Diamond is replaced by the (strictly) forwards looking operator $\Diamond_{>}$ and the (strictly) backwards looking operator $\Diamond_{<}$. These four operators can be semantically modeled analogously to \Box and \Diamond as follows:

- $\Box_{>}\phi$ is true at w iff ϕ is true at every w -accessible world (other than w).
- $\Box_{<}\phi$ is true at w iff ϕ is true at every world that accesses w (other than w).
- $\Diamond_{>}\phi$ is true at w iff ϕ is true at some w -accessible world (other than w).
- $\Diamond_{<}\phi$ is true at w iff ϕ is true at some world that accesses w (other than w).

Using these operators in a suitable S4 modal logic, it’s possible to formulate the core thesis that any sets can form a set in several ways: (i) as a single formula in a second-order language (Parsons); (ii) schematically in a singular, first-order language (Parsons and Studd); (iii) as a single formula in a plural, first-order language (Linnebo). I will briefly discuss each in turn.

(i): In ch. 4., I presented Parsons’s second-order formulation of the core thesis as

$$\mathbf{FC}^{\Diamond} \Box(\forall F)\Diamond(\exists y)(\forall x)(x \in y \leftrightarrow Fx).$$

Intuitively, \mathbf{FC}^{\Diamond} says that (necessarily) any property F determines the possible existence of the set of all F s. To avoid paradox, ‘all F s’ must be read in such a way that it refers to all the F s there actually are (at w), not to all the F s that there would be (at w) if these F s formed a set. Parsons enforces such a reading by “fully rigidifying” properties relative to worlds, as I showed (pp. 123–124).

(ii): The idea that (necessarily) any formula ϕ determines the possible existence of the set of all ϕ s is naturally formalized as (the necessitation of):

$$(12) \ \diamond(\exists y)(\forall x)(x \in y \leftrightarrow \phi(x)).$$

If ϕ is rigid, or as Studd calls it “modally invariant” ($\text{INV}[\phi]$), so that at any stage in the process of set formation, everything is necessarily ϕ or necessarily not ϕ , this idea may be strengthened to the idea that any formula ϕ determines the possible existence of the set of all the ϕ s *there could possibly be*. This is naturally formalized as:

$$(13) \ \diamond(\exists y)\Box(\forall x)(x \in y \leftrightarrow \phi(x)).$$

Unfortunately, both (12) and (13) are false. Under a substitution of ‘ $x \notin x$ ’ for ϕ , they each imply the possible existence of the Russell set R from which the contradiction $R \in R \wedge R \notin R$ is derivable. Additional contradictions arise under substitutions of the formulas that define the sets: U , V , V_P , O and K .

In an effort to diagnose the problem, Studd notes that these contradiction-inducing formulas share a property, which is intuitively characterized by the observation that at every stage, ϕ defines new sets that do not exist at earlier stages. Studd dubs this property “indefinite extensibility” which he formally defines as follows:

- A formula $\phi(x)$ is extensible ($\text{EXT}_x[\phi(x)]$) =_{df.} $\phi(x)$ is satisfied by sets that are merely possible (intuitively, they will be formed only at a later stage)
- A formula $\phi(x)$ is indefinitely extensible ($\Box\text{EXT}_x[\phi(x)]$) =_{df.} $\phi(x)$ is necessarily extensible (intuitively, $\phi(x)$ is extensible at every stage).³⁵

He proceeds to modify (13) by restricting ϕ to those formulas that are not only invariant but also not indefinitely extensible. He dubs the modification “Max” for the core thesis (which he calls “the maximality thesis”) that any sets can form a set.

$$\mathbf{Max} \ \text{INV}[\phi] \wedge \neg\Box\text{EXT}_x[\phi(x)] \rightarrow \diamond(\exists y)\Box(\forall x)(x \in y \leftrightarrow \phi(x)).³⁶$$

³⁵See 1.4.3 and 3.1 for discussion of indefinite extensibility.

³⁶In the context of a modal set theory, statements of the form ‘ $x = y$ ’ and ‘ $x \in y$ ’ are only true if both x and y actually exist. In order to avoid treating instances of such statements in the language

Intuitively, the restriction to formulas that are not indefinitely extensible means that any sets that might be determined by such formulas, such as (all) the non-self-membered sets, are excluded from consideration: how then can Studd maintain that Max captures the core thesis that *any* things can form a set? His answer (p. 699) is that the proper interpretation of ‘any sets’ in the context of the iterative conception of set is ‘any sets formed at some stage in the iterative process’.

No matter how far we’ve proceeded up the hierarchy, the sets formed so far are *all* the sets there are: (otherwise) unrestricted quantification over sets ranges only over the ones formed so far.

Isn’t this a type of restrictivism? The answer depends on perspective. From a non-modal perspective, according to which absolutely all sets are actual, Studd’s “proper interpretation” imposes a quantificational restriction under which ‘any sets’ is understood as ranging over only those sets that are formed at some stage. However, from the modal perspective, according to which the set-theoretic hierarchy is potential, Studd’s “proper interpretation” does not impose a quantificational restriction. The fact that ‘any sets’ cannot take certain pluralities as its value—such as *absolutely all the sets*, i.e., all the sets that will ever be formed—is simply due to the modal fact that for any indefinitely extensible formula ϕ , there is no point at which the plurality of all ϕ s is actual.

(iii): Linnebo formulates the core thesis in a plural, first-order modal logic as:

$$\mathbf{FPC}^\diamond \quad \Box(\forall zz)\diamond(\exists y)\Box(\forall x)(x \in y \leftrightarrow x \prec zz).$$

\mathbf{FPC}^\diamond is a modalized version of the inconsistent full plural comprehension principle:

$$\mathbf{FPC} \quad (\forall zz)(\exists y)(\forall x)(x \in y \leftrightarrow x \prec zz)^{37}$$

of standard set theory as being contingent (true or false depending on the stage at which they are evaluated), Studd translates them into the language of modal set theory by adding a \diamond operator out front: $\diamond(x = y)$ and $\diamond(x \in y)$. Taking this into account, his official statement of Max, is: $\text{INV}[\phi] \wedge \neg\Box\text{EXT}_x[\phi(x)] \rightarrow \diamond(\exists y)\Box(\forall x)(\diamond(x \in y) \leftrightarrow \phi(x))$.

³⁷FPC is the pluralization of FC (i.e., the result of replacing the quantification over properties in FC with plural quantification over objects). It is not surprising, therefore, that FPC leads to

and a “pluralization” of FC^\diamond :

$$\mathbf{FC}^\diamond \quad \Box(\forall F)\diamond(\exists y)(\forall x)(x \in y \leftrightarrow Fx).$$

Like FC^\diamond , FPC^\diamond involves quantification into the scope of the modal operator \diamond (though in this case the quantification is plural, not second-order). To avoid paradox, the plural variables must be “fully rigidified” so that they refer to the same things both inside and outside the scope of this operator (like the property variables in FC^\diamond). The definition of full rigidity for plurals is analogous to Parsons’s definition of full rigidity for properties.

- The xx are *rigidified* relative to a world $w =_{df}$. anything that is among the xx at w is necessarily among the xx and anything that is not among the xx at w is necessarily not among the xx .
- The xx are *fully rigidified* relative to a world $w =_{df}$. (i) the xx are rigidified relative to w and (ii) the xx at w are all the possible xx .³⁸

FPC^\diamond is consistent if the zz are (fully) rigidified relative to the world w to which the outermost \Box in FPC^\diamond is instantiated. The result is that the occurrence of ‘ zz ’ in ‘ $x \prec zz$ ’ is read as falling outside the scope of \diamond . This blocks the derivation of Russell’s paradox when ‘ $\forall zz$ ’ is instantiated to the non-self-membered sets (at w).

Russell’s paradox when $\forall zz$ is instantiated to the non-self-membered sets. The derivation, which requires the plural comprehension principle (p.191), parallels the derivation of Russell’s paradox from FC presented in chapter 4. See Linnebo (2010, 146–147) for a clear presentation.

³⁸Conditions (i) and (ii) correspond to Studd’s notions of invariance and (the denial of) extensibility. Linnebo (2013, 211) uses the terms “stability” and “inextensibility” and shows that each can be expressed formally in the language of plural modal logic. The modal “stability” principles for pluralities are:

$$\mathbf{Nec}_\prec \quad x \prec xx \rightarrow \Box(x \prec xx)$$

$$\mathbf{Nec}_\not\prec \quad x \not\prec xx \rightarrow \Box(x \not\prec xx)$$

Inextensibility for pluralities is expressed by “relativizing” the Barcan formula, which guarantees that the domains of accessible worlds contain no additional objects, to the condition $x \prec xx$:

$$\mathbf{BF}_\prec \quad \forall x(x \prec xx \rightarrow \Box\phi) \rightarrow \Box\forall x(x \prec xx \rightarrow \phi).$$

Instead of referring to the non-self-membered sets that exist at the world v to which \diamond is instantiated, the occurrence of ‘ zz ’ in ‘ $x \prec zz$ ’ now refers to the non-self-membered sets that exist at w .

As with Studd’s restriction to formulas that are not indefinitely extensible, Linnebo’s rigidification of the plural variables excludes any “contradiction-inducing pluralities” from consideration. Like Studd, Linnebo argues that this is justified by the potential character of the set-theoretic hierarchy. He acknowledges that this has the unintuitive result that, for example, we cannot refer plurally to absolutely all the sets; however, such limitations are not quantificational in nature; rather, they are necessary consequences of the modal fact that certain pluralities are never actual. Thus, Linnebo (2010, 159–160) writes that “you cannot lose something you never had.” He goes on to explain:

A plurality consists of a fixed range of objects, but the set-theoretic hierarchy is inherently potential and thus resists being summed up by a fixed range of objects.

6.4.2.1 Modal set theory

The core thesis, together with supplemental modal principles, provides the basis for a modalized version of iterative set theory (MST) in which all quantified formulas of ordinary iterative set theory are rewritten according to the translation scheme:

- $(\forall x)\phi(x) \mapsto \Box(\forall x)\phi(x)$
- $(\exists x)\phi(x) \mapsto \Diamond(\exists x)\phi(x)$.³⁹

When ϕ is a predicate satisfied by every set in the language of standard set theory, the modal theorist will say that ϕ is necessarily satisfied by every set in the language of modal set theory, i.e., ϕ is satisfied by every set at every stage. Intuitively, this allows

³⁹Putnam (1979, 56–59) may have been the first to suggest this strategy. See also Hellman (1989, 65–79).

the modal theorist to make general claims about “all sets” without presupposing their actual existence.

Of particular interest are the modal translations of the iterative existence axioms:

Empty Set [◇]	$\diamond(\exists y)\Box(\forall x)(x \notin y)$
Pairing [◇]	$\Box(\forall x)\Box(\forall z)\diamond(\exists y)(y = \{x, z\})$
Union [◇]	$\Box(\forall x)\diamond(\exists y)\Box(\forall z)(z \in y \leftrightarrow \diamond\exists w(w \in x \wedge z \in w))$
Power Set [◇]	$\Box(\forall x)\diamond(\exists y)\Box(\forall z)(z \in y \leftrightarrow z \subseteq x)$
Infinity [◇]	$\diamond(\exists y)(\emptyset \in y \wedge \Box(\forall x)(x \in y \rightarrow x \cup \{x\} \in y))$
Separation [◇]	$\Box(\forall x)\diamond(\exists y)\Box(\forall z)(z \in y \leftrightarrow z \in x \wedge \phi z)$
Replacement [◇]	$\Box(\forall x)[\Box(\forall z \in x)\diamond(\exists! w)(\phi(z) = w) \rightarrow$ $\diamond(\exists y)\Box(\forall w)(w \in y \leftrightarrow \diamond(\exists z \in x)(\phi(z) = w))]$

The project of deriving these axioms from the core thesis, supplemented with additional axioms that express the iterative conception of *set* in the modal context parallels the project described in 6.3 of deriving the existence axioms of (non-modal) iterative set theory from axioms that express the iterative conception in a non-modal context (Ax1–Ax4, SFS, ISE). Following upon earlier work by Parsons, Linnebo and Studd execute this project with great technical skill.⁴⁰ However, the significance of

⁴⁰Linnebo adds four axioms: the first two are (the necessitation of) the ordinary axioms of Extensionality and Foundation:

$$\mathbf{(Ext)} \quad \Box(\forall x)(\forall y)[(\forall z)(z \in x \leftrightarrow z \in y) \rightarrow x = y]$$

$$\mathbf{(F)} \quad \Box(\forall x)(x \neq \emptyset \rightarrow (\exists y)(y \in x \wedge y \cap x = \emptyset)).$$

The second two exploit the full rigidity of plurals to express the full rigidity—or, as he calls it, “extensional definiteness”—of membership and subsethood:

$$\mathbf{ED-}\in \quad \Box(\forall x)(\exists yy)\Box(\forall z)(z \prec yy \leftrightarrow z \in x)$$

$$\mathbf{ED-}\subseteq \quad \Box(\forall x)(\exists yy)\Box(\forall z)(z \prec yy \leftrightarrow z \subseteq x).$$

Intuitively (ED- \in) says that any set x has the same members at any world in which it exists and (ED- \subseteq) says that any set x has the same subsets at any world in which it exists. From these, together with FPC[◇], Linnebo is able to derive all the axioms above except Infinity[◇] and Replacement[◇]. To

their achievement rests in large part upon the prior question of the intelligibility of the modal notions in play. To this question, I now turn.

6.4.2.2 Understanding the modal operators

Insofar as we think of worlds as stages of set formation and of the modal operators as quantifiers over worlds, we may be inclined to define the latter in terms of quantification over the stages of set formation as follows:

D1. $\Box\phi$ is true at stage $s =_{df.}$ ϕ is true at s and at every later stage t .

D2. $\Diamond\phi$ is true at stage $s =_{df.}$ ϕ is true at some later stage t .

Clearly, these definitions are inadequate for the modal theorist. His hope is to use modality to provide a non-constructivist interpretation of the priority relation. This is impossible if the modal operators he employs are defined in terms of the stages in a constructivist process. Of course, D1 and D2 do not offer an analysis of stages (constructivist or not) and so they are compatible with a non-constructivist analysis of stage theory. But, at best, what this shows is that D1 and D2 do not go far enough.

The modal theorist might replace quantification over stages in D1 and D2 with a tensed language and use this to define the modal operators as temporal operators as follows:

D3. $\Box\phi$ is true (now) $=_{df.}$ ϕ is true (now) and will remain true, no matter how many sets are formed.

D4. $\Diamond\phi$ is true (now) $=_{df.}$ ϕ will become true when enough sets have been formed.

derive Infinity $^\diamond$, he adds a limitation of size principle for pluralities. To derive Replacement $^\diamond$, he adds a modal variant of (Ref).

Studd adds three axioms: a modalization of the axiom of Extensionality (intuitively: sets are identical if they have the same members at all stages) and axioms of Priority (intuitively: any set formed at a stage has its members formed at earlier stages) and Plenitude (intuitively: any sets ever formed will form a set at every later stage). From these, together with Max, Studd is able to derive all the axioms above except Infinity $^\diamond$ and Replacement $^\diamond$. To derive these, he adds a modal variant of (Ref) (which is stronger than Linnebo's).

But this does not help. Rather than replacing “the language of time and activity” with “the language of potentiality and actuality,” D3 and D4 proceed in the opposite direction, *reducing* “the language of potentiality and actuality” to “the language of time and activity”.

The modal theorist might seek to preserve whatever insight D1–D2, or D3–D4, provide, by “downgrading” them from official definitions to unofficial glosses. Both Linnebo and Studd are attracted by this idea. For example, while Linnebo (2010, 155) acknowledges that “officially” the modal operators are “new primitives governed only by a modal logic” he suggests that—“for heuristic purposes”—‘ $\Box\phi$ ’ may be read as: “no matter what sets we go on to form it will remain the case that ϕ ,” and ‘ $\Diamond\phi$ ’ may be read as “it is possible to go on to form sets so as to make it the case that ϕ .” Studd (2013, 701) suggests a similar unofficial reading of ‘ $\Box_{>}\phi$ ’ and ‘ $\Box_{<}\phi$ ’ as: “it will be the case at every later stage that ϕ ” and “it was the case at every earlier stage that ϕ ”. (Under the obvious extension to $\Diamond_{>}$ and $\Diamond_{<}$, ‘ $\Diamond_{>}\phi$ ’ is unofficially read as ‘it will be the case at some later stage that ϕ ’ and ‘ $\Diamond_{<}\phi$ ’ as ‘it was the case at some earlier stage that ϕ ’.)

Such paraphrases mislead in certain cases. Consider the (non-modal) statement ‘no sets include every set’, which may be expressed in a plural first-order language as:

$$(14) (\forall xx)(\exists y)(y \not\subseteq xx).$$

Following Linnebo, let ϕ^\diamond denote the result of applying the modal translation scheme to the non-modal formula ϕ of (plural) iterative set theory. The modal translation of (14) is:

$$(14)^\diamond \Box(\forall xx)\Diamond(\exists y)(y \not\subseteq xx).$$

Applying Linnebo’s paraphrase, this may be expressed as: ‘no matter what sets we go on to form, it will remain the case that, for any sets xx , it is possible to go on

to form a set y such that $y \not\prec xx'$. The original sentence (14) is false in standard iterative set theory, since the open formula $'(\exists y)(y \not\prec xx)'$ is not satisfied when $'xx'$ is assigned to all the sets. But $(14)^\diamond$ is true in the corresponding modal set theory, in which FPC^\diamond holds (and so, intuitively, it's always possible to form new sets).⁴¹

How serious an objection is this? It might be noted that (14) involves plural quantification, and consequently, even if it provides grounds for questioning Linnebo's gloss on modal operators, which is intended as a bridge between non-modal and modal formulations of set theory *in a plural language*, it does not apply to Studd's gloss on the modal operators, which is limited to singular formulations of set theory. On the other hand, it is now well known that iterative set theory can be formalized in a first-order plural language in which the use of schematic variables is replaced by plural quantification. Doing so has a number of advantages and has been subsequently adopted by a number of philosophers of set theory.⁴² Consequently, I maintain that it would be unacceptably arbitrary to adopt a modal interpretation of singular set theory and refuse to extend this to plural set theory. Studd cannot evade the objection by simply ignoring plural set theory.

Instead of explicating the operators by means of an official definition or an unofficial paraphrase, we might attempt to describe the intended modality more directly. For starters, we should note that this modality cannot be metaphysical. It is a widely held belief that mathematical objects are either metaphysically necessary or metaphysically impossible, and consequently that there is no middle ground which would allow for any sets to be metaphysically possible, but *not* metaphysically ac-

⁴¹Note that the same result holds if we replace $(14)^\diamond$ with its interpretation under a standard semantic model. The interpretation of $(14)^\diamond$ (at a world w_0) may be expressed as: 'for any (w_0 -accessible) world w , for any xx in w , there is a (w -accessible) world v at which there exists a set y such that $y \not\prec xx'$ '. The interpretation of $(14)^\diamond$ is true since every world (in the model) accesses worlds containing additional sets.

⁴²These include Lewis (1991), Uzquiano (2003), Burgess (2008), Linnebo (2010, 2013) and Oliver and Smiley (2013).

tual. Without such a middle ground, FPC^\diamond leads to contradictions. To see this, it suffices to note that relative to any sets, xx , considered to be actual, FPC^\diamond asserts the possible existence of their universal set U_{xx} . If U_{xx} 's possible existence implies its actual existence, then when xx are all sets considered to be actual, FPC^\diamond entails the actual existence of U_{xx} , which in this context is the (actual) universal set. This leads to familiar derivations of Cantor's and Russell's paradoxes. (Cantor's paradox is derived by considering U_{xx} 's subsets, which must paradoxically be more numerous than U_{xx} 's members; Russell's paradox is derived by substituting ' $x \notin x$ ' for ' ϕ ' in the Separation schema.) To restore consistency, we must understand \diamond in a way in which sets are (merely) possible relative to their elements. Doing so requires a special non-metaphysical modality.

A number of proposals have been made. Studd (2013, 706) suggests a linguistic modality, according to which the possible existence of a set relative to its members is to be understood in terms of a possible expansion of the set-theoretic vocabulary.⁴³ Since, under the intended interpretation, the set-theoretic vocabulary can *always* be expanded to accommodate new sets, there can be no interpretation of 'set' that is absolutely unrestricted and therefore the linguistic interpretation amounts to a form of modal restrictivism. This makes it unhelpful in the present context, in which discussion of the iterative conception, and hence the priority relation, is driven primarily by the goal of providing an absolutist reply to the set-theoretic paradoxes. Such a reply must provide an explanation for why the paradoxes are blocked that is *not* based on the premise that quantification over sets is restricted.

Linnebo (2010, 158) writes that he favors understanding the modalities in terms of a process of "individuating mathematical objects." He describes this process as follows:

⁴³See also Williamson (1998).

To individuate a mathematical object is to provide it with clear and determinate identity conditions. This is done in a stepwise manner, where at any stage we can make use of any objects already individuated in our attempts to individuate further objects.

I find this description to be rather puzzling. Is the individuation of mathematical objects something that *we* do? If so, then the modal account under Linnebo's interpretation would appear to devolve into a constructivist account, according to which sets are formed (or individuated) by some sort of mental activity. On the other hand, if the individuation of mathematical objects is not something that we do, then how might it be understood?

Simon Hewitt (2015, 325) writes that Linnebo's process of individuating mathematical objects should be understood in terms of possible extensions of a language.

What is meant by 'individuation'? Just this: I can individuate an x such that $\phi(x)$ iff there is a possible extension of my language such that there is a singular term ' a ' such that ' $\phi(a)$ ' is true.

Whatever its merits, Hewitt's interpretation reverts to the linguistic model for modality suggested by Studd, and we've already seen the inadequacy of this in the present context.

Alternatively, one might seek to understand the notion of individuation in a more robust, extra-linguistic sense, in terms of the metaphysical relation of *grounding* or *dependence*.⁴⁴

Let's grant for the sake of argument that these metaphysical relations provide a viable account of the priority relation. The question is whether this should be considered as an interpretation of the *modal account* of the priority relation. I think the answer must be no. The modal profile of dependence (or grounding) does not

⁴⁴In other writing, Linnebo expresses skepticism about such an approach. See Linnebo (2009, 222).

fit that of the modalities required by the modal account.⁴⁵ This is evidenced by the fact that in saying that one thing depends on some other things, or that one thing’s identity is grounded in the identity of some other things, we are not generally taken to imply that the one thing is in any sense merely possible relative to these others. Thus, one would not be accused of inconsistency if he said both (a) that Socrates’s singleton *depends* on Socrates and (b) that Socrates’s singleton must *actually* exist if Socrates does; but one would likely be so accused if instead of (a) he said (a*) that Socrates’s singleton is merely possible relative to Socrates.⁴⁶ I conclude that while it may be possible to provide a metaphysical account of the priority relation, this is best viewed as a non-modal alternative to the modal account of priority. I discuss this alternative in the next section.

In related work on the question of unrestricted quantification in mathematics, Fine (2006a,b) recommends a special “postulational” mathematical modality according to which the possible existence of a set relative to its members is to be understood in terms of a rule, or “procedural postulate”, for introducing it. An instance of such a rule is the procedural postulate

(15) Introduce the set of all sets (in the domain D).

Fine shows how (15) forces the expansion of D by introducing a universal set that (on pain of paradox) falls outside D . In Fine’s terms, this universal set is a (mere) postulational possibility relative to D .

In his (2013), Linnebo also advocates for a special mathematical modality. He likens this to the modality employed in the historical distinction between the notions of potential and actual infinity. He goes on to remark (207–208):

⁴⁵In speaking of “the modal profile” of dependence, I do not mean to suggest that dependence can be given a modal analysis.

⁴⁶For a modal set theory that seeks to capture the metaphysical modality expressed by claims such as (b), see Fine (1981).

This modality is tied to a process of building up larger and larger domains of mathematical objects. A claim is possible in this sense, if it can be made to hold by a permissible extension of the mathematical ontology; and it is necessary if it holds under any permissible such extension.

(Notice that here he does not speak about individuating objects but of extending mathematical ontology.)

For both Fine and Linnebo, the intelligibility of mathematical modality appears to be tied to the ability to understand a particular process (for Fine, this is the process by which new sets are postulated; for Linnebo it is the process by which mathematical domains are extended). Taken literally, these views are constructivist. However, any constructivist process effects a change in the world and Fine, at least, denies that postulation has this consequence. By postulating objects, we do not change the world; rather we change which world is under consideration. Fine explains:

On the present view, there is no such thing as *the* ontology, one that is privileged as genuinely being the sum-total of what there is. There are merely many different ontologies, all of which have the same right (or perhaps we should say no right) to be regarded as the sum-total of what there is.

This explanation is troubling insofar as it seems to involve a form of ontological relativism, according to which there are multiple (mathematical) ontologies and no principled way to privilege one over the others. In response, it might be argued that this is precisely how we should understand the modal account; that part of what it is for the set-theoretic hierarchy to be potential is for there to be no fact of the matter as to just how many sets are actual. Even if this is right, we are left with an interpretive dilemma: on the one hand, postulation is not to be understood in a constructivist sense, as effecting a genuine change in the world; on the other hand, it cannot be described in more familiar objectivist terms, such as grounding or metaphysical possibility. Similar remarks apply to Linnebo's notion of the extension of mathematical domains. If this is not a constructivist process and it is not to be understood in terms of grounding or metaphysical possibility, then how is it to be

understood? Perhaps there is a way that I have not considered; but I cannot think of what this might be. I conclude that none of the attempts by Studd, Linnebo and Fine to elucidate the modal operators succeeds.

6.4.3 The dependence account of priority

The dependence account of priority may be motivated by reflection on the extensional identity conditions for sets. Recall that according to the axiom of extensionality, two sets are identical iff their members are identical. Viewed metaphysically, extensionality may be taken to suggest that sets depend (fully) on their members, or equivalently, that the identity of sets is (fully) grounded in the identity of their members.⁴⁷ One may then claim that the priority of the members to their set consists in the metaphysical fact that while their set depends fully on them (and partially on each one of them), they do not depend on it (either fully or partially).

Unfortunately, this account of priority entails the existence of the contradiction-inducing sets. To see this, note that if x fully depends on yy (perhaps related in some way), it is natural to think of x as nothing over and above the yy (and the relations holding between them). All God has to do to create x is to create yy (and to relate them in the right way). Thus, plausibly, a table fully depends on its parts (related in the right way). It is nothing over and above them (arranged in the right way). All God must do to create the table is to create the parts (and to arrange them in the right way). This presents an immediate problem for the iterative conception. If the set x fully depends on its members yy , then all God has to do to create x is to create the yy . (Unlike the table, any additional relations between the yy are irrelevant, since the identity conditions for sets do not mention them.) Thus, all God must do

⁴⁷Intuitively, x depends on y if what it is to be x involves y . See Fine (1994), Fine (1995a), Fine (1995b), and Tahko and Lowe (2015) for more detail. For a somewhat different formulation, see Schaffer (2009) and Rosen (2010). When an individual x depends on a (possibly infinite) number of distinct individuals, yy , perhaps related in some way, and on nothing else, x *fully depends* on yy so related. If x fully depends on yy , then x *partially depends* on each y among yy .

to create the universal set is create all the sets. But the quantifiers in iterative set theory are taken to range over all the sets. So if priority is understood as dependence, iterative set theory is committed to the existence of a universal set.

One might reply that the universal set cannot possibly exist, since if it did, there would be a set which contained itself as a member and which therefore depended (partially) upon itself. Since this is impossible (for dependence is generally held to be irreflexive), the universal set cannot possibly exist. However, what this shows is not that the derivation of the universal set from the dependence account of priority fails, but rather that the dependence account is inconsistent, entailing both that the universal set exists and that it does not exist. It seems, therefore, that the only suitable response to the argument above is to reject the claim that a set is nothing over and above its members. Creating the members of a set is *not* all that God must do to create the set. What else must God do? *Form* the set? *Collect* its members? It seems that any natural answer to this question forces a new metaphor on us. If dependence can make sense of priority only by forcing us to accept another metaphor, it is an unsatisfactory answer to our initial project of freeing the iterative conception from metaphor.

6.5 Conclusion

The iterative conception of *set* purports to motivate a consistent set theory in which quantification over sets is absolutely unrestricted. I have argued that it fails in two respects: (1) it is unable to satisfactorily motivate the Empty Set axiom of iterative set theory (and consequently is also unable to motivate the axiom of Infinity); (2) it is unable to provide a satisfactory account of the priority relation. As a result, the iterative conception cannot explain the extensional correctness of iterative set theory: it can neither explain why all the sets that exist according to iterative set theory exist, nor why these are all the sets that exist. In particular, it cannot provide

a satisfactory explanation for why none of the contradiction-inducing sets exists. I conclude that the iterative conception fails to provide the absolutist with a theory of sets that is superior to the theory of set restrictivism presented in chapter 4.

CHAPTER 7

SETLESS VARIANTS OF THE PARADOXES

In this chapter, I argue that unrestricted quantification together with “naive” comprehension principles governing the existence of interpretations, propositions and ordinals lead to setless variants of the paradoxes of Russell, Cantor and Burali-Forti. One might seek to block these paradoxes by extending the orthodox strategy for sets (discussed in ch. 6). This approach would consist of replacing these naive comprehension principles with restricted comprehension principles. However, such restrictions are unintuitive and ad hoc. Alternatively, one might seek to block these paradoxes by restricting quantification. This latter strategy is attractive since it provides a uniform solution to the paradoxes (in both their original and setless forms) that eludes the absolutist.

The outline for this chapter is the following. In 7.1, I apply unrestricted singular quantification over interpretations to derive a setless (semantic) version of Russell’s paradox. In 7.2, I apply unrestricted plural quantification over objects to derive a setless version of Cantor’s paradox. In 7.3, I apply unrestricted plural quantification over ordinals to derive a setless version of Burali-Forti’s paradox. These derivations are not new and much of the work in this chapter consists in applying previous work to the question of absolutism and restrictivism. In particular, 7.1 is based on Williamson (2003); 7.2 is based (more loosely) on Spencer (2012) and 7.3 is organized around Hellman (2011).

7.1 Russell’s Paradox: Semantic Version

Timothy Williamson (2003) presents a setless version of Russell’s paradox, which employs a “naive” comprehension principle for interpretations of the predicate expressions of a language. The principle, which I dub ‘Interpretation Comprehension’ (IC), states that if ‘P’ is a predicate letter of an uninterpreted object language L and ϕ is a meaningful predicate (of the right form) in L’s metalanguage, then there exists an interpretation of L on which ‘P’ means ϕ . IC is officially defined as follows:

IC. If (i) L is an object language containing the n -place predicate letter ‘P’, (ii) \mathcal{L} is a metalanguage used to interpret L, and (iii) \mathcal{L} contains the meaningful n -place predicate letter \mathcal{P} , then there is an interpretation, $I_{\mathcal{P}}$, under which ‘P’ means \mathcal{P} .

Several quick observations are in order. First, IC is limited in scope to interpretations of the predicate expressions of a language. In what follows, I will refer to these simply as ‘interpretations’. Second, IC does not purport to offer *necessary* conditions on the existence of interpretations. In this sense, IC is not as informative about interpretations as FC is about sets. Third, IC is relativized to metalanguages. As a result, if there are possible meanings for the predicate vocabulary in L which are not captured by the predicates in any metalanguage, \mathcal{L} , then, for all IC says, there may be fewer interpretations of L than possible assignments of meanings to its nonlogical expressions.¹ Fourth, if we restrict ourselves to extensional languages, we can define (in \mathcal{L}) what it is for a predicate ‘P’ to *mean* \mathcal{P} under an interpretation I.

D1. ‘P’ *means* \mathcal{P} under I =_{df.} $(\forall x)(x \text{ satisfies ‘P’ under I iff } \mathcal{P}x)$

¹This seems wrong. But it would refute IC only if IC purported to offer necessary conditions on the existence of interpretations. These last two observations suggest that a strengthened version of IC, one that would provide necessary and sufficient conditions for the existence of interpretations, is possible only given two modifications. First, the strengthened principle must say something about semantic assignments to names. Second, there must be some guarantee that the metalanguages contain enough predicates. Would it be acceptable to simply postulate a maximal metalanguage $\mathcal{L}+$ containing a predicate for every possible meaning?

IC is quite intuitive. Suppose you're considering possible interpretations of the uninterpreted predicate letter 'P' relative to English. Is it possible to interpret 'P' as meaning 'is a philosopher'? Is it possible to interpret 'P' as meaning 'runs'? It's difficult to see why not. Generalizing, it follows that for any meaningful English predicate ϕ of the correct form, it's possible to interpret 'P' as meaning ϕ . In addition to its intuitiveness, Williamson (426) argues that this principle is presupposed by the model-theoretic (Tarskian) definition of logical consequence in terms of truth-preservation under all interpretations.

In principle, when we apply the definition of logical consequence, it must be possible to interpret a predicate letter according to any contentful predicate, since otherwise we are not generalizing over all the contentful arguments of the right form. Thus, whatever predicate we substitute for [' ϕ '], some legitimate interpretation (say, $I[(\phi)]$) interprets the predicate letter ['P'] accordingly.

In what follows, I will assume that L is an extensional first-order language, containing the one-place predicate letter 'P', and that English (or a simple extension of English) is the metalanguage \mathcal{L} in which L is interpreted. These assumptions satisfy conditions (i)–(ii) of IC. If, in addition, a meaningful one-place predicate of English (or a simple extension of English) is substituted for \mathcal{P} , then condition (iii) of IC is satisfied. For example, if the English predicate 'is a cat' is substituted for \mathcal{P} , then IC entails that there is an interpretation, I_{cat} , under which 'P' means 'is a cat'. Applying D1, IC entails:

- $(\forall x)(x \text{ satisfies 'P' under } I_{\text{cat}} \text{ iff } x \text{ is a cat}).$

Or suppose that we select the English predicate 'runs'. Then IC entails that there is an interpretation, I_{runs} , under which 'P' means 'runs'. Applying D1, IC entails:

- $(\forall x)(x \text{ satisfies 'P' under } I_{\text{runs}} \text{ iff } x \text{ runs}).$

In general, for any meaningful one-place English predicate ϕ , IC entails the schema

$$(1) (\exists y)(\forall x)(x \text{ satisfies 'P' under } y_{\phi} \leftrightarrow \phi x).$$

7.1.1 Derivation of the contradiction

To derive the paradox, we introduce an extension of English (English+) by the addition of the predicate ‘ \mathcal{R} ’ which is defined by the condition:

$$(2) (\forall x)(\mathcal{R}x \leftrightarrow x \text{ does not satisfy ‘P’ under } x).$$

Since \mathcal{R} is a meaningful predicate in English+, it can be substituted for ‘ ϕ ’ in (1) to get:

$$(3.1) (\exists y)(\forall x)(x \text{ satisfies ‘P’ under } y_{\mathcal{R}} \leftrightarrow \mathcal{R}x).$$

Since L is extensional, there is a unique value of y that make (3.1) true. This is the interpretation under which ‘P’ means ‘ \mathcal{R} ’, which I dub ‘ $I_{\mathcal{R}}$ ’. We then have

$$(3.2) (\forall x)(x \text{ satisfies ‘P’ under } I_{\mathcal{R}} \leftrightarrow \mathcal{R}x).$$

By (2) and (3.2), it follows that:

$$(3.3) (\forall x)(x \text{ satisfies ‘P’ under } I_{\mathcal{R}} \leftrightarrow x \text{ does not satisfy ‘P’ under } x).$$

Instantiation to $I_{\mathcal{R}}$ then yields the inconsistent

$$(3.4) I_{\mathcal{R}} \text{ satisfies ‘P’ under } I_{\mathcal{R}} \leftrightarrow I_{\mathcal{R}} \text{ does not satisfy ‘P’ under } I_{\mathcal{R}}.$$

7.1.2 Evaluation

One might seek to block the argument by claiming that \mathcal{R} is not a meaningful predicate and therefore (3.1) is not a permissible substitution instance of the schema (1). Bennet and Karlsson (2008) take this approach, arguing that since IC says nothing about what predicates are meaningful, we are free to view the derivation of the contradictory (3.4) as a reductio of the meaningfulness of \mathcal{R} . On their view, Williamson’s argument is a “pseudo-paradox” similar to Russell’s “pseudo-paradox” of the barber.²

²The paradox the barber consists in defining some individual, b , as the barber who shaves all and only those who don’t shave themselves and proving that if b exists, the contradiction follows that he

[IC] does not seem to have anything to do with $[\mathcal{R}]$ being a contentful predicate. We may deny this without giving up [IC] as a true principle . . . In this sense Williamson’s argument is more like the non-paradox of the barber, i.e., it is a *reductio* of \mathcal{R} being a contentful predicate (324)

It is true that IC does say nothing about whether \mathcal{R} is meaningful. But note that in the derivation of Russell’s original paradox, FC likewise says nothing about the meaningfulness of the predicate ‘ $x \notin x$ ’. If Williamson’s argument is a pseudo-paradox, then why isn’t Russell’s original argument a pseudo-paradox as well? The original derivation of the contradictory $R \in R \leftrightarrow R \notin R$ is startling because, prior to recognizing its implications, no one would doubt that ‘ $x \notin x$ ’ is meaningful. We do not require a principle for the meaningfulness of predicates in order for this to strike us as a genuine paradox. It just seems so obvious that the contradiction-inducing predicates are meaningful. This appearance is confirmed by the fact that we can easily run through cases and see what sets would satisfy ‘ $x \notin x$ ’ and what sets would not. The set of all Greek philosophers is not a Greek philosopher and therefore is not a member of itself. But the set of all sets is a set and therefore (if it exists) is a member of itself. Similarly with the set of all not-men (which is not a man) and the set of all abstract objects (which, presumably, is abstract).

In Williamson’s argument, (2) clearly states the conditions under which something \mathcal{R} s (and the conditions under which something does not \mathcal{R}). Something \mathcal{R} s if (a) it is an interpretation of ‘P’ that does not fall under ‘P’ so interpreted (this is true of the interpretation I_{cat} , which is not itself a cat) or (b) it is not an interpretation (if x is not an interpretation, then a fortiori, x is not an interpretation under which x satisfies ‘P’). On the other hand, something doesn’t \mathcal{R} if it is an interpretation of ‘P’

shaves himself if and only if he does not. However, there is no independent reason for believing in the existence of b . In particular, there is no “comprehension principle” for barbers to which we find ourselves antecedently committed. In the absence of some good reason for believing that b exists, we can take the argument as a *reductio* of this assumption.

that does fall under ‘P’ so interpreted. An example is the interpretation I_{int} (which interprets ‘P’ as meaning ‘is an interpretation’). For I_{int} *is* an interpretation.

I conclude that (2) provides us with powerful evidence—which is independent of IC—that \mathcal{R} is a meaningful predicate. Consequently, Bennet and Karlsson have failed to provide a convincing case for denying (3.1). But once we grant this premise, the rest of the reasoning is simply a matter of applying the rules of quantified logic. For if (3.1) is true, then $I_{\mathcal{R}}$ exists and is within the range of the unrestricted quantifier \forall in (2) and (3.2). Consequently, $I_{\mathcal{R}}$ is within the range of the unrestricted quantifier in (3.3) as well. This is all that is required to justify the instantiation to the contradictory (3.4).

Williamson’s own solution is to treat quantification over interpretations as second-order. Thus, (1) is to be rewritten as

$$(4) (\exists Y)(\forall x)(x \text{ satisfies ‘P’ under } Y_\phi \leftrightarrow \phi x)$$

in which Y is a second-order variable. In order to preserve the absolutist thesis that absolutely unrestricted quantification is possible, Williamson argues that interpretations are not objects of any kind; the second-order quantification used to interpret L is ontologically innocent. However, this strategy leads to a version of the inexpressibility problem levied against restrictivism. For, in order to understand this second-order quantification, we must have some way of providing a semantics for it. If we attempt to do this by second-order quantification over interpretations of this second-order language, the contradiction re-emerges. Williamson must therefore appeal to interpretations from a higher (third-order language). In order to provide a semantics for this third-order language, he must appeal to interpretations of a still higher (fourth-order) language and so on. At each type, the contradiction would re-emerge if interpretations of the n th-order variables had type n . To avoid this, we continually promote them to higher and higher types. The generalization of Williamson’s solution therefore amounts to the adoption of a simple type theory in which type-restrictions

prevent the paradox from being reinstated. But then there are certain semantic facts about this type theory that go inexpressible. Linnebo (2006) describes two:

- Infinity: there are infinitely many different kinds of semantic value, and
- Compositionality: the semantic value of a complex expression is determined as a function of the semantic values of the expression's simpler constituents.

Within the type restrictions needed to block the paradox we cannot express either (since these involve quantification over semantic values of different logical types).

Insofar as the absolutist's solution to this paradox commits her to semantic inexpressibility, she seems to fare no better than the restrictivist. Indeed, she may fare worse. For the restrictivist is able provide a uniform solution to Russell's paradox in both its original and semantic forms; whereas the absolutist must now rely on a two very different solutions. She solves Russell's original paradox by rejecting the intuitive principle of FC; but solves the semantic variant by adopting a simple type theory.

The restrictivist solves Williamson's paradox by restricting quantification over interpretations in (1) and therefore also its instance (3.1). The mirrors the solution to Russell's original paradox, which consisted in restricting quantification over sets in FC (and in its instances). (3.1) is rewritten as:

$$(3.1^*) (\exists y)(\forall_D x)(x \text{ satisfies 'P' under } y_{\mathcal{R}} \leftrightarrow \phi x)$$

Since quantification is restricted, we are free to infer that the interpretation $I_{\mathcal{R}}$ falls outside of D in (3.1) and consequently that the instantiation to $I_{\mathcal{R}}$ at (3.4) is illegitimate.

7.2 Cantor's Paradox: Plural Version

I have used plural quantification and plural logic previously in my discussions of extensions (4.3.1) and iterative set theory (ch. 6); however a brief restatement is in

order here. Plural logic enriches singular logic with devices for plural reference. Add first-order plural variables ($'xx'$, $'yy'$), plural quantifiers ($'\exists xx'$ and $'\forall xx'$) and a logical connective $'\prec'$. The plural variables can take individuals as minimal (singular) values. This accords with English, in which plural expressions sometimes refer to individuals (e.g., 'the even prime numbers'). When an individual is understood as the singular value of a plural variable, I will call it a *singular plurality*. Read $'\exists xx \dots'$ as 'there are some—i.e., one or more— $xx \dots$ '. Read $'\forall xx'$ as 'for any—i.e., one or more— $xx \dots$ '. Officially, the connective \prec takes only singular terms in its first argument place and plural terms in its second argument place. Read $'x \prec yy'$ as ' x is one of (among) yy '. Unofficially, \prec can take plural terms in its first argument place. Read $'xx \prec yy'$ as an abbreviation for $'(\forall x)(x \prec xx \rightarrow x \prec yy)'$. Thus, we read 'The even numbers \prec the numbers' as an abbreviation of 'for any x : if x is one of the even numbers, then x is one of the numbers'.

7.2.1 Cantor's theorem: plural version

The plural variant of Cantor's paradox (henceforth, simply Cantor's paradox) consists of applying a plural variant of Cantor's theorem (henceforth, simply Cantor's theorem) to the plurality of all things. Call the yy a *sub-plurality* of the xx if $(\forall z)(z \prec yy \rightarrow z \prec xx)$ ³ and let $\mathcal{P}(xx)$ denote the plurality of all sub-pluralities of the xx , i.e., $(\forall yy)(yy \prec \mathcal{P}(xx) \leftrightarrow (\forall z)(z \prec yy \rightarrow z \prec xx))$. We might read $\mathcal{P}(xx)$ as 'the *power plurality* of the xx '. If, for example, $xx =$ the odd numbers, then $\mathcal{P}(xx)$ is the plurality of all pluralities of odd numbers. The odd numbers less than 10 are one of these pluralities; the odd numbers identical to 15 are another; the odd prime numbers a third. Intuitively, Cantor's theorem tells us that any plurality of (individual) things contains more sub-pluralities than things and so there can be no bijection between

³Note that according to this definition, every plurality is an (improper) sub-plurality of itself.

the plurality and its power plurality. The proof has two features that are worth mentioning.

First, it makes use of higher-order plural quantification (plural quantification over pluralities, i.e., plural quantification over things plurally). The cogency of such quantification has been called into question and this makes the proof controversial.

Second, the plural variant of Cantor's theorem is not necessarily as general as the singular version. In particular, if there is no empty plurality, it does not hold for "singular" pluralities, which consist of a single element. Such pluralities have the same size as their power plurality, namely 1.

Let us then assume that the aa are a plurality of at least two elements. Assume, in addition, that the notion of a cardinality is extended to pluralities in the natural way, as a measure of size, so that:

D2. $\text{Card}(xx) = \text{Card}(yy) =_{df}$ there is a bijection between xx and yy .

D3. $\text{Card}(xx) \leq \text{Card}(yy) =_{df}$ there is an injection from xx to yy .

D4. $\text{Card}(xx) < \text{Card}(yy) =_{df}$ (i) $\text{Card}(xx) \leq \text{Card}(yy)$ and

(ii) $\text{Card}(xx) \neq \text{Card}(yy)$.⁴

To show $\text{Card}(aa) < \text{Card}(\mathcal{P}aa)$ it suffices to show (i) $\text{Card}(aa) \leq \text{Card}(\mathcal{P}aa)$ and (ii) $\text{Card}(aa) \neq \text{Card}(\mathcal{P}aa)$. Since every individual a is a minimal plurality, the first part is quite simple. Map every $a \prec aa$ to itself. To prove the second part, we will show that there can be no bijection between $\mathcal{P}aa$ and aa .

Assume for reductio that f is a bijection between $\mathcal{P}aa$ and aa , i.e., f pairs every sub-plurality $xx \prec aa$ with a distinct individual $f(xx) = y$, where $y \prec aa$. Define the ww as the sub-plurality consisting of every $y \prec aa$ that is not among the xx with

⁴See discussion of cardinalities in ch. 2.

which it is paired by f . If we let ‘ y ’ range (singly) over the aa , we can state this definition more succinctly as follows:

$$(5.1) \quad (\forall y)(y \prec ww \leftrightarrow y \not\prec f^{-1}(y))$$

The proof has two parts. First, we show that the ww exist, i.e., that the plurality defined by (5.1) is nonempty. Next, we instantiate the ww to the individual paired with them to derive a contradiction.

Consider the singular pluralities $xx_1 = y_1$ and $xx_2 = y_2$. Either $y_1 \prec f^{-1}(y_1)$ or $y_1 \not\prec f^{-1}(y_1)$. Similarly, either $y_2 \prec f^{-1}(y_2)$ or $y_2 \not\prec f^{-1}(y_2)$. If either $y_1 \not\prec f^{-1}(y_1)$ or $y_2 \not\prec f^{-1}(y_2)$, then there is at least one $y \not\prec f^{-1}(y)$ and consequently, the ww are nonempty. On the other hand, if $y_1 \prec f^{-1}(y_1)$ and $y_2 \prec f^{-1}(y_2)$, then consider the plurality $xx_3 = y_1, y_2$. Since f is a bijection, $f(xx_3) = y_3$, where y_3 is distinct from both y_1 and y_2 . Since $y_3 \not\prec xx_3$ and $xx_3 = f^{-1}(y_3)$, it follows that $y_3 \not\prec f^{-1}(y_3)$. Again, there is at least one $y \not\prec f^{-1}(y)$ and consequently, the ww are nonempty.

Since the ww are a sub-plurality of the aa , they are included in the mapping f . Consider the individual—call it y_{ww} —that f pairs with the ww . Instantiating (5.1) to y_{ww} , it follows that

$$(5.2) \quad y_{ww} \prec ww \leftrightarrow y_{ww} \not\prec f^{-1}(y_{ww}).$$

But $f^{-1}(y_{ww}) = ww$, so

$$(5.3) \quad y_{ww} \prec ww \leftrightarrow y_{ww} \not\prec ww.$$

Contradiction. By reductio, we conclude that f is not a bijection. Consequently, $\text{Card}(aa) < \text{Card}(\mathcal{P}aa)$.

7.2.2 Derivation of the contradiction

The plural variant of Cantor’s paradox follows the reasoning of Cantor’s paradox I.⁵ It consists of a proof that a particular plurality of propositions has the same size as its power plurality. This proof is based on an existence principle for propositions, which states that for any things, there is a distinct proposition which states that these are all the things there are. Following Joshua Spencer (2012), I will refer to this as the Existential Propositions Thesis (EPT).⁶ Officially, PT is formulated as follows:

(PT) $(\forall xx)((i)$ there is a proposition p stating that xx are all the things there are and (ii) p is distinct from any proposition q stating that $(\exists yy)(yy \neq xx \wedge yy$ are all the things there are)).

I will refer to the propositions whose existence PT asserts as ‘existence propositions’. PT implies that there is an injection f from the power plurality of all things, $\mathcal{P}(uu)$, into the plurality of all things, uu . Define f so that it pairs every sub-plurality $xx \prec uu$ with the existence proposition $f(xx) = p$, which states that the xx are all the things there are. Given PT, distinct sub-pluralities must be paired with distinct propositions. Hence, by D3,

$$(6.1) \text{Card}(\mathcal{P}(uu)) \leq \text{Card}(uu).$$

This contradicts Cantor’s theorem, according to which

$$(6.2) \text{Card}(uu) < \text{Card}(\mathcal{P}uu).^7$$

⁵It should also be possible to derive a plural variant following the reasoning of Cantor’s paradox III. Since (by D2–D4) cardinal numbers are attached to pluralities of the requisite cardinality (not sets), the argument works by (i) replacing each cardinal number, κ , among the plurality of all cardinals, $\kappa\kappa$, with a plurality $f(\kappa)$ of cardinality κ and (ii) forming the union of these pluralities, $\bigcup(\forall\kappa)f(\kappa)$. This union has the greatest cardinal number; yet, by Cantor’s theorem, its power plurality has an even greater cardinal number. Therefore, there can be no plurality $\kappa\kappa$ of all cardinal numbers.

⁶Spencer’s principle is slightly different. He expresses it as follows (77): “for any xs , (i) there is the proposition that those xs exist and (ii) for any ys which are not the same as the xs , the proposition that those ys exist is distinct from the proposition that those xs exist.”

⁷I am assuming the natural extension of SBT (see ch. 2) from sets to pluralities.

7.2.3 Evaluation

What does this argument show? Spencer takes it as a proof of the ontological thesis that there is no all-inclusive plurality. If there is no all-inclusive plurality, then it is impossible to quantify plurally over all things and consequently impossible to define the injection from all pluralities of things into all things that is needed to derive the contradiction above. The restrictivist solution I recommend is importantly different. I do not deny the existence of an all-inclusive plurality; but directly deny the possibility of unrestricted quantification over all things.

The crux of the argument is PT. What reasons—other than the threat of paradox—are there for doubting it? Given that the universe is populated by a unique plurality, all existence propositions but one are false; but this is not a reason to doubt PT. False propositions exist just as much as true propositions. Moreover, there are numerous other principles that entail the existence of an injection from $\mathcal{P}(uu)$, into uu . Joshua Spencer (2012, 78–79) mentions two: the Compositional Propositions Thesis (CPT) and the Doxastic Propositions Thesis (DPT), which he formulates as follows:

- (CPT) For any xs , (i) there is the proposition that those xs compose something and (ii) for any ys which are not the same as the xs , the proposition that those ys compose something is distinct from the proposition that these xs compose something.
- (DPT) For any xs , (i) there is the proposition that those xs are thought about by someone and (ii) for any ys which are not the same as the xs , the proposition that those ys are thought about by someone is distinct from the proposition that those xs are thought about by someone.

It is difficult to see how these principles might be reasonably denied. Any restriction on them would appear to be ad hoc. To illustrate, if (EPT) is false, then there are some peculiar things such that there is no proposition asserting that they are all the things. Why is there no such proposition for these things, when there are such propositions for other pluralities? Certainly, we would not deny the existence

of existential propositions for most familiar pluralities. For example, we would not deny that there is a proposition that asserts that the items on my desk are all the things that there are. Of course, this proposition is false, but that does not matter. Similar remarks apply to CPT and DPT. A more principled approach is to adopt a sparse account of propositions, according to which propositions are identified with sets (or pluralities) of possible worlds. But this view has some rather unintuitive consequences. Perhaps the most serious is that it implies that every necessarily true proposition is identical (and that every necessarily false proposition is identical).

There is an argument (due to Spencer) that PT is inconsistent. Some terminology is needed. Say that an existential proposition p *selects* an individual x iff x is among the things that p says there are. Define a *self-effacing* proposition as an existential proposition that does not select itself. Clearly, some existential propositions are self-effacing. An example is the existential proposition mentioned above which selects only items on my desk. Let the pp refer to all the self-effacing propositions and let r name the existential proposition which selects all and only those propositions that do not select themselves, i.e., r selects the pp . Given the obvious parallels between r and the Russell set, call r the Russell proposition. Then the following is true:

$$(7.1) \quad (\forall p)(r \text{ selects } p \leftrightarrow p \text{ does not select } p).$$

By instantiation on (7.1), it follows that:

$$(7.2) \quad r \text{ selects } r \leftrightarrow r \text{ does not select } r.$$

Given standard propositional reasoning, (7.2) entails the contradiction:

$$(7.3) \quad r \text{ selects } r \wedge r \text{ does not select } r.$$

This argument is clearly a variant of Russell's paradox. If the reasoning is sound, it would show that PT was inconsistent on its own, which would undermine the derivation of Cantor's paradox above.

Russell’s paradox in its original form is typically understood as a proof that the Russell set does not exist and consequently that FC is false. This is the No Set solution to described in ch. 2. We might be tempted to regard this argument in the same way: as a proof that the Russell proposition does not exist and consequently that PT is false. This is not what Spencer does. To rescue PT from this argument, Spencer recommends abandoning the relevant instance of excluded middle—either r selects r or r does not select r —needed to derive (7.3) from (7.2). His solution is a version of the No Function solution to Russell’s paradox described in ch. 2.

In ch. 4, I recommended a restrictivist solution to Russell’s paradox, according to which FC may be regarded as true, provided it is always interpreted in contexts in which quantification over sets is restricted. It’s possible to extend this solution to the present case so that PT may be regarded as true, provided it is always interpreted in contexts in which quantification over propositions is restricted. Whenever we speak of ‘all’ self-effacing propositions, we are speaking of all the self-effacing propositions that belong to a restricted domain. To represent this, replace the (purportedly) unrestricted plural term ‘ pp ’ with the (explicitly) restricted plural term ‘ pp_D ’ and replace the unrestricted plural quantifier $\forall xx$ with the restricted plural quantifier $\forall_D xx$. PT may then be rewritten as:

PT* $(\forall_D xx)(\text{there is a proposition } p \text{ stating that } xx \text{ are all the things there are and } p \text{ is distinct from any proposition } q \text{ stating that: } (\exists_D yy)(yy \neq xx \wedge yy \text{ are all the things there are})).$

Given any domain, D , PT* determines a corresponding domain, D_p , of existential propositions. Relative to D_p , it is possible to define the Russell proposition, r_D , by:

$$(8.1) (\forall_{D_p} p)(r_D \text{ selects } p \leftrightarrow p \text{ does not select } p).$$

But here the derivation comes to a halt. If r_D was in the range of \forall_{D_p} we could instantiate ‘ p ’ to r_D in (8.1) to get:

(8.2) r_D selects $r_D \leftrightarrow r_D$ does not select r_D .

However, since \forall_{Dp} is restricted, this assumption is open to question. We conclude by reductio that r_D lies outside the range of \forall_{Dp} ; consequently that:

(8.3) r_D does not select r_D .

As before, we can shift to a more inclusive context, in which PT^* determines a more inclusive domain of existential propositions, which includes r_D ; relative to this domain, (8.1) defines a different Russell proposition. We conclude by reductio that this Russell proposition lies outside the expanded range of quantification.

7.3 Burali-Forti Paradox: Plural Version

The plural variant of Burali-Forti's paradox is based on the natural extension of well-orderings to pluralities and of ordinal numbers to well-ordered pluralities, as measures of length. Say that any xx are *well-ordered* by the relation R just in case:

- R is transitive on the xx (i.e., for any $x_1, x_2, x_3 \prec xx$: if $x_1 R x_2$ and $x_2 R x_3$, then $x_1 R x_3$)
- The xx satisfy trichotomy (i.e., for any $x_1, x_2 \prec xx$: exactly one of the cases (a) $x_1 R x_2$, (b) $x_2 R x_1$, (c) $x_1 = x_2$ holds)
- Any $yy \prec xx$ have a least element under R

In what follows, I will use ' $<$ ' (with or without subscripts) as a symbol for well-ordering relations. When speaking of the ordinal numbers, $<$ (without subscripts) will be used to denote the "less than" relation defined on the ordinals. This is appropriate since this relation can be shown to be a well-ordering of the ordinals (see PBF1 below). I write ' $<_{xx}$ ' to denote the well-ordering relation $<$ on the xx and ' $(xx, <_{xx})$ ' to denote the xx when they are well-ordered by $<_{xx}$. Relations are typically identified with sets; however, here they are identified with pluralities of ordered pairs. Thus,

$<_{xx}$ is identified with the plurality of ordered pairs $\langle x_1, x_2 \rangle$ where $x_1, x_2 \prec xx$ and $x_1 < x_2$.⁸

Ordinal numbers are governed by the following two definitions:

D5. $\text{Ord}(xx, <_{xx}) = \text{Ord}(yy, <_{yy}) =_{df}$ there is an isomorphism between $(xx, <_{xx})$ and $(yy, <_{yy})$.

D6. $\text{Ord}(xx, <_{xx}) < \text{Ord}(yy, <_{yy}) =_{df}$ there is a $y \prec yy$ such that $(xx, <_{xx}) \cong \text{seg}(y)$.

Intuitively, an isomorphism is an order-preserving bijection. This can be defined more formally in the plural context as follows: $(xx, <_{xx})$ are isomorphic to $(yy, <_{yy}) =_{df}$ there is a bijection f between xx and yy that satisfies the condition $\langle x_1, x_2 \rangle \prec (xx, <_{xx})$ iff $\langle f(x_1), f(x_2) \rangle \prec (yy, <_{yy})$. An initial segment is defined in the plural context as follows: if $x \prec (xx, <_{xx})$, then the zz defined by the condition $z \prec zz \leftrightarrow z <_{xx} x$ are an initial segment of $(xx, <_{xx})$, written ‘ $\text{seg}(x)$ ’. Intuitively, if $(xx, <_{xx})$ are isomorphic to an initial segment of $(yy, <_{yy})$, then $(xx, <_{xx})$ are “longer” and have a greater ordinal number. From D5 and D6, it’s possible to establish plural variants of BF1 and BF2. Using ‘ α ’ and ‘ $\alpha\alpha$ ’ as singular and plural variables for ordinals, these may be stated as:

PBF1. $(\forall \alpha\alpha)(\alpha\alpha \text{ are well-ordered by } <)$

PBF2. $(\forall \alpha)(\alpha = \text{Ord}(\alpha\alpha, <))$.⁹

7.3.1 Derivation of the contradiction

The derivation of a contradiction follows the pattern of Burali-Forti III. Let ‘ oo ’ denote all the ordinals. By instantiation of PBF1 to oo , it follows that the oo are

⁸See Hellman (2011). We needn’t take a stand on what ordered pairs are. If we want, we can adopt the standard Kuratowski definition of ordered pairs with sets.

⁹For any ordinal number α , ‘ $\alpha\alpha$ ’ denotes the ordinals $< \alpha$.

well-ordered by $<$, so $(oo, <)$ have an ordinal number. Call this ‘ Ω ’. So we have it that $\text{Ord}(oo, <) = \Omega$. By instantiation of PBF2 to Ω , it follows that $\Omega = \text{Ord}(\Omega\Omega, <)$. These results establish the identity

$$(9.1) \quad \text{Ord}(\Omega\Omega, <) = \Omega = \text{Ord}(oo, <).$$

But $(\Omega\Omega, <)$ is an initial segment of $(oo, <)$ (intuitively, $(oo, <)$ are longer than $\Omega\Omega_{<}$ by having Ω tacked on at the end.) So, by D6

$$(9.2) \quad \text{Ord}(\Omega\Omega, <) < \text{Ord}(oo, <).$$

(9.1) and (9.2) immediately imply

$$(9.3) \quad \Omega < \Omega,$$

which contradicts PBF1, according to which

$$(9.4) \quad \Omega \not< \Omega.$$

7.3.2 Evaluation

As I noted in ch. 2, Burali-Forti’s paradox has many moving parts and each of these may be singled out as the source of the contradiction. The immediate premises used in the derivation are PBF1 and PBF2. However, rejecting either of these is tantamount to abandoning a theory of ordinals in this plural context. If there is no theory of ordinals, then there is no inconsistent theory of ordinals, but surely it is preferable to find a less draconian solution.

In the set-theoretic case, there are two familiar solutions. One is to deny that the set of all ordinals exists. This is the No Set solution favored by iterative theorists. The analogue in this case would be to deny that all the ordinals exist. Arguably, the idea that all the ordinals do not exist may be sussed out, along the lines suggested by modal interpretations of the iterative conception of set as the idea that all the ordinals do not actually exist. However, we’ve seen the problems faced by such modal accounts.

Alternatively, one might adopt a version of the limitation of size doctrine and claim that some pluralities of well-ordered ordinal numbers are absolutely infinite and do not have an ordinal number. This is a version of the No Function solution. It would amount to denying PBF2 and is therefore highly revisionist. In addition, it strikes me as unacceptably ad hoc. Intuitively, all that is needed to have an ordinal is to be well-ordered. Being well-ordered seems to have very little to do with size.

Alternatively, the argument may suggest that it is the linguistic assumption that ‘all the ordinals’ are a permissible value of the plural quantifier ‘any ordinals’ in PBF1 that is at fault. Whenever we speak of ‘all’ ordinals, we are speaking of all the ordinals that belong to a restricted domain. To represent this, replace the (purportedly) unrestricted plural term ‘ oo ’ with the (explicitly) restricted plural term ‘ oo_D ’ and replace the unrestricted quantifiers $\forall\alpha$ and $\forall\alpha\alpha$ with the restricted quantifiers $\forall_D\alpha$ and $\forall_D\alpha\alpha$. PBF1 and PBF2 may then be rewritten as:

PBF1*. $(\forall_D\alpha\alpha)(\alpha\alpha \text{ are well-ordered by } <)$

PBF2*. $(\forall\alpha)(\alpha = \text{Ord}(\alpha\alpha, <)).$

If all the ordinals were a permissible value of $\alpha\alpha$ in PBF1*, we could derive the problematic $\text{Ord}(oo, <) = \Omega$ by instantiation. But $\alpha\alpha$ is restricted and so all the ordinals are not a permissible value of $\alpha\alpha$ in PBF1*. We cannot say that all the ordinals are well-ordered. But it would be a mistake to infer that all the ordinals are not well-ordered. Rather, the restrictivist solution consists in the doctrine that we simply cannot speak of all the ordinals at all.

BIBLIOGRAPHY

- Balaguer, M. (1998). *Platonism and Anti-Platonism in Mathematics*. Oxford University Press.
- Beaney, M. (1997). *The Frege Reader*. Blackwell.
- Benacerraf, P. (1965). What Numbers Could Not Be. In P. Benacerraf and H. Putnam (Eds.), *Philosophy of Mathematics: Selected Readings*, pp. 272–294. Cambridge University Press (1983).
- Bennet, C. and M. F. Karlsson (2008). Williamson’s Barber. *Analysis* 68(300), 320–326.
- Bernays, P. (1961). Zur Frage der Unendlichkeitsschemata in der axiomatischen Mengenlehre. In Y. Bar-Hillel, E. Poznanski, M. Rabin, and A. Robinson (Eds.), *Essays on the Foundations of Mathematics*, pp. 3–49. Jerusalem: Magnes Press.
- Black, M. (1971). The Elusiveness of Sets. *Review of Metaphysics* 24, 614–636.
- Boolos, G. (1984). To Be is to Be a Value of a Variable (or to Be some Values of some Variables). In *Logic, Logic, Logic*, pp. 54–72.
- Boolos, G. (1998). *Logic, Logic, and Logic*. Harvard University Press.
- Burge, T. (1984). Frege on Extensions of Concepts, From 1884 to 1903. In *Truth, Thought, Reason: Essays on Frege*, pp. 273–298. Oxford University Press (2005).
- Burgess, J. (2008). *Mathematics, Models, and Modality: Selected Philosophical Essays*. Cambridge University Press.
- Burgess, J. and G. Rosen (1997). *A Subject with No Object*. Oxford: Clarendon Press.
- Button, T. (2010). Dadaism: Restrictivism as Militant Quietism. *Proceedings of the Aristotelian Society* 110(3), 387–398.
- Cantor, G. (1883). Grundlagen einer allgemeinen Mannigfaltigkeitslehre. In E. Zermelo (Ed.), *Gesammelte Abhandlungen*, pp. 165–208. Springer (1932).
- Cantor, G. (1885). Review of Frege 1884. *Deutsche Literaturzeitung* VI(20), 728–729.
- Cantor, G. (1895). Beiträge zur Begründung der transfiniten Mengenlehre I. *Mathematische Annalen* 46, 481–512.

- Cantor, G. (1899). Letter to Dedekind. In J. van Heijenoort (Ed.), *From Frege to Gödel*, pp. 113–117. Harvard University Press (1967).
- Carnap, R. (1950). Empiricism, Semantics, and Ontology. *Revue Internationale de Philosophie* 4(11), 20–40.
- Cartwright, R. (1994). Speaking of Everything. *Noûs* 28(1), 1–20.
- Cocchiarella, N. (1987). *Logical Studies in Early Analytic Philosophy*. Ohio State University Press.
- Coffa, J. A. (1979). The Humble Origins of Russell’s Paradox. *Russell* 33, 31–37.
- Copi, I. M. (1958). The Burali-Forti Paradox. *Philosophy of Science* 25(4), 281–286.
- Dieveney, P. (2013). Anything and Everything. *Erkenntnis* 78(1), 119–140.
- Dummett, M. (1963). The Philosophical Significance of Gödel’s Theorem. In M. Dummett (Ed.), *Truth and Other Enigmas (1978)*, pp. 186–214. Duckworth.
- Dummett, M. (1981). *Frege: Philosophy of Language* (Second ed.). Duckworth.
- Dummett, M. (1991). *Frege: Philosophy of Mathematics*. Harvard University Press.
- Ebert, P. and M. Rossberg (2009). Cantor on Frege’s Foundations of Arithmetic: Cantor’s 1885 Review of Frege’s *Die Grundlagen der Arithmetik*. *History and Philosophy of Logic* 30(4), 341–348.
- Eklund, M. (2005). Fiction, Indifference, and Ontology. *Philosophy and Phenomenological Research* 71(3), 557–579.
- Enderton, H. (1977). *Elements of Set Theory*. Academic Press.
- Enderton, H. (2015). Second-order and Higher-order Logic. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2015 ed.).
- Ewald, W. B. (2005). *From Kant to Hilbert Volume 2*. Oxford University Press.
- Fine, K. (1981). First-order modal theories i-sets. *Noûs* 15, 177–205.
- Fine, K. (1994). Essence and Modality. *Philosophical Perspectives* 8, 1–16.
- Fine, K. (1995a). The Logic of Essence. *Journal of Philosophical Logic* 24(3), 241–273.
- Fine, K. (1995b). Ontological Dependence. *Proceedings of the Aristotelian Society* 95, 269–290.
- Fine, K. (2006a). Our Knowledge of Mathematical Objects. In T. Z. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, Volume 1, pp. 89–109. Clarendon Press.

- Fine, K. (2006b). Relatively Unrestricted Quantification. In A. Rayo and G. Uzquiano (Eds.), *Absolute Generality*, pp. 20–44. Oxford University Press.
- Fraenkel, A. (1922). Zu den Grundlagen der Cantor-Zermeloschen Mengenlehre. *Mathematische Annalen* 86, 230–237.
- Fraenkel, A. A., Y. Bar-Hillel, and A. Levy (1973). *Foundations of Set Theory*. North Holland.
- Frege, G. (1884). *The Foundations of Arithmetic*. Northwestern University Press (Second ed., 1980).
- Frege, G. (1891). Function and Concept. In M. Beaney (Ed.), *The Frege Reader*, pp. 130–148. Blackwell (1997).
- Frege, G. (1892a). Comments on Sinn and Bedeutung. In M. Beaney (Ed.), *The Frege Reader*, pp. 172–180. Blackwell (1997).
- Frege, G. (1892b). Concept and Object. In M. Beaney (Ed.), *The Frege Reader*, pp. 181–193. Blackwell (1997).
- Frege, G. (1892c). On Sinn and Bedeutung. In M. Beaney (Ed.), *The Frege Reader*, pp. 151–171. Blackwell (1997).
- Frege, G. (1893). Grundgesetze der Arithmetik, Volume I. In M. Beaney (Ed.), *The Frege Reader*, pp. 194–223. Blackwell (1997).
- Frege, G. (1895). A Critical Elucidation of some Points in E. Schröder’s “Lectures on the Algebra of Logic”. In B. McGuinness (Ed.), *Collected Papers on Mathematics, Logic and Philosophy*, pp. 210–228. Blackwell (1984).
- Frege, G. (1903). Grundgesetze der Arithmetik, Volume II. In M. Beaney (Ed.), *The Frege Reader*, pp. 258–289. Blackwell (1997).
- Frege, G. (1906). On Shoenflies: “Die logischen Paradoxien der Mengenlehre”. In *Postumous Writings*, pp. 176–183. University of Chicago (1979).
- Fritz, P., H. Lederman, T. Liu, and D. Scott (n.d.). Can Modalities Save Naive Set Theory? *Review of Symbolic Logic*, to appear.
- Geach, P. (1962). *Reference and Generality*. Cornell University Press.
- Gillies, D. (2011). *Frege, Dedekind, and Peano on the Foundations of Arithmetic*. Routledge.
- Glanzberg, M. (2004). Quantification and Realism. *Philosophy and Phenomenological Research* 69(3), 541–572.
- Glanzberg, M. (2006). Context and Unrestricted Quantification. In A. Rayo and G. Uzquiano (Eds.), *Absolute Generality*, pp. 45–74. Oxford University Press.

- Gödel, K. (1947). What is Cantor's Continuum Problem? *American Mathematical Monthly* 54(9), 515–525.
- Hale, B. (1988). *Abstract Objects*. Basil Blackwell.
- Hale, B. and C. Wright (2001). *The Reason's Proper Study: Essays Towards a Neo-Fregean Philosophy of Mathematics*. Oxford University Press.
- Hallett, M. (1984). *Cantorian Set Theory and Limitation of Size*. Oxford University Press.
- Hazen, A. (1997). Relations in Lewis's Framework without Atoms. *Analysis* 57(4), 243–248.
- Hellman, G. (1989). *Mathematics Without Numbers: Towards a Modal-Structural Interpretation*. Oxford University Press.
- Hellman, G. (2011). On the Significance of the Burali-Forti Paradox. *Analysis* 71(4), 631–637.
- Hewitt, S. (2015). When Do Some Things Form a Set? *Philosophia Mathematica* 23(3), 311–337.
- Hirsch, E. (1982). *The Concept of Identity*. Oxford University Press.
- Incurvati, L. (2012). How to be a Minimalist about Sets. *Philosophical Studies* 159, 69–87.
- Incurvati, L. (n.d.). Maximality Principles in Set Theory. *Philosophia Mathematica*, to appear.
- Jech, T. (2003). *Set Theory*. Springer.
- Jourdain, P. (1905). On Transfinite Cardinal Number of the Exponential Form. *Philosophical Magazine* 9(6), 42–56.
- Keränen, J. (2006). The Identity Problem for Realist Structuralism II: A Reply to Shapiro. In F. MacBride (Ed.), *Identity and Modality*, pp. 146–163. Oxford.
- Klement, K. (2010). Russell, His Paradoxes, and Cantor's Theorem: Part I. *Philosophy Compass* 5(1), 16–28.
- Klement, K. (2012). Frege's Changing Conception of Number. *Theoria* 78(2), 146–167.
- Lavine, S. (1994). *Understanding the Infinite*. Harvard University Press.
- Lavine, S. (2006). Something about Everything: Universal Quantification in the Universal Sense of Universal Quantification. In A. Rayo and G. Uzquiano (Eds.), *Absolute Generality*, pp. 98–148. Oxford University Press.

- Lewis, D. (1983). New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61(December), 343–377.
- Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell Publishers.
- Lewis, D. (1991). *Parts of Classes*. Blackwell.
- Lewis, D. (1993). Many, but Almost One. In *Papers in Metaphysics and Epistemology*, pp. 164–182. Cambridge University Press 1999.
- Linnebo, O. (2006). Sets, Properties and Unrestricted Quantification. In A. Rayo and G. Uzquiano (Eds.), *Absolute Generality*, pp. 45–97. Oxford University Press.
- Linnebo, O. (2009). The Individuation of the Natural Numbers. In O. Bueno and O. Linnebo (Eds.), *New Waves in Philosophy of Mathematics*, pp. 220–238. Palgrave.
- Linnebo, O. (2010). Pluralities and Sets. *Journal of Philosophy* 107(3), 144–164.
- Linnebo, O. (2013). The Potential Hierarchy of Sets. *Review of Symbolic Logic* 6(2), 205–228.
- Linnebo, O. (2014). Plural Quantification. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2014 ed.).
- MacBride, F. (2005). Structuralism Reconsidered. In S. Shapiro (Ed.), *The Oxford Handbook of Philosophy of Mathematics and Logic*, pp. 563–589. Oxford University Press.
- Maddy, P. (1997). *Naturalism in Mathematics*. Oxford University Press.
- Martin, D. A. (2001). Multiple Universes of Sets and Indeterminate Truth Values. *Topoi* 20(1), 5–16.
- McGee, V. (2000). ‘Everything’. In G. Sher and R. L. Tieszen (Eds.), *Between Logic and Intuition: Essays in Honor of Charles Parsons*, pp. 54–78. Cambridge University Press.
- McKay, T. (2006). *Plural Predication*. Oxford: Clarendon Press.
- Mendelson, E. (1997). *Introduction to Mathematical Logic* (Fourth ed.). Chapman and Hall.
- Menzel, C. (2014). Wide Sets, ZFCU, and the Iterative Conception. *Journal of Philosophy* 111(2), 57–83.
- Montague, R. (1957). *Contributions to the Axiomatic Foundation of Set Theory*. University of California, Berkeley: PhD thesis.

- Montague, R. (1961). Fraenkel's Addition to the Axioms of Zermelo. In Y. Bar-Hillel, E. Poznanski, M. Rabin, and A. Robinson (Eds.), *Essays on the Foundations of Mathematics*, pp. 91–114. Jerusalem: Magnes Press.
- Moore, G. and A. Garciadiego (1981). Burali Forti's Paradox: A Reappraisal of its Origins. *Historia Mathematica* 8, 319–50.
- Oliver, A. and T. Smiley (2006). What Are Sets and What Are They For? *Philosophical Perspectives* 20(1), 123–155.
- Oliver, A. and T. Smiley (2013). *Plural Logic*. Oxford University Press.
- Parsons, C. (1977). What is the Iterative Conception of Set? In *Mathematics in Philosophy: Selected Essays*, pp. 268–297. Cornell University Press (1983).
- Parsons, C. (1983). *Mathematics in Philosophy: Selected Essays*. Cornell University Press.
- Parsons, C. (1997). Structuralism and the Concept of Set. In E. Agazzi and G. Darvas (Eds.), *Philosophy of Mathematics Today*, pp. 171–194. Kluwer Academic Publishers.
- Parsons, T. (1986). Why Frege Should Not Have Said “The Concept Horse is Not a Concept”. *History of Philosophy Quarterly* 3(4), 449–465.
- Paseau, A. (2007). Boolos on the Justification of Set Theory. *Philosophia Mathematica* 15(1), 30–53.
- Potter, M. (1993). Iterative Set Theory. *Philosophical Quarterly* 44(171), 178–193.
- Potter, M. (2004). *Set Theory and its Philosophy: A Critical Introduction*. Oxford University Press.
- Priest, G. (1994). The Structure of the Paradoxes of Self-Reference. *Mind* 103(409), 25–34.
- Priest, G. (2002). *Beyond the Limits of Thought*. Oxford University Press.
- Priest, G. (2013). Indefinite Extensibility—Dialetheic Style. *Studia Logica* 101(6), 1263–1275.
- Putnam, H. (1977). Realism and Reason. *Proceedings and Addresses of the American Philosophical Association* 50(6), 483–498.
- Putnam, H. (1979). *Mathematics, Matter and Method* (Second ed.). Cambridge University Press.
- Putnam, H. (1980). Models and Reality. *Journal of Symbolic Logic* 45(3), 464–482.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge University Press.

- Putnam, H. (1987). *The Many Faces of Realism*. Open Court.
- Putnam, H. (2000). Paradox Revisited II: Sets—A Case of All or None. In G. Sher and R. Tieszen (Eds.), *Between Logic and Intuition: Essays in Honor of Charles Parsons*, pp. 16–26. Cambridge University Press.
- Quine, W. V. (1937). New Foundations for Mathematical Logic. *The American Mathematical Monthly* 44(2), 70–80.
- Quine, W. V. (1960). *Word and Object*. MIT Press.
- Quine, W. V. (1969). *Ontological Relativity and Other Essays*. Columbia University Press.
- Quine, W. V. (1986). *Philosophy of Logic*. Harvard University Press.
- Rayo, A. (2002). Word and Objects. *Noûs* 36, 436–464.
- Rayo, A. (2006). Beyond Plurals. In A. Rayo and G. Uzquiano (Eds.), *Absolute Generality*, pp. 220–254. Oxford University Press.
- Rayo, A. (2007). Ontological Commitment. *Philosophy Compass* 2(3), 428–444.
- Resnik, M. (1980). *Frege and the Philosophy of Mathematics*. Cornell University Press.
- Resnik, M. (1997). *Mathematics as a Science of Patterns*. Oxford University Press.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In Bob Hale and Aviv Hoffmann (Eds.), *Modality: Metaphysics, Logic, and Epistemology*, pp. 109–36. Oxford University Press.
- Ruffino, M. (2003). Frege’s View on Vagueness. *Manuscrito* 26(2), 253–277.
- Russell, B. (1903). *The Principles of Mathematics*. Cambridge University Press.
- Russell, B. (1906). On Some Difficulties in the Theory of Transfinite Numbers and Order Types. In D. Lackey (Ed.), *Essays in Analysis*, pp. 135–164. George Braziller (1973).
- Russell, B. (1908). Mathematical Logic as Based on the Theory of Types. In R. C. Marsh (Ed.), *Logic and Knowledge*, pp. 59–102. Macmillan (1964).
- Russell, B. (1919). *Introduction to Mathematical Philosophy*. Dover Publications.
- Russell, B. (1959). *My Philosophical Development*. George Allen and Unwin.
- Schaffer, J. (2009). On What Grounds What. In D. Manley, D. J. Chalmers, and R. Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 347–383. Oxford University Press.

- Scott, D. (1974). Axiomatizing Set Theory. In T. Jech (Ed.), *Proceedings of Symposia in Pure Mathematics*, Volume 13 (Part 2), pp. 207–214. American Mathematical Society.
- Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press.
- Shapiro, S. (2003). All Sets Great and Small: And I Do Mean ALL. *Philosophical Perspectives* 17(1), 467–490.
- Shapiro, S. (2006). Structure and Identity. In F. MacBride (Ed.), *Identity and Modality*, pp. 109–145. Oxford.
- Shapiro, S. and C. Wright (2006). All Things Indefinitely Extensible. In A. Rayo and G. Uzquiano (Eds.), *Relatively Unrestricted Quantification*, pp. 255–304. Oxford University Press.
- Shoenfield, J. (1967). *Mathematical Logic*. Addison-Wesley Pub. Co.
- Skolem, T. (1922). Some Remarks on Axiomatized Set Theory. In J. van Heijenoort (Ed.), *From Frege to Gödel*, pp. 290–301. Harvard University Press (1967).
- Spencer, J. (2012). All Things Must Pass Away. *Oxford Studies in Metaphysics* 7, 67–92.
- Stoll, R. and H. Enderton (2016). Set Theory. In *Encyclopdia Britannica Online*.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Routledge.
- Studd, J. P. (2013). The Iterative Conception of Set: A (Bi-)Modal Axiomatisation. *Journal of Philosophical Logic* 42(5), 697–725.
- Tahko, T. and J. Lowe (2015). Ontological Dependence. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2015 ed.).
- Tait, W. (1997). Frege Versus Cantor and Dedekind: On the Concept of Number. In W. Tait (Ed.), *Early Analytic Philosophy, Frege, Russell, Wittgenstein: Essays in Honor of Leonard Linsky*, pp. 213–248. Open Court.
- Tait, W. (2005). *The Provenance of Pure Reason: Essays in the Philosophy of Mathematics and its History*. Oxford University Press.
- Uzquiano, G. (2003). Plural Quantification and Classes. *Philosophia Mathematica* 11(1), 67–81.
- Uzquiano, G. (2015). Varieties of Indefinite Extensibility. *Notre Dame Journal of Formal Logic* 56(1), 147–166.
- van Heijenoort, J. (1986). Frege on Vagueness. In L. Haaparanta and J. Hintikka (Eds.), *Frege Synthesized*, pp. 31–45. Reidel Publishing.

- von Neumann, J. (1925). An Axiomatization of Set Theory. In J. van Heijenoort (Ed.), *From Frege to Gödel*, pp. 393–413. Harvard University Press 1967.
- Wang, H. (1974). The Concept of Set. In P. Benacerraf and H. Putnam (Eds.), *Philosophy of Mathematics: Selected Readings* (2nd ed.), pp. 530–570. Cambridge University Press (1983).
- Welch, P. and L. Horsten (2016). Reflecting on Absolute Infinity. *Journal of Philosophy* 113(2), 89–111.
- Weyl, H. (1949). *Philosophy of Mathematics and Natural Science*. Princeton University Press.
- Wiggins, D. (1980). *Sameness and Substance*. Harvard University Press.
- Williamson, T. (1998). Indefinite Extensibility. *Grazer Philosophische Studien* 55, 1–24.
- Williamson, T. (2003). Everything. *Philosophical Perspectives* 17(1), 415–465.
- Wittgenstein, L. (2001). *Tractatus Logico-Philosophicus*. Routledge.
- Wright, C. (1983). *Frege's Conception of Numbers as Objects*. Aberdeen University Press.
- Xu, F. (1997). From Lot's Wife to a Pillar of Salt: Evidence That Physical Object is a Sortal Concept. *Mind and Language* 12(3-4), 365–392.
- Yablo, S. (2006). Circularity and Paradox. In T. Bolander, V. F. Hendricks, and S. A. Pedersen (Eds.), *Self-Reference*, pp. 165–183. CSLI Publications.
- Yi, B. (2005). The Logic and Meaning of Plurals. Part I. *Journal of Philosophical Logic* 34(5-6), 459–506.
- Yi, B. (2006). The Logic and Meaning of Plurals. Part II. *Journal of Philosophical Logic* 35, 239–288.
- Zermelo, E. (1908). Investigations in the Foundations of Set Theory I. In J. van Heijenoort (Ed.), *From Frege to Gödel*, pp. 199–215. Harvard University Press 1967.
- Zermelo, E. (1930). Über Grenzzahlen und Mengenbereiche: Neue Untersuchungen über die Grundlagen der Mengenlehre. *Fundamenta Mathematicae* 16, 29–47.