Notre Dame Journal of Formal Logic
Volume 41, Number 4, 2000

ABSTRACTION AND SET THEORY

BOB HALE

Abstract The neo-Fregean program in the philosophy of mathematics seeks a foundation for a substantial part of mathematics in abstraction principles—for example, Hume's Principle: The number of Fs = the number of Gs iff the Fs and Gs correspond one-one—which can be regarded as implicitly definitional of fundamental mathematical concepts—for example, cardinal number. This paper considers what kind of abstraction principle might serve as the basis for a neo-Fregean set theory. Following a brief review of the main difficulties confronting the most widely discussed proposal to date—replacing Frege's inconsistent Basic Law V by Boolos's New V which restricts concepts whose extensions obey the principle of extensionality to those which are small in the sense of being smaller than the universe—the paper canvasses an alternative way of implementing the limitation of size idea and explores the kind of restrictions which would be required for it to avoid collapse.

1. Preliminaries

The neo-Fregean program in the philosophy of mathematics aims to provide a foundation for a substantial part of mathematics in abstraction principles which can be regarded as implicitly definitional of fundamental mathematical concepts. By abstraction principles we mean, roughly,[1] principles of the shape,

$$\forall\alpha\forall\beta(\Sigma(\alpha) = \Sigma(\beta) \leftrightarrow \alpha \approx \beta)$$

where $\approx$ is an equivalence relation on entities of the type of $\alpha, \beta, \ldots$, and Σ is a function from entities of that type to objects. Prominent examples are the

Direction Equivalence: The direction of line a = the direction of line b $\leftrightarrow$ a and b are parallel,

in terms of which much of Frege's original discussion of such principles is conducted;[2]

Hume's Principle: $\forall F\forall G[\mathrm{N}x : Fx = \mathrm{N}x : Gx \leftrightarrow \exists R(F1_{\_R\_}G)]$,

Received March 29, 2001; printed October 28, 2002
2001 Mathematics Subject Classification: Primary, 03A05; Secondary, 00A30
Keywords: abstraction principle, set, limitation of size, sortal concept, definiteness

which he considers, but eventually rejects, as a means of defining (cardinal) *number*; and, of course,

Basic Law V: $\forall F \forall G[\{x : Fx\} = \{x : Gx\} \leftrightarrow \forall x(Fx \leftrightarrow Gx)]$

—the set or class form of Frege's ill-fated axiom on value-ranges. As is well known, in the neo-Fregean view, Hume's Principle may—Frege's and other misgivings notwithstanding—be taken as a means of implicitly defining the concept of number and can serve as a foundational principle for, at least, elementary arithmetic. In its philosophical aspect, this claim is of course controversial. But it is not my purpose to engage further in that controversy here. Instead I would like to explore—rather tentatively, I should stress—the prospects for developing a version of set theory along abstractionist lines. Can we find an abstraction principle, or principles, which might serve as the foundation for an interesting theory of sets?

An abstraction principle which can plausibly be seen as implicitly defining the concept of *set* will do so by fixing the identity conditions for its instances, and will—on pain of changing the subject—take the identity of sets to consist in their having the same members. If we assume—what does not seem seriously disputable—that any plausible candidate will be a higher-order abstraction (i.e., will involve abstraction over an equivalence relation among concepts, rather than objects), then it seems clear, further, that the equivalence relation involved will have to be either coextensiveness of concepts or a close relative of it. In other words, what we are looking for is, broadly speaking, a consistency-preserving restriction of Basic Law V. I think we can also assume that a suitable restriction—if one can be found—will, in effect (and in contrast with the restriction Frege himself tried), be a restriction on what concepts (can) have sets corresponding to them.

If we schematically represent the sought after restriction using a second-level predicate 'Good', then the most obvious ways to restrict BLV are

(A) $\forall F \forall G[\mathrm{Good}(F) \vee \mathrm{Good}(G) \rightarrow (\{x \mid Fx\} = \{x \mid Gx\} \leftrightarrow \forall x(Fx \leftrightarrow Gx))]$

and

(B) $\forall F \forall G[\{x \mid Fx\} = \{x \mid Gx\} \leftrightarrow (\mathrm{Good}(F) \vee \mathrm{Good}(G) \rightarrow \forall x(Fx \leftrightarrow Gx))]$.

The main difference between these is that (A) is a conditionalized abstraction principle, whereas (B) is unconditional, with the restriction built into the relation required to hold between F and G for them to yield the same set—fairly obviously, the resulting relation is an equivalence relation.[3] Consequently, (B) yields a "set" for every F, regardless of whether it is Good or not, whereas (A) yields a set $\{x \mid Fx\}$ only if we have the additional premise that F is Good. If neither F nor G is Good, the right-hand side of (B) holds vacuously, so we get that $\{x \mid Fx\} = \{x \mid Gx\}$—regardless of whether F and G are coextensive. That is, we get the same "set" from all Bad concepts. We get real sets via (B)—that is, objects whose identity is determined by their membership—only from Good concepts.

2. Goodness as Smallness (1)—New V

What is it for a concept to be Good? Various suggestions have been canvassed. One general approach picks up on the well-entrenched "limitation of size" idea, that the set-theoretic paradoxes stem from treating as sets "collections" which are in some sense "too big"—the collection of all sets, of all sets that are not members of themselves, of all ordinals, and so forth. On one version of this (Goodness is Smallness) approach, we define a concept to be Small if it is smaller (i.e., has fewer instances) than some concept under which everything, or at least every object, falls. And the favored universal concept has been that of self-identity. Following Boolos ([3], p. 178), we say that a concept G "goes into" a concept F if and only if there is a one-one function taking the Gs into the Fs, and that F is Small if and only if the concept self-identical does not go into F. If we frame our restricted set-abstraction in the style (B), the result is what Boolos called New V (and Wright ([12], p. 300) calls VE):

New V: $$\forall F \forall G[\{x \mid Fx\} = \{x \mid Gx\} \leftrightarrow (\text{Small}(F) \vee \text{Small}(G) \rightarrow \forall x(Fx \leftrightarrow Gx))].$$

Although, as is now well known, a certain amount of set-theory can be obtained by adding New V to second-order logic, there are some problems with it. The most serious of these is that we don't get *enough* set theory. As Boolos showed, neither an axiom of infinity nor the power set axiom can be obtained as theorems on this basis, so the theory is rather weak—and certainly weaker than a neo-Fregean requires, if he is to have a set-theoretic foundation for analysis.[4] I shall return to this below.

A further, quite different, difficulty relates to the constraints needed to differentiate between good or acceptable abstraction principles and bad or unacceptable ones. It is clear that some constraints are needed since not all abstractions are acceptable—as is dramatically illustrated by Basic Law V. Obviously consistency is one requirement. But it does not seem that it can be the only one since—as Boolos also showed—one can formulate abstraction principles which are severally consistent but mutually incompatible. For example,[5] we can take as our equivalence relation the relation which holds between concepts F and G just when their symmetric difference is finite (i.e., when there are just finitely many objects which are either F-but-not-G or G-but-not-F). Writing this briefly as $\Delta(F, G)$, we can frame the abstraction which Wright calls

Nuisances: $$\forall F \forall G[\nu(F) = \nu(G) \leftrightarrow \Delta(F, G)].$$

As Wright shows, this is a consistent abstraction but is satisfiable only in domains containing finitely many objects. Hume's Principle by contrast, though likewise provably consistent, is satisfiable only in domains containing at least a countable infinity of objects. Since they are thus mutually incompatible, Hume's Principle and Nuisances cannot both be acceptable. The problem for the neo-Fregean is to justify rejecting the latter as unacceptable. To this end, Wright proposes a constraint—his *first conservativeness constraint*[6]—which has obvious affinities to the notion of conservativeness deployed by Field in his defense of nominalism.[7] His plausible thought is, roughly, that a satisfactory explanation of a concept—whether by means of an abstraction principle or other form of definition—should do no more than fix the truth-*conditions* of statements involving that concept. It should have nothing to say about the truth-*values* of statements which already have determinate truth-conditions independently of the introduction of that concept, and in particular, it should carry no implications for the extensions of other concepts unconnected with the concept the explanation seeks to introduce. If we think of an abstraction principle as added to an existing theory, the requirement can be expressed, still somewhat roughly, as that the abstraction should carry no implications regarding the "old" ontology—the ontology of the given theory; it should be conservative with respect to that theory in the sense that its addition to the theory does not settle the truth-values of any statements expressible in the old language which are left unsettled by that theory.[8]

Precisely because it is satisfiable only in finite universes, Nuisances *does* carry such implications—it implies, for example, that there are at most finitely many aardvarks, or subatomic particles, or space-time points, and so forth—it fails this constraint and should therefore be rejected as unacceptable, even though consistent. Hume's Principle, by constrast, while implying that there are at least countably infinitely many objects, places no restrictions—either upper or lower bounds—upon the extensions of concepts other than the concept it is intended implicitly to define. That it places no upper bound is obvious. But it is crucial that it places no lower bound either, since the requirement that there be at least infinitely many objects is satisfied by the abstracts—numbers—which it itself serves to introduce, so that it makes no demand that any other concepts should have infinitely many instances.

So far, so good. But how do things stand with regard to New V? Shapiro and Weir [11], exploiting a point made originally by Boolos ([2], p. 102), have argued that New V violates Wright's own first constraint. In essence, the argument is simple enough. If the concept *ordinal* is Small, New V yields a set (and not just a "set") of all the ordinals and we have the Burali-Forti contradiction. Hence *ordinal* must be Big. But in that case it is exactly as big as the universe, that is, there is a one-one correspondence between *ordinal* and a(ny) universal concept, say self-identity. But this, together with the fact that the ordinals are well-ordered by membership, entails Global Well-Ordering—the existence of a well-ordering of the universe. Given that the existence, or otherwise, of such a well-ordering may reasonably be taken to be independent of existing theory, New V must be reckoned nonconservative.

3. Goodness as Smallness (2)—Small^2 V

Since the difficulty for New V just explained crucially exploits the definition of Small as, in effect, smaller than the (or a) universal concept, it is possible that it could be avoided by a suitable redefinition of Smallness. As a first step in the direction I have in mind, one might, in preference to defining Smallness in terms of being smaller than some specified universal concept, say that a concept is Small if it is *smaller than some concept or other*—where $F < G$ if there is a bijection of F into G but not-($F \sim G$)—and take as our set-abstraction either New V with Small so defined, or perhaps a conditional ((A)-type) abstraction:

Small V: $$\forall F \forall G[\text{Small}(F) \vee \text{Small}(G) \rightarrow (\{x \mid Fx\} = \{x \mid Gx\} \leftrightarrow \forall x(Fx \leftrightarrow Gx))].$$

However, it is obvious that, at least so long as some universal concept V is in play and can serve as providing an upper bound, so to speak, on the potential sizes of concepts, this simple suggestion makes no advance over New V as originally understood.[9] For then, since any concept can be no bigger than our universal concept V, a concept F will be smaller than some concept G only if smaller than V, and if smaller than V, will certainly be smaller than some concept or other—so $\exists G\ F < G$ if and only if $F <$ V. But now if *ordinal* is Small in this sense, and so Good, we shall have the Burali-Forti again, so that *ordinal* must be Bad, that is, *ordinal* $\sim$V, and we have Global Well-Ordering, just as before.

Flawed as our simple proposal is, there is a refinement of it which really does avoid the Global Well-Ordering problem.[10] This is to interpret Goodness as double Smallness, where a concept is doubly small if and only if it is (strictly) smaller than some concept which is itself (strictly) smaller than some concept, that is,

$$\mathrm{Small}^2(F) \leftrightarrow \exists G \exists H(F < G < H).$$

Interpreting Goodness as Smallness^2 blocks the reasoning which shows that New V, as originally understood with Small as meaning "smaller than the universe," implies Global Well-Ordering. We can still show, of course, that *ordinal* cannot be Good—that is, now Small^2—since if it were, we would have the Burali-Forti paradox, just as before. So we have to agree that *ordinal* is Bad. But that just means that it is not Small^2, and from this we cannot infer that it is bijectible onto any universal concept (or indeed onto any concept).[11]

As has already been indicated, Wright's first conservativeness constraint is one of a pair. The *second* constraint he proposes (see [12], pp. 300–302) concerns abstraction principles which, as he puts it, embed a paradoxical component—centrally, abstractions of the type,

(D) $\quad \forall F \forall G[\Sigma(F) = \Sigma(G) \leftrightarrow ((\varphi(F) \wedge \varphi(G)) \vee \forall x(Fx \leftrightarrow Gx)))],$

of which New V is an instance.[12] In general, by exploiting the reasoning that leads from BLV to contradiction, we can prove, from any instance of (D), that $\exists F \varphi(F)$. For example, from New V, one can prove, via the Russell contradiction, that there are Bad concepts (i.e., that not all concepts are Small). In particular, we can prove that *self-identity* is Big. But as Wright observes, that is a result which we can prove *independently* of New V, as a theorem of second-order logic. The second constraint proposes that this last condition should be met by any (D)-type abstraction or, perhaps more generally, any abstraction which embeds a paradoxical component. As Wright at one point expresses it: "any consequences which may be elicited [from the abstraction] by exploiting its paradoxical component should be, a priori, in independent good standing."[13] The precise force of this constraint depends, obviously, on what is to be understood by consequences being in independent good standing. Being independently provable in logic alone would clearly suffice, but Wright does not wish to accept that as a necessary condition: ". . . 'independent good standing' might also reasonably be taken to cover the case where a consequence elicited from such an abstraction by 'fishy'—paradox-exploitative—means can also be obtained not from logic alone but, as it were, innocently from additional resources provided by that very abstraction" ([12], p. 303). Given this qualification, it is at least not clear that the derivability of Global Well-Ordering via the Burali-Forti constitutes a violation of conservativeness in Wright's second sense, as distinct from his first. And for essentially the same reason, the fact that Small^2 V, although not implying Global Well-Ordering, does imply that there can be no concept larger than ordinal but smaller than the universe is not clearly in breach of the second constraint either.[14] But, at least pending further clarification of the key notion of independent good standing—and especially of what it is for a result to be establishable only in a viciously paradox-exploitative way—it is anyway not clear what form, if any, of the second constraint should be respected. Certainly "paradox-exploitative" had better not be understood so liberally as to render any proof by reductio ad absurdum as such. No acceptable constraint—and certainly none that Wright intended—should require that we may accept a result established by reductio ad absurdum only when we can independently prove it by constructive means.

The other, and almost certainly the most serious, of the two problems New V faces, as we noted, is that it suffices for only a rather weak set theory. In particular, it doesn't give us either an axiom of infinity or a power set axiom as theorems. And the same goes, of course, for New V reinterpreted with Good as Small^2 and Small^2 V. As Wright has observed, however, this need not be a crippling drawback from the neo-Fregean's point of view, if he can justify supplementing New V, or Small^2 V, with other principles—perhaps other abstraction principles—which compensate for its weakness. On this more catholic approach, we separate two distinct roles one might ask a set-abstraction principle to discharge—*fixing the concept* of set, on the one hand and, on the other, serving as a *comprehension principle*. The claim would be that New V's—or Small^2 V's—shortcomings as a comprehension principle need not debar it from successfully discharging a concept-fixing role—of serving as a means of introducing the concept, while leaving its extension to be determined, largely or even entirely, by other principles.[15] I want to discuss a couple of ways in which this might be done. I'll concentrate on the possibility of supplementing Small^2 V with other *abstraction* principles—I don't think it is obvious that the neo-Fregean could not justify using supplementary principles other than abstractions, but I won't pursue that alternative here.[16]

The general strategy then is to look for other abstraction principles which might be used to set up sortal concepts in the presence of which we get an interesting range of Small^2 sortal concepts which have sets corresponding to them.

4. Cut Principles

The first approach I want to consider makes essential use of a kind of abstraction principle which plays a key role in a neo-Fregean construction of the real numbers I developed a little while ago (in Hale [8]). In outline, the leading idea was to get the real numbers in broadly the way that Frege proposed to do, by defining them as *ratios of quantities*. The concept of a ratio of quantities is itself to be introduced by means of an abstraction corresponding to the ancient *equimultiples* principle:

> The ratio $a : b =$ the ratio $c : d$ iff for all positive integers m and n, ma is equal to, greater, or less than nb according as mc is equal to, greater, or less than nd.

Here a, b are quantities of some single kind, and likewise c, d. Crucially, c and d need not be of the same kind as a and b.[17] Quantities are themselves abstract objects (defined by abstraction over quantitative equivalence relations).[18] To get all the positive reals (and, with a little extra work, all the reals), ratio abstraction has to be applied to a sufficiently rich abstract structure—what I called a *complete quantitative domain*. A kind of quantity Q constitutes a complete domain if and only if it is closed under a commutative and associative operation $\oplus$ such that exactly one of $a = b, \exists c(a = b \oplus c), \exists c(b = a \oplus c)$ holds for any "elements" of Q and the following further conditions are met:

[Archimedean condition]	$\forall a, b \in Q \exists m(ma > b)$,
[Fourth proportionals]	$\forall a, b, c, \in Q \exists d \in Q(a : b = d : c)$,
[Completeness]	Every bounded nonempty property P on Q has a least upper bound.

Given a complete domain of quantities, it is hardly surprising that one gets the (positive) reals by ratio abstraction over it. The obvious question is: Can the neo-Fregean justify the assumption that there exists at least one such domain? If attention is restricted to domains of physical quantities (i.e., quantities "belonging" to physical objects), then the answer is almost certainly No. But nothing in the definition of quantity, or that of quantitative domains, precludes recognizing numbers themselves as a kind of quantity. In particular, the positive natural numbers—whose existence the neo-Fregean can justify by appeal to Hume's Principle—form a quantitative domain meeting all but the last two conditions. And the ratios of positive natural numbers, R^{N+}, form a domain meeting all but the last condition—completeness. To demonstrate, a priori, the existence of a complete domain, the neo-Fregean can mimic Dedekind's construction. Define a property P of ratios of positive natural numbers to be a *cut property* just in case P is nonempty, bounded above, downward closed, and having no greatest instance. Then we can abstract over the cut properties on R^{N+} using

Cut abstraction: $\forall F \forall G[\mathrm{Cut}(F) = \mathrm{Cut}(G) \leftrightarrow \forall x(Fx \leftrightarrow Gx)]$ where x varies over just R^{N+} and F, G over just cut properties on R^{N+}.

The Cuts thus obtained can then be shown to form a complete domain.

For present purposes, it is Cut abstraction that is of primary interest to us. On the basis of Hume's Principle, we can define the property of being a natural number, and show that it has no end of instances. Appealing then to Cut abstraction, we can define the property of being a Cut on R^{N+} and show that the property of being a natural number is smaller than this property, and so is Small. If an abstractionist set theory could be based on Small V, this would suffice to give us an infinite set—the set of natural numbers. That is, Small V combined with Hume's Principle and Cut abstraction would give us the effect of an axiom of infinity. However, if we are working with Small^2 V, we need to show that *natural number* is *doubly* small before we can obtain the corresponding set. Can we do so? Consideration of an objection that has been brought against my use of Cut abstraction suggests a way in which we might.[19]

Cut abstraction may be viewed as an instance of a general schema:

Cut schema: $\forall F \forall G[\mathrm{Cut}(F) = \mathrm{Cut}(G) \leftrightarrow \forall x(Fx \leftrightarrow Gx)]$ where x varies over a suitable domain Q and F, G over cut properties on Q.

What counts as an instance of this schema is, of course, unclear until we say what counts as a *suitable* domain. Although the definition of cut property makes sense whatever domain we assume, there only exist cut properties if the domain is suitably structured—it will need to be at least densely ordered. It is certainly implausible to suppose that the particular Cut principle I've used is the only acceptable instance of this schema.

Cut principles, like Hume's Principle and unlike the direction Equivalence, are *second-order* abstractions—they abstract over an equivalence relation on concepts, rather than objects. But in other respects, they differ significantly from Hume's Principle. Let us assume that we are concerned with the application of abstraction principles to domains of definite cardinal size, finite or infinite. With Hume and Cut principles, the underlying domain is a domain of concepts. There will be more of these than there are objects, however concepts are individuated. If there are κ objects, and concepts are individuated purely extensionally, there will be exactly 2^κ concepts. Thus a second-order abstraction can "generate" up to 2^κ abstracts, when the initial domain of objects is κ-sized. Applied to any domain of concepts, BLV generates the maximum collection of abstracts—one for each concept. By contrast, Hume's Principle is quite modest. Because its equivalence relation partitions the concepts into equivalence classes by equinumerosity, and there are just $\kappa+1$ such classes when there are κ objects over which the concepts are defined, Hume generates more abstracts than there are other objects when, but only when, κ is finite. When κ is infinite, Hume generates $\kappa+1$ abstracts from the 2^κ concepts, and $\kappa+1=\kappa$. Hume inflates on finite domains but not on infinite domains. This ensures that Hume has no finite models. But it is stable at infinite cardinalities. Cut principles behave quite differently. As noted, if a cut principle is applied to a finite domain, it generates no abstracts at all, as there are no cut properties on the domain. If the underlying domain is infinite and at least densely ordered, what happens depends upon whether the domain is *strictly* densely ordered (like the rationals) or *completely* ordered (like the reals). Applied to a strictly dense domain, cut abstraction inflates, giving a completely ordered domain of abstracts. Applied to a completely ordered domain, however, it gives a domain of abstracts isomorphic to the underlying domain, and so does not inflate. That is, what happens with the rationals and reals when cut abstraction is applied and then reapplied is representative of what happens in general. So any one cut principle inflates on a strictly dense domain, but is not rampantly inflationary (in the sense that its iterated application leads to unlimited inflation).

Various people[20] have observed that on certain standard set-theoretic assumptions, any set of definite cardinal size can be put in a (strictly) dense linear order, so that while it is true—as I have claimed—that the reapplication of a cut principle to the (complete and so not strictly dense) domain of abstracts generated by its application to a strictly dense domain does not inflate, there will be a strictly dense ordering of the domain of cut abstracts thus generated on which a cut principle does inflate. So that if we start with a countable strictly dense ordering and apply a cut principle to get a $2^{\aleph_0}$-sized domain C of cuts, there is a strictly dense ordering of C , call it C^*, to which another cut principle may be applied to get a new domain of cuts of size $2^{2^{\aleph_0}}$, and so on. In particular, Cook has proved the following.

Theorem 4.1 (Cook's Theorem) *For any infinite cardinal κ, there is a linear order $(A,<)$ such that $|A|\leq\kappa$ and $|\mathrm{Comp}(A,<)|>\kappa$, where $\mathrm{Comp}(A,<)$ is the set of Dedekind Cuts on $(A,<)$.*[21]

Proof: Given an infinite cardinal κ, let λ be the least cardinal $\leq\kappa$ such that $2^\lambda>\kappa$. Let A be the subset of functions from λ (as an ordinal) into $\{0,1\}$ such that $f\in A$ if and only if there is an ordinal $\gamma<\lambda$ such that for all ordinals $\alpha\geq\gamma$, $f(\alpha)=0$. For $f,g\in A$, let $f<g$ if and only if, at the least γ where $f(\gamma)\neq g(\gamma)$, $f(\gamma)=0$. Then $|A|\leq\kappa$ by the following computation:

$$|A|=|\bigcup_{\gamma<\lambda}^{\gamma}2|\leq\sum_{\gamma<\lambda}2^{|\gamma|}\leq\sum_{\gamma<\lambda}\kappa\leq\lambda\times\kappa=\kappa.$$

But $|\mathrm{Comp}(A,<)|=2^\lambda>\kappa$, since $\mathrm{Comp}(A,<)$ is isomorphic to the set of all functions from λ to $\{0,1\}$. □

Cook thinks his result is disastrous for the neo-Fregean logicist because he thinks that the neo-Fregean should only endorse abstraction principles that are "epistemologically modest," but that certain "natural generalizations" of my cut principle are clearly epistemologically extravagant. Specifically, he argues that the conjunction of Hume's Principle with a generalized cut principle,

GCA $\forall P \forall Q \forall \{H, <\}[\mathrm{Cut}(P, \{H, <\}) = \mathrm{Cut}(Q, \{H, <\}) \leftrightarrow \forall x((Hx \wedge P \text{ and } Q \text{ are cut properties on } \{H, <\}) \rightarrow (Px \leftrightarrow Qx))]$,

has, at best, only proper class-sized models, and that its conjunction with a generalized cut schema

GCA *Schema:* All formulas of the form $\forall P \forall Q[\mathrm{Cut}(P, \{H, <\}) = \mathrm{Cut}(Q, \{H, <\}) \leftrightarrow \forall x(Hx \wedge P \text{ and } Q \text{ are cut properties on } \{H, <\} \rightarrow (Px \leftrightarrow Qx))]$

may have set-sized models, but if so, can have only models of cardinality infinitely many times up from that of the continuum.

I don't have space to discuss this objection in the detail it deserves, so I shall be brief and somewhat dogmatic. I am unmoved by it for two main reasons.

First, Cook seems to me to give no compelling reason why a neo-Fregean abstractionist must endorse either his generalized cut principle or even all instances of his generalized cut schema. GCA is not itself an abstraction principle, and it is not clear why an abstractionist should be committed to it. No doubt there are many instances of the generalization to which the abstractionist should have no objection, but that does not amount to a reason for thinking that he must assert the generalization itself. After all, there are doubtless many instances of the general abstraction schema,

$$\forall \alpha \forall \beta[\Sigma(\alpha) = \Sigma(\beta) \leftrightarrow \alpha \approx \beta],$$

with which the abstractionist should have no quarrel, but he can hardly be expected to endorse *its* generalization,

$$\forall \approx \forall \Sigma \forall \alpha \forall \beta[\Sigma(\alpha) = \Sigma(\beta) \leftrightarrow \alpha \approx \beta],$$

which implies, inter alia, BLV, and is therefore outright inconsistent! I would of course agree that an abstractionist shouldn't reject *any* instance of GCA Schema without good reason, but that is not the same thing as being committed to *all* instances.

Second, Cook's understanding of epistemological modesty seems to me flawed, and indeed, simply question-begging. Cook takes it that an abstraction will be immodest if it "generates too many objects." If "too many" meant "too many to avoid inconsistency," there could be no disagreeing with him. But he doesn't—if he did, he would have an objection only if he'd shown that the generalizations, coupled with HP, lead to contradiction. In fact, it's not clear that he means anything more precise than "rather a lot, by set-theoreticians' standards." The short answer to that is a question: Why should that be *objectionable*? Wouldn't it actually be a rather *good* result, from the neo-Fregean's perspective, if it turned out that his principles are mathematically quite powerful?

Indeed—to return to our main business—it might seem that the neo-Fregean can turn Cook's result to his own advantage. We saw how, by appealing to the fact that the concept *Natural number* is smaller than the concept *Cut on* R^{N+}, he can show that the former concept is at least (singly) Small. But if something sufficiently close to Cook's result were at his disposal, why should he not apply it to the cuts on R^{N+}, together with a suitable linear ordering, to obtain a further cut concept strictly larger than *Cut on* R^{N+}? He would then be in position to apply Small^2 V to obtain the countably infinite set of natural numbers. And if it can be done once, why shouldn't it be done again, and again, to obtain larger and larger uncountably infinite sets?

Before he succumbs to euphoria at the prospect of a quite powerful Small^2 set theory, however, the neo-Fregean should remind himself that he does not get Cook's result for free. We have already noted that its proof relies upon Choice. It is not obvious that Choice must be out of bounds for the neo-Fregean—that he could not argue for it as a logical principle, or secure its effect by means of a suitable abstraction. But it is equally not obvious that he *could* do so. I have not yet been able to get a clear view on the matter, and so must leave this question for further investigation. But there is, in any case, another—glaringly obvious and seemingly more troublesome—fly in the ointment. Cook's proof begins: "Given an infinite cardinal κ, let λ be the least cardinal $\leq \kappa$ such that $2^{\lambda} > \kappa$ . . . ." But how do we know that there is such a λ? What, in other words, justifies the assumption that there *are* cardinals $> \kappa$? I can see no way of justifying it without appealing to the Power Set Axiom and Cantor's Theorem, or something at least as problematic, from the neo-Fregean's point of view. If that is right, then it would seem that Cook's result cannot after all be the blessing in disguise that it may at first appear to be.

The difficulty is not decisive. It might be suggested that a neo-Fregean who looks askance at Cook's proof because it is a proof in set theory is being unduly fastidious.[22] Consider Boolos's proof of the equiconsistency of Frege Arithmetic (HP+2nd-order logic) with second-order arithmetic—this is likewise a proof *in set theory*, but that need not mean that the assurance it provides is unavailable to neo-Fregeans. We need a distinction between the "internal perspective"—which is concerned with what results can be obtained using *only resources available to the neo-Fregean*—and the "external perspective"—which is concerned with what results can be obtained, perhaps by making indispensable use of *other* resources, *about* the neo-Fregean enterprise. Why shouldn't the neo-Fregean welcome Cook's proof *as it stands*, as demonstrating "from the outside" that a neo-Fregean set theory based on Small^2 V + HP + (a suitably restricted) Cut schema is agreeably powerful?

This suggestion raises delicate issues. Their resolution depends, in part at least, upon what principled attitude the neo-Fregean can take toward reasoning that makes essential use of principles which cannot be justified on a neo-Fregean basis, and to what extent he can justify reliance on such reasoning. I don't think we know how much mathematics—and in particular how much set theory—is amenable to neo-Fregean reconstruction. My guess—and I imagine just about everyone's—is that there may be quite severe limits on what can be so reconstructed. That begs the question: What should the neo-Fregean say about the parts that neo-Fregeanism doesn't reach? I think he is bound to regard them as having a significantly different epistemological and ontological status from the reachable parts. But that need not mean that he must dismiss them as worthless. Perhaps he can find an indirect justification for relying on them. If so, then there may be a way to uphold the present suggestion. But even if there is, the difficulty is serious enough to warrant exploration of an alternative strategy.

5. Power Concepts

Since the doubt whether the neo-Fregean can exploit Cook's result turns on the need to appeal to the Power Set Axiom, one might wonder whether one can secure some of the effect of Cantor's Theorem in a higher-order logic without using the Power Set

Axiom—enough to secure a significant range of concepts as Small^2, and so as having sets corresponding to them.

For any first-level concept F, we can form a second-level concept F^P—the concept: *subconcept of* F—defined by $F^P(G) \leftrightarrow \forall x(Gx \rightarrow Fx)$. Define $F \leq G$ if and only if $F \sim H$ for some subconcept H of G, and define $F < G$ if and only if $F \leq G$ but $\neg F \sim G$. Then we can prove, by an obvious adaptation of the usual proof of Cantor's Theorem, that $\forall F\ F < F^P$.

Proof: (i) For each x falling under F, there is a unitary subconcept of F—the concept: $=x$—under which x alone falls. Denote the (second-level) concept under which all these unitary subconcepts fall by F^{Unit}. Obviously $\forall G(F^{\text{Unit}}(G) \rightarrow F^P(G))$. Define the relation S by

$$S(x, G) \leftrightarrow Fx \wedge \forall y(Gy \leftrightarrow y = x).$$

Then x bears S to G if and only if G is that unitary subconcept of F under which x alone falls, and, since S is obviously one-one, we have $F \sim F^{\text{Unit}}$ under R. So $F \leq F^P$.

(ii) Suppose $F \sim F^P$ under some one-one R. Define a subconcept D of F by

$$Dx \leftrightarrow Fx \wedge \forall G(R(x, G) \rightarrow \neg Gx).$$

By the assumption that $F \sim F^P$ under R, we have $R(x, d)$ for some x falling under F. Suppose $R(d, D)$. Suppose Dd. Then by definition of D, we have

$$Fd \wedge \forall G(R(x, G) \rightarrow \neg Gx)$$

whence $R(d, D) \rightarrow \neg Dd$

whence $\neg Dd$.

Suppose then, that $\neg Dd$. Since $R(d, D)$ and R is one-one, $R(d, G)$ if and only if G is D (i.e., $R(d, G) \leftrightarrow \forall x(Gx \leftrightarrow Dx)$), it follows that $\forall G(R(x, G) \rightarrow \neg Gx)$. Hence, again by the definition of D, Dd. So $Dd \leftrightarrow \neg Dd$. Contradiction! Hence $\neg(F \sim F^P)$. $\square$

This is a *restricted* form of Cantor's Theorem, asserting that any first-level concept is strictly smaller than its (second-level) power concept. To state and prove it, we need third-order logic. If we ascend to fourth-order logic, we can prove that any second-level concept is strictly smaller that its (third-level) power concept. And, presumably, so on . . . . The prospect opens up of obtaining each *finite* restriction of Cantor's Theorem—that is, each instance of the schema,

$$\forall \varphi\ \varphi < \varphi^P \text{ for } \varphi \text{ of level } n \text{ and } \varphi^P \text{ of level } n+1$$

in a logic of order ω. Of course, even going up this far doesn't give us anything approaching the full strength of the Power Set Axiom, but it does suggest a method of establishing the Smallness^2 of a significant series of larger and larger concepts by noting that they are doubly smaller than the power concepts of their own power concepts.

Before we turn to what difficulties may stand in the way of this approach, it is worth noticing that the recourse to logic of order ω may be avoidable—we may not need to go above fifth-order logic. Let F be a first-level concept for which we can show, as above, that $F < F^P < F^{P^P}$. Then F is Small^2. Since any subconcept of a Small^2 concept is Small^2, any subconcept G of F is Small^2, and so has a set corresponding to it by Small^2 V. Define F^{P^*} to be the first-level property which an object y has if and only if $y = \{x \mid Gx\}$ for some subconcept G of F. Let R be the relation which holds between x and y if and only if x is G and $y = \{x \mid Gx\}$ for $G \subseteq F$. Then obviously $F^P \sim_R F^{P^*}$. Since we can prove $F^P < F^{P^P} < (F^{P^P})^P$ in fifth-order logic, F^P is Small^2, so that F^{P^*} is also Small^2. So by Small^2 V, we have $\{x \mid F^{P^*}(x)\}$—the power set corresponding to F^P, that is, the set of all sub*sets* of $\{x \mid Fx\}$.

Since F^{P^*} is first level, we have in third-order logic that it is smaller than its power concept $F^{P^{*P}}$, which in turn can be shown (in fourth-order logic) to be smaller than *its* power concept. There is a first-level concept, $F^{P^{**}}$, corresponding to $F^{P^{*P}}$ as F^{P^*} does to F^P. So we can repeat the foregoing reasoning to get the power set $\{x \mid F^{P^{**}}(x)\}$ corresponding to $F^{P^{*P}}$, that is, the set of all subsets of $\{x \mid F^{P^*}(x)\}$. And generally, for any *set* of *objects*, X, we have each of the ascending sequence of *powersets*—$\wp(X), \wp(\wp(X)), \wp(\wp(\wp(X))), \ldots$.

Since Nat is smaller than Nat^P, which is in turn smaller than Nat$^{P^P}$, Nat is Small^2, so we have $N = \{x \mid \mathrm{Nat}(x)\}, \wp(N), \wp(\wp(N)), \ldots$. So, taking Nat as our starting point, we can obtain first-level concepts and corresponding sets of increasing transfinite cardinality $\aleph_0, 2^{\aleph_0}, 2^{2^{\aleph_0}}, \ldots$.

Perhaps then, we may be able to get a small but nonnegligible theory of sets by supplementing fifth-order logic with Small^2 V. Perhaps . . . , but there is, once again, a more or less obvious fly in the ointment. For on the face of it, our special cases of Cantor's Theorem in higher-order logic are, in one crucial respect, perfectly general. When we proved , in third-order logic, that (the first-level concept) F is strictly smaller than its (second-level) power concept F^P, F could be any first-level concept. But with no restriction on our choice of F, we can let it be, say, *self-identical*, and following our route, show that that concept is Small^2, since it is twice smaller than the power concept of its power concept. So applying Small^2 V, we have a universal set of all self-identicals. But we shall also, by the same route, be able to show that *self-identical*^{P} is Small^2, whence we shall also have the powerset of that set, and Cantor's paradox. And similar moves with ordinal will get us the Burali-Forti. Clearly then, some further restriction is needed, if anything like the last proposal we've been reviewing is to have any chance of getting anywhere useful. In my closing section I want to indicate two rather different ways in which one might try to frame and motivate a suitable restriction.

6. Definiteness and Restricted Cardinality Relations

The first suggestion I shall discuss has its origin in a third—and if well taken—fundamental misgiving one might feel about New V as originally understood. Initially, this focuses on the suitability of "smaller than self-identity" as an explication of Goodness.

On the face of it, it makes good sense to think of one concept F as having as many instances as, or fewer instances than, another concept G only if F and G are both *sortal* concepts—that is, roughly, concepts with which are associated both criteria of application and criteria of identity. Thus on the widely accepted assumption that *brown* is a merely adjectival, nonsortal concept, it makes no sense to speak of the number of brown objects, or of there being as many brown objects as there are Fs, for any bona fide sortal F. The worry about "smaller than self-identity" stems, initially, from a doubt on this score. To get it into focus, it will be helpful to digress briefly to reconsider an objection Boolos made ([1], pp. 313–14) to Hume's Principle, turning upon the existence of the *universal number*—antizero—the number of all the objects there are, defined as $Nx : x = x$. Boolos claimed that since neo-Fregeans are happy to define zero as $Nx : x \neq x$, they can hardly refuse to admit the existence of antizero, defined as proposed. But that, he argued, is disastrous, since it puts the neo-Fregean reconstruction of arithmetic in direct conflict with ZF plus standard definitions, from which it follows that there can be no such number.

A crucial part of Wright's [13] reply to this objection was that, contrary to what Boolos claimed, the neo-Fregean has very good reason to deny that there is such a number as antizero. For the question How many Fs are there? to be in good order (and so for 'the number of Fs' to have determinate reference), F has to be a sortal concept. But *self-identical* is, Wright argued, no sortal. It seems undeniable that if F is any sortal concept, then so will be its restriction by any other concept G, irrespective of whether G is sortal or merely adjectival. For example, given that *horse* is sortal, *brown horse*, for example, must likewise be sortal, even though *brown* (or *brown thing*) is itself no sortal. But now if *self-identical* were sortal, *brown self-identical* would likewise have to be so. But since every object is necessarily self-identical, *brown self-identical* is equivalent to *brown* simpliciter—necessarily an object is brown and self-identical just in case it is brown. Since *brown* is not a sortal, neither can *brown self-identical* be one. Nor therefore, can *self-identical* be one. If this is right, then the seemingly good question How many self-identicals are there? has no determinate answer, and '$Nx : x = x$' has no determinate reference. There is no universal number. There is also space, I think, for a further doubt, about whether the contexts 'There are just as many Fs as Gs' and 'There are fewer Fs than Gs' are well defined, or have determinate truth-conditions, when one or both of F and G is nonsortal, and hence whether *self-identical* can be a suitable filler for G in those contexts. If not, then there is a further reason to deny that the proposed explication of Goodness as smaller than self-identity is satisfactory.

I think an objector might concede that a concept F must be sortal for the *how many* question and talk of the number of Fs to be in good order, and agree that self-identical is therefore, as it stands, unsuitable, but argue that we can get around this and reinstate antizero, by defining it slightly differently. First note that if F is sortal, then so is *self-identical F* (i.e., the concept for which the predicate 'x is the same F as x'—briefly '$x =_F x$'—stands). Of course, one can't get around Wright's objection to antizero, or the related difficulty I've raised, just by picking some particular sortal concept F and using *self-identical F* in place of *self-identical*. More precisely, *self-identical F* will—though sortal—fail to apply to every object unless F itself does so; but if F itself is a universal sortal, then the detour through self-identity is a waste of time, since antizero could then be just defined as $Nx : Fx$, and we could simply explain Good as *smaller than F*. We may, however, form the complex predicates 'For all F, $x =_F y$' and 'For some F, $x =_F y$'. And from these in turn we may form 'For all F, $x =_F x$' and 'For some F, $x =_F x$'. Presumably the first of these last two is true of no object whatever, and it would seem that every object whatever must satisfy the second. And this—or so it might be supposed—gives us a way out of both difficulties: just define antizero as $Nx : \exists F x =_F x$, and 'smaller than the universe' as 'smaller that the concept $\exists F x =_F x$'.

Of course, this way out is good only if the concept $\exists F x =_F x$ is itself a genuine sortal concept. The mere fact that 'For some F, $x =_F x$' is true of every object is certainly not enough to make it a sortal predicate—any more than the fact that 'x has mass' is true of every physical object is enough to make it a sortal predicate of physical objects. For being, for some F, the same F as itself to be a genuine sortal, there needs to be a criterion of identity for the objects falling under it. But so, it seems, there is. Let us abbreviate our predicate 'For some F, $x =_F x$' by 'Vx'. Suppose b and c both satisfy 'Vx'. What condition is both necessary and sufficient for b and c to be one and the same V? Well, the obvious answer is that b and c are one just in case for some single F, $b =_F c$. 'V' has thus both a criterion of application—Vx if and only if for some F, $x =_F x$—and a criterion of identity—$x =_V y$ if and only if for some F, $x =_F y$. It thus appears that V is a genuine sortal concept.

Does that show that Boolos was after all right, and Wright wrong? I don't myself think so. A concept F's being sortal is a *necessary* condition for the how many question to be in good order and for the corrresponding term '$Nx : Fx$' to have determinate reference. But I think it is arguably *not sufficient*. Indeed, it is fairly obviously insufficient, if there are—as there certainly seem to be—concepts which are sortal but indefinitely extensible in Dummett's sense (however one thinks that difficult notion is best to be explicated). The concepts of *ordinal number*, *cardinal number*, and *set* all seem to be in this case. And, since the ordinals, cardinals, and sets are among the objects that there are, it is plausible that any universal sortal concept must likewise be indefinitely extensible.[23] But in any case, there is a particular reason to doubt that '$Nx : Vx$' can have a determinate reference. For—given that our proposed definition of the universal concept V involves quantification over (sortal) concepts—it could do so only if it were already determinate what sortal concepts there are. It can scarcely be that there is a determinate answer to the question How many objects are there?—where this is construed as 'For how many x do we have $\exists F\ x =_F x$?'—unless there is a determinate answer to the question What sortal concepts are there? It is at least not obvious that there can be a determinate answer to that question.

Someone might protest, "There is no difficulty over that. For any given domain of objects, the corresponding domain of concepts is fixed. For each and every way of dividing the domain of objects, there is a concept, and those are all the concepts. If the domain of objects comprises k objects, there are thus 2^k concepts." But there is an obvious difficulty with this answer. The use of the phrase "For any given domain of objects" gives the game away. Whether there is or is not a determinate domain of objects (i.e., all objects whatever) is precisely our problem—clearly if there is, there is nothing amiss in the assumption that it has a definite cardinality, even if we are unable to determinate what that cardinality is. Thus to assume a domain of objects "given" is simply to fail to engage with the problem, or to assume it somehow solved. We cannot both assume a given domain of objects as a means of fixing the range of the quantifier 'For some F', and at the same time use that quantifier to define the sortal concept 'Vx' (i.e., 'x is an object').

If a domain of objects is already somehow fixed as comprising k objects, then it is, of course, quite right that there are 2^k concepts on the domain—at least provided that concepts are individuated extensionally. However, while there is no general objection to treating concepts extensionally—as, in effect, determined simply by what objects fall under them—it is questionable whether they are appropriately so treated in the present context, for at least two, and perhaps three, reasons.

The first, more general, reason is that it makes sense to think of concepts extensionally only if it is already determinate what objects belong to the domain on which they are to be thought of as defined. That condition may well be met in a particular case—it will be met if, for example, we are considering the domain comprising exactly the natural numbers. But it clearly cannot be assumed met in the present case.

Secondly, and more specifically, the whole point of insisting that an identity-statement $x = y$ has to be understood as asserting that x and y are one and the same F, for some appropriate sortal F, is lost, if the covering sortal F is thought of as determined purely extensionally. The point is—at least in part—that objects cannot be individuated save as instances of some sortal concept or other, so that unless some appropriate sortal is specified or understood from the context, it is simply not determinate what is being asserted, when it is said that $x = y$. If objects could be individuated simply as objects, there would be no justification for insisting that '$x = y$' must be understood as elliptical for '$x =_F y$' for some specific sortal F, such as *horse*, *person*, *number*, or the like—any identity-statement $x = y$ could be understood as claiming simply that x is the same *object* as y.

The third reason—which should, I think, weigh with the neo-Fregean, but may not be felt compelling by others—is that taking the sortal concepts to comprise just the extensionally individuated concepts on some supposed fixed domain of objects seems, in effect, simply to beg the question against the idea that abstraction principles give a way of introducing "new" sortal concepts, with a "new" range of objects falling under them.

If what I've said is right, the universal concept self-identical under F, for some sortal F, exhibits something akin to the property of indefinite extensibility. I'm not sure that it is indefinitely extensible in the usual sense, which requires, for a concept G to be indefinitely extensible, that given any definite collection of Gs, there is an object satisfying the intuitive requirements for being G which cannot be one of that collection. But even if the universal concept isn't strictly indefinitely extensible, it seems clear that it has a similar kind of indeterminacy—leaving open the question whether this coincides with indefinite extensibility, I shall say that it is *sortally indeterminate*, and for brevity say that a concept is *indefinite* if it is either indefinitely extensible or sortally indeterminate. Like Wright, I think that no determinate number can be associated with any indefinitely extensible concept. And the same goes, in my view, for sortally indeterminate concepts like the universal concept, even if they are not indefinitely extensible in the usual sense (however exactly that is to be explained).

Even if it is right that no determinate number can be assigned to any indefinite concept, it does not straightforwardly follow from this that where F is an indefinite concept, there cannot be functions from F into other concepts. It is of course true that if a concept F is indefinitely extensible, there can be no functions from F (to other concepts, definite or not) which are not themselves indefinitely extensible. But that is only to be expected and constitutes no clear objection to the idea that there may be functions from an indefinitely extensible concept to others—anyone who accepts that there are indefinitely extensible concepts will have no principled reason to deny that there are indefinitely extensible relations, including indefinitely extensible functions. If sortal indeterminacy coincides with indefinite extensibility, then the point applies to indefinite concepts quite generally. But it is not clear either that sortal indeterminacy is just indefinite extensibility under another name, or that, if it is not, there can anyway be no difficulty in principle with the idea of functions from a sortally indeterminate concept. In our only putative example of sortal indeterminacy—self-identity under F for some sortal F—the source of indeterminacy lies in the indeterminacy of the range of the second-order quantifier, and is to that extent a higher-order matter, in contrast with indefinite extensibility, which consists in the fact that no definite first-level concept can have all instances of an indefinitely extensible concept in its extension. Perhaps it could be shown that this makes no essential difference, so that we may have sortally indeterminate *functions* just as we can have indefinitely extensible ones.

I shall not here try to determine whether the concerns aired in the last few paragraphs are, in the end, well founded. Anyone in sympathy with them ought, it seems, to view the free-wheeling, unrestricted talk of double smallness involved in our original formulation of Small^2 V with some suspicion, at least unless she can view the unqualified statement that $F < G < H$ as a merely heuristically useful way of expressing the idea that F is a definite (i.e., not indefinite) concept. But even one not moved by those concerns ought to be able to discern the shape of a possible restriction on the application of Small^2 V—to the effect that we may take the fact that for some concepts G and H, $F < G < H$ as entitling us to conclude that there is a set of Fs only when G, and hence F, is a definite concept. If such a restriction can be imposed, it will straightforwardly block the paradoxes—both Cantor's and the Burali-Forti—that threaten the proposal sketched in the preceding section. However, if the paradoxes are to be blocked by imposing a restriction to definite concepts in the application of the set-theoretic abstraction, it is no longer clear that the shift to interpreting Good as Small^2 is doing useful work. The original point of that shift was to block the derivation of Global Well-Ordering which convicted New V of a violation of Wright's first conservativeness constraint. But restricting Good concepts to definite ones would seem by itself enough to achieve that result, since the derivation of the Burali-Forti from the assumption that *ordinal* is definite would then force only to the conclusion that *ordinal* is indefinite, which does not yield Global Well-Ordering.

The second of the two suggested restrictions I want to mention, by contrast, leaves Small^2 V playing a significant role in the enterprise and can be stated rather more briefly. This exploits two thoughts. The first is that the notion of what it is for one concept to be smaller than another, involved in the definition of Small^2, need not be taken as fixed in advance and independently of the neo-Fregean enterprise. The neo-Fregean is, so far as I can see, perfectly free to stipulate a meaning for it that suits his purposes. The second is that, on the neo-Fregean approach to set theory which I have been exploring in the last few sections, there is no aspiration to develop that theory as a free-standing theory, based exclusively on distinctively set-theoretic abstraction principles. On the contrary, we are already embracing the idea that much of the ontology of the theory, and hence much its power, is to be provided by other abstraction principles—such as Hume's Principle and Cut principles—which do not specifically concern sets at all, but objects of other kinds. Crucially, these other abstractions are, when acceptable, to be conceived of as in good standing independently of the development of any abstractionist set theory. Their acceptability is to be thought of, rather, as a matter of their compliance with whatever constraints—some of which we have touched on above—govern legitimate abstraction in general.[24] In the context of these two thoughts, a natural proposal is that the neo-Fregean may go a step further and take the sortal concepts introduced via independently acceptable abstraction principles as his basis for the identification of a privileged class of concepts which may serve to anchor, as it were, a restricted $<$ relation for the purposes of Small^2 V. In a little more detail, the idea would be that $F < G < H$ holds, in the relevantly restricted sense, *only when* G and H are concepts independently in good standing courtesy of other acceptable abstraction principles, or power concepts of such concepts, or power concepts of power concepts of such concepts, and so on.

Obviously both of the suggestions canvassed here are merely directions for further investigation, without which one can have little confidence that either of them will withstand closer scrutiny or result in a satisfactory and agreeably powerful abstractionist theory of sets. And there may, of course, be other possible ways to impose the restriction(s) on Small^2 V which we have seen to be needed. I must leave that work for another occasion. I hope, at least, that the present discussion will have served to identify some of the difficulties facing an abstractionist development of set theory and, perhaps, some strategies for dealing with them worth further thought.

Notes

1. “Roughly” because it is desirable to count as abstractions some principles which don’t, as they stand, have precisely this form. In the only case that matters for present purposes, Σ is a function of two arguments, not one, and the RHS relation is 4- rather than 2-termed. One can easily deal with this, either by introducing ordered pairs by abstraction or by generalizing the notion of an equivalence relation.

2. In Frege [7], §§ 62–67.

3. Reflexivity and Symmetry are obvious. For Transitivity, suppose
   (a) $\mathrm{Good}(F) \vee \mathrm{Good}(G) \rightarrow \forall x(Fx \leftrightarrow Gx)$ and
   (b) $\mathrm{Good}(G) \vee \mathrm{Good}(H) \rightarrow \forall x(Gx \leftrightarrow Hx)$.
   If the antecedents of both (a) and (b) are both false, then $\neg(\mathrm{Good}(F) \vee \mathrm{Good}(H))$, whence
   (c) $\mathrm{Good}(F) \vee \mathrm{Good}(H) \rightarrow \forall x(Fx \leftrightarrow Hx)$.
   Likewise, if the consequents of both (a) and (b) are true, then $\forall x(Fx \leftrightarrow Hx)$, whence (c). If (a)’s antecedent is false but (b)’s consequent is true, then $\neg\mathrm{Good}(H)$, whence $\neg(\mathrm{Good}(F) \vee \mathrm{Good}(H))$ and so (c) again. Similarly if (a)’s consequent is true but (b)’s antecedent is false. Essentially this is the proof given in [12], fn. 32.

4. This last point may not, in itself, be as damaging as might be at first supposed, since there are known to be abstractionist methods of obtaining the real numbers which avoid any essential reliance upon an underlying set theory. But a neo-Fregean should be concerned to develop as powerful a set theory as can be done using the resources—centrally, but not necessarily only, abstraction principles—at his disposal, and so may hope to do better than New V.

5. The example is not Boolos’s own but a very similar one taken from [12], pp. 289–91.

6. “First,” because Wright introduces a second quite distinct conservativeness constraint which will be discussed briefly later.

7. See Field [6] and subsequent papers collected together in Field [5].

8. For a fuller discussion and more precise articulation of the proposed constraint, see [12], pp. 295–97, especially fn. 49.

9. The proviso is not merely decorative—we shall later consider possible reasons to feel misgivings about it.

10. I am heavily indebted to Wright for this suggestion and much useful discussion of it.

11. Cook has pointed out that while the Global Well-Ordering result is blocked, one still gets a signficant result—that, since *ordinal* is not Small^2, there cannot be any concept that is bigger than *ordinal* but smaller than the universe. Though a weaker result, this is, Cook remarks, independent of second-order ZFC. Perhaps so, but it is not a nonconservativeness result for a set theory based on New V with Small interpreted as Small^2 (or Small^2 V, that is, Small V with Small so reinterpreted) in the sense in which Global Well-Ordering is a nonconservativeness result for original New V. Global Well-ordering implies well-ordering for the "old" ontology—that is, the universe of objects as a whole, including those not included among the abstracts provided by New V—in violation of Wright's first conservativeness constraint (see [12], p. 296 ff.). By contrast, if our set-theory is to be based on Small^2 V, then the theory of ordinals will be naturally construed as a part of it, and the ordinals themselves will be a species of the new abstracts so introduced. That this species is at most singly small would seem to have no bearing on anything—cardinal or ordinal—essentially to do with the old ontology.

12. On (D)-type abstractions, see [12], Section IV. φ is any property of concepts for which coextensiveness is a congruence, that is, $\varphi(F)$ and $\forall x(Fx \leftrightarrow Gx)$ jointly entail $\varphi(G)$. New V comes from the schema by reading φ as Big. New V as formulated here is not strictly of the form (d), but is obviously equivalent to an abstraction of that form.

13. [12], p. 303. A more precise formulation is given at p. 304.

14. See note 10.

15. Largely, but not entirely, if one works with a (B)-type abstraction, such as New V with Good understood as Small^2 but entirely if one works with a conditionalized, (A)-type abstraction such as Small^2 V. I shall make no attempt to adjudicate here whether there are compelling reasons to favor one approach over the other. Very roughly, the fragment of standard set theory which one can recover without appeal to non-set-theoretic abstraction principles—though existentially very weak—is larger if one works with a (B)-type principle such as New V rather than an (A)-type principle such as Small^2 V. This might be thought a reason for preferring New V over Small^2 V.

16. A good deal of what I shall be saying applies, with relatively minor adjustments, to a development based on New V with Good interpreted as Small^2. It would unduly complicate the discussion to keep both alternatives in play throughout.

17. Crucially, because we want the same real numbers to be applicable in measuring quantities of different kinds.

18. For some explanation of this, see [8], Section II.

19. By Cook in [4]. I am grateful to him both for letting me see earlier versions of this paper, and for helpful discussion of it.

20. Including Shapiro, Cook, and one (still) anonymous referee of the paper in which some of these ideas were first put forward.

21. This statement of the theorem and its proof are taken verbatim from [4].

22. I am grateful to Potter for this suggestion.

23. This will be so, if we can assume that if F is indefinitely extensible and $\forall x(Fx \to Gx)$, then G is likewise indefinitely extensible.

24. For further discussion of which, see, as well as [12], Hale and Wright [9].

References

[1] Boolos, G., "Is Hume's Principle analytic?" pp. 301–15 in *Logic, Logic, and Logic*, edited by R. Jeffrey, Harvard University Press, Cambridge, 1998. Zbl 0972.03509. MR 2000b:03005. 391

[2] Boolos, G., "Iteration again," pp. 88–104 in *Logic, Logic, and Logic*, edited by R. Jeffrey, Harvard University Press, Cambridge, 1998. Zbl 0972.03511. MR 2000b:03005. 382

[3] Boolos, G., "Saving Frege from contradiction," pp. 171–82 in *Logic, Logic, and Logic*, edited by R. Jeffrey, Harvard University Press, Cambridge, 1998. Zbl 0972.03502. MR 2000b:03005. 381

[4] Cook, R. T., "The state of the economy: Neo-logicism and inflation," *Philosophia Mathematica. Series 3*, vol. 10 (2002), pp. 43–66. Zbl 01779438. MR 1875536. 397

[5] Field, H., *Realism, Mathematics and Modality*, Basil Blackwell Inc., New York, 1989. MR 92b:03003. 396

[6] Field, H. H., *Science without Numbers. A Defence of Nominalism*, Princeton University Press, Princeton, 1980. Zbl 0454.00015. MR 82d:03008. 396

[7] Frege, G., *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, Harper & Brothers, New York, 1960. Zbl 0123.24604. MR 23:A2287. 395

[8] Hale, B., "Reals by abstraction," *Philosophia Mathematica*, vol. 8 (2000), pp. 100–123. Reprinted in [10], pp. 399–420. Zbl 0968.03010. MR 2001i:03015. 384, 397

[9] Hale, B., and C. Wright, "Implicit definition and the a priori," pp. 286–319 in *New Essays on the A Priori*, edited by P. Boghossian and C. Peacocke, Oxford University Press, Oxford, 2000. Reprinted in [10], pp. 117–50. 397

[10] Hale, B., and C. Wright, *The Reason's Proper Study*, Oxford University Press, Oxford, 2001. Zbl 01764654. 397, 398

[11] Shapiro, S., and A. Weir, "New V, ZF and abstraction," *Philosophia Mathematica*, vol. 7 (1999), pp. 293–321. The George Boolos Memorial Symposium (Notre Dame, IN, 1998). Zbl 0953.03061. MR 2000j:03006. 382

[12] Wright, C., "The philosophical significance of Frege's Theorem," pp. 201–45 in *Language, Thought, and Logic. Essays in honour of Michael Dummett*, edited by R. G. Heck, Jr., Oxford University Press, Oxford, 1997. Reprinted in [10], pp. 272–306. Zbl 0938.03508. MR 2000h:03005. 381, 383, 395, 396, 397

[13] Wright, C., "Is Hume's Principle analytic?" *Notre Dame Journal of Formal Logic*, vol. 40 (1999), pp. 6–30. Special issue in honor and memory of George S. Boolos (Notre Dame, IN, 1998). Reprinted in [10], pp. 307–32. Zbl 0968.03009. MR 2002a:03014. 391

Acknowledgments

I am especially grateful to Adam Rieger, Stewart Shapiro, and Crispin Wright for critical and constructive discussion of the ideas in this paper.

Department of Philosophy
University of Glasgow
65–69 Oakfield Avenue
Glasgow G12 8QQ
Scotland
UNITED KINGDOM
rlh@arts.gla.ac.uk