Philosophical Studies (2006) 129:545–574
DOI 10.1007/s11098-004-3388-1

ØYSTEIN LINNEBO

EPISTEMOLOGICAL CHALLENGES TO MATHEMATICAL PLATONISM

ABSTRACT. Since Benacerraf's "Mathematical Truth" a number of epistemological challenges have been launched against mathematical platonism. I first argue that these challenges fail because they unduely assimilate mathematics to empirical science. Then I develop an improved challenge which is immune to this criticism. Very roughly, what I demand is an account of how people's mathematical beliefs are responsive to the truth of these beliefs. Finally I argue that if we employ a semantic truth-predicate rather than just a deflationary one, there surprisingly turns out to be logical space for a response to the improved challenge where no such space appeared to exist.

The claims of mathematics purport to refer to mathematical objects. And most of these claims are true. Hence there exist mathematical objects. Though extremely brief, this argument has great force. Ever since it was introduced by Gottlob Frege, more than a century ago, it has been our best argument for *mathematical platonism* – the view that there exist mathematical objects.[1]

Despite the force of this argument, mathematical platonism has been found to be epistemologically problematic. Unlike ordinary objects, such as tables and chairs, mathematical objects cannot be perceived. Nor can they be observed with the help of modern technology, as electrons and distant celestial objects can. So how can we have knowledge of mathematical objects? This epistemological challenge is inspired by Paul Benacerraf's famous 1973 article "Mathematical Truth", which argues that two fundamental desiderata on the philosophical analysis of mathematics conflict with one another. On the one hand, we want "a homogeneous semantical theory in which semantics for the propositions of mathematics parallel the semantics for the rest of the language" (p. 403). This would mean that the truth of a mathematical sentence S requires the existence of suitable mathematical objects to which S's singular terms refer and over which its quantifiers range. On the other hand, we want "the account of mathematical truth [to] mesh with a reasonable epistemology" (*ibid*.). And according to Benacerraf, "a reasonable epistemology" will include some causal theory according to which knowledge requires a causal connection between the knower and the known. But since mathematical objects are abstract and therefore causally inefficacious, no such connection would be possible. So Benacerraf's two desiderata appear to imply that mathematical knowledge is impossible.

Some philosophers (such as Benacerraf) find this conclusion unattractive and therefore regard the above considerations as a challenge to mathematical platonism that can and must be met. Other philosophers (such as Hartry Field) think that these considerations form the core of a devastating objection to mathematical platonism. But either way, it is natural to protest that Benacerraf's considerations are biased against mathematics at the very outset. (This was certainly *my* response, when I first read "Mathematical Truth".) By asking for a causal connection between the epistemic agent and the object of knowledge, Benacerraf treats platonistic mathematics much like physics and the other garden-variety empirical sciences. But mathematics is different. So philosophers have no right to subject it to epistemological standards that have their home in the domain of contingent empirical knowledge.[2] Since mathematics does not purport to discover contingent empirical truths, it deserves to be treated differently. And when we only treat it as it deserves – as mathematics rather than as physics – the epistemological problem Benacerraf points to will surely go away. I will call this *the Natural Response*. This response has great initial plausibility, and I believe it suffices to undermine the best existing epistemological challenges to mathematical platonism.

In this paper I develop an improved epistemological challenge which is immune to the Natural Response. I also make some suggestions about how this improved challenge may be met. Let me be a bit more specific. In Section 1 I argue that it is not enough for the advocate of the Natural Response to assert that mathematics is autonomous but that she owes us some positive account of mathematical knowledge, or, alternatively, some explanation why no such account can be required. In Sections 2 and 3 I discuss some attempts, inspired by John P. Burgess, David Lewis, and Gideon Rosen,[3] to supplement the Natural Response with some rather minimal accounts of mathematical knowledge. I find serious faults with these attempts. Exploiting these faults, I propose, in Section 4, an improved epistemological challenge. This improved challenge is based on the demand for a particular brand of explanation that is carefully designed to be sufficiently non-trivial to represent a real challenge, while remaining neutral enough to be immune to the Natural Response. Very roughly, what I demand is an account of how people's mathematical beliefs are responsive to the truth of these beliefs. Finally, in Section 5, I argue that if we employ a semantic truth-predicate rather than just a deflationary one, there surprisingly turns out to be logical space for a response to the improved challenge where no such space appeared to exist.

The topic of this paper is just one instance of a much greater philosophical problem concerning knowledge of truths that are not straightforwardly contingent. In addition to mathematical truths, I have in mind truths about what is possible and what is necessary, as well as purported truths of ethics or even theology. These truths are naturally taken to be about possible worlds, moral properties, or the divine. These entities and their properties do not naturally fit into our best understood epistemological model, which is based on knowledge of contingent empirical truths. So in each of these areas, there is an epistemological challenge similar to Benacerraf's. And to each of these challenges, philosophers have responded by some version of the Natural Response. The debate about the epistemology of mathematical platonism has important structural similarities with other instances of this greater philosophical debate. The similarities are particularly close in the case of modal knowledge. Most of what I say in this paper carries over directly to the debate about such knowledge and about the ontological status of possible worlds.

1. THE BEST EXISTING EPISTEMOLOGICAL CHALLENGE

As mentioned above, Benacerraf's appeal to causal theories of knowledge is problematic. But this appeal is just a superficial feature of Benacerraf's challenge. The philosophical puzzle that Benacerraf called to our attention is far more robust. In this section, I will present a version of Benacerraf's challenge that eliminates this problematic appeal to causal theories of knowledge. I believe this version, due primarily to Hartry Field, is the best epistemological challenge that can be patched together from the existing literature.

The structure of the challenge is as follows. First the challenger grants, for the sake of the argument, that there are abstract mathematical objects. Then she argues that even so, our beliefs in platonistic mathematics would not be epistemically justified. Hence she infers that we can never be justified in believing in abstract mathematical objects. Based on this, she demands that the platonist stop claiming that mathematical platonism is true. Even if it succeeds, this challenge will not quite prove that mathematical platonism is false. But it will do something equally effective: it will show that no one can ever be justified in asserting that mathematical platonism is true.[4]

Clearly, the hard part of this challenge is to argue that even if platonism is true, our beliefs in platonistic mathematics will not be justified. Field suggests that an argument to this effect can be given, based on the question why mathematicians' beliefs are reliable.[5] Mathematicians take themselves to be highly reliable about mathematical matters. That is, they take the schema

If mathematicians accept S, then S is true to be true in nearly all instances where S is a mathematical sentence. I will call claims of this sort *reliability claims*.[6] If mathematicians are reliable in this sense, this cannot be just by chance. Admittedly, it is *conceivable* that there be a person who, when presented with a mathematical sentence, decides whether to accept it by tossing a coin, but who happens always to get it right. Call this person *the Lucky Fool*. But we are convinced that competent mathematicians are very different from the Lucky Fool. We are convinced that their tendency to hold true mathematical beliefs, unlike the Lucky Fool's, has some explanation.

Based on these considerations, the challenger can argue that the beliefs of the mathematical platonist are not justified. First, she argues, the platonist has to accept the mathematical reliability claim. Second, unless it is at least in principle possible to explain this reliability claim, the platonist's mathematical beliefs cannot be justified. And third, no such explanation is possible. This new challenge avoids commitment to any particular epistemological doctrine – at least so far. As Field observes, this challenge can be formulated without even using the "term of art 'know'" (1989, p. 230).

It is hard to deny the first claim on which this new challenge is based. We do indeed take our mathematical beliefs to be reliable. So if the mathematical platonist is to object to the reliability claim, this objection must be directed, not at the claim itself, but at the way it is formulated. Two questions about its interpretation need to be addressed.

The first question concerns the truth-predicate that is employed in the consequent of the reliability claim. Should this be taken to be a deflationary or a semantic truth-predicate? A truth-predicate is said to be *deflationary* if it essentially just disquotes. With a deflationary truth-predicate, 'S is true' is cognitively synonymous, for those who understand it, with S itself.[7] By contrast, a *semantic* truth-predicate takes into account semantic properties of the sentence S. On this construal, 'S is true' makes a claim, not only about the subject matter of S, but about the meaning of S, to the effect that it is a correct account of this subject matter. So with a semantic truth-predicate, 'S is true' is not cognitively synonymous with S.[8]

It is widely assumed that a deflationary truth-predicate suffices for the epistemological challenge.[9] All that is needed for this challenge to work, it seems, is that mathematicians' beliefs be correlated with the mathematical facts. And in order to express this correlation, a deflationary truth-predicate will do. Moreover, a deflationary truth-predicate uses smaller philosophical resources than a semantic one. So if the challenge can be adequately formulated with a deflationary truth-predicate, this would be advantageous. I will therefore begin by formalizing the mathematical reliability claim by means of a deflationary truth-predicate. In Section 5, however, I will suggest that it may in the end be more fruitful to employ a semantic truth-predicate.

The second question concerns the 'if-then' conditional. Should this be formalized as a material conditional or as some stronger, non-truth-functional conditional? Field favors the former option.[10] He argues that this provides a sufficient basis for an epistemological challenge. All the challenger needs is that the mathematical platonist be committed to an *actual* correlation between mathematicians' beliefs and the mathematical truths. This correlation can be expressed using only a material conditional. And since this is a correlation between two very different systems of facts, it is sufficiently puzzling to call for an explanation.

Moreover, the material conditional rendering of the reliability claim provides the least problematic basis for the epistemological challenge. Were we to formalize the reliability claim by means of a counterfactual conditional, say, this counterfactual would have to be of the form 'Had the mathematical facts been different, then so would the psychological ones'. For among these two systems of facts, it is the mathematical ones that are more fundamental. If any of these systems of facts is determined by the other, it is the psychological ones by the mathematical ones, not vice versa. It is instructive to compare this correlation with that between a fuel gauge and the level of fuel in a tank. Clearly, it is the fuel gauge that depends counterfactually on the level of fuel, not vice versa. If, for instance, the gauge were to malfunction, this would not affect the level of fuel in the tank. Likewise, if a mathematician were to "malfunction", this would not affect the mathematical truths.

However, the correlation we're interested in is complicated by the fact that mathematical truths, unlike truths about the level of fuel in a tank, are necessary. To make sense of how mathematical beliefs track the mathematical facts, we would have to make sense of *varying* these facts. But since mathematical facts are necessary, we have no idea how to do this. So it appears that we have no choice but to formalize the relevant reliability claims by means of the material conditional. Since the material conditional rendering of the mathematical reliability claim seems both adequate and inevitable, I adopt it for the time being.[11]

Summing up, I will formalize the mathematical reliability claim as

(1_d) $\forall S$(mathematicians accept $S \rightarrow S$ is true$_d$)[12].

I will use subscripts 'd' and 's' to distinguish between formalizations using a deflationary and a semantic truth-predicate. Since this formalization uses no contentious concepts, the platonist has to accept it.

The second claim on which this epistemological challenge is based says that, for beliefs in platonistic mathematics to be justified, it must be possible to explain the mathematical reliability claim. What sort of explanation does the challenger have in mind? Field does not explicitly say. But he gives an example of what he takes to be a successful explanation of a reliability claim, namely one concerned with the reliability of physicists' beliefs about electrons:[13]

(2_d) $\forall S$(physicists accept S) $\rightarrow S$ is true$_d$).

Field suggests that (2_d) can be explained by undertaking a scientific study of particle physicists, their experimental procedures, and their laboratory equipment. This explanation will be based on a physical account of how electrons affect the measuring apparatus and of how this apparatus in turn affects the eyes of the physicists. It will also involve a psychological account of how these sensory stimulations give rise to physicists' beliefs about electrons and of how these beliefs in turn inform their decision to accept or reject sentences concerned with electrons.

What this example makes clear is that Field is not demanding some absolute justification of science, for instance of the sort sought by foundationalist epistemology. Nor is he demanding a defense of science against skeptical arguments. Were any of these his demand, his own account of (2_d) would be viciously circular because of its reliance on physics. Rather, what Field demands is a *scientific* explanation of the reliability claims. Though still vague, this shows at least which ballpark Field is playing in.

But is it reasonable to require that, for mathematical beliefs to be justified, it be possible to give a scientific explanation of the mathematical reliability claim? I believe it is. To see why, it is instructive to contrast this condition with the related but stricter condition that we *actually give* a scientific explanation of the associated reliability claim. This stricter condition faces relatively uncontroversial counterexamples. For instance, people were presumably justified in trusting at least some of their perceptual beliefs before they knew enough about the human perceptual system to give a scientific explanation of its reliability. All our condition requires, however, is that it be *in principle possible* to explain this reliability claim. And this condition seems eminently plausible. In particular, long before people knew much about our perceptual system, they had reason to believe that it was possible to explain the reliability of their perceptual beliefs.

Field's third claim is that it is impossible to give a scientific explanation of the mathematical reliability claim. Since mathematical objects do not participate in the causal order, (1_d) clearly cannot be explained in the same way as (2_d). But Field needs a stronger claim than that: he needs that there can be no scientific explanation of (1_d) whatsoever. This is where Field's argument becomes problematic. To defend this stronger claim, Field claims that mathematical objects don't just happen to be causally inaccessible, as distant regions of space–time do according to special relativity, but that "the truth values of our mathematical assertions depend on facts involving platonic entities that reside in a realm outside of space–time'' (1989, p. 68) and thus are causally isolated from us even in principle. According to Field, this radical separation of the platonic entities from our physical universe makes it impossible to give any kind of scientific explanation of (1_d).

Field's argument assumes that a scientific explanation of a correlation must involve a causal connection between at least some of the correlated items. But it is hard to see why the platonist should accept this assumption. For one thing, there appear to be perfectly good scientific explanations that have no causal component. Consider, for instance, the correlation between a first-order theory's being consistent and its having a model. This correlation is explained by the completeness theorem, which does not say a word about causality. For another, mathematics is very different from physics and the other empirical sciences. So *prima facie*, the fact that the mathematical reliability claim cannot be explained in the same way as the physical one should be neither surprising nor particularly worrisome. For we had no reason in the first place to expect that it could be.

So Field's challenge fails as an *objection* to mathematical platonism. But this failure does not undermine its force as a *challenge*. Even if it is inappropriate to demand an ordinary causal explanation of the mathematical reliability claim, it remains an open question how this claim *is* to be explained, or, if no explanation is possible, *why* none is required. In order to succeed, the Natural Response must answer this open question.

2. THE BORING EXPLANATION

In this section and the next I will consider some attempts to answer this question, inspired by John P. Burgess, David Lewis, and Gideon Rosen. These attempts consist of some minimal, non-causal explanations of the mathematical reliability claim, reinforced by arguments that no further explanation can be required.

To see what the first minimal explanation is, imagine someone asks you how it is that your neighbor, Jones, is a reliable mathematician. A natural response would be the following. Jones went to school, where he took courses in mathematics. Being a good student, Jones learnt a good deal of mathematics, and he learnt how to apply it to problems that are given to him. Moreover, the mathematical theory Jones was taught is true: its axioms are mathematical truths, and its rules preserve mathematical truth.[14] In fact, for every competent mathematician, there is a story of this sort. Every competent mathematician has learnt some true mathematical theory. And this explains why his mathematical beliefs are reliable.

I will call this *the Boring Explanation*. Obviously, it is very boring compared to the explanation I just sketched of the physical reliability claim. It does not at all attempt to explain *the connection* between mathematicians' dispositions to accept sentences and these sentences' being true. But as we've seen, it is illegitimate to object to the Boring Explanation merely on the ground that it is different from the kind of explanations we find in the empirical sciences. Instead, I will now argue that the Boring Explanation is inadequate because it fails adequately to distinguish competent mathematicians from the Lucky Fool, who decides whether to accept a mathematical sentence by tossing a coin but who happens always to get it right.

To develop this objection, I need to analyze the situation in bit more detail. I will idealize slightly and assume that every competent mathematician operates with the same set of axioms and rules. Since this idealization only strengthens the Boring Explanation, I can safely adopt it.

With this idealization in place, the Boring Explanation proceeds as follows. It begins by "factoring" the mathematical reliability claim (1_d) into two components:

(1P) $\forall$S(mathematicians accept S $\rightarrow$ S follows from Σ),

($1M_d$) $\forall$S(S follows from $\Sigma \rightarrow$ S is true_d).

These components, which I will call *the psychological fact* and *the mathematical fact*, jointly imply (1_d). Next, the Boring

Explanation explains these two components individually. The psychological fact is explained by mathematicians' somehow having learnt the theory Σ. And the mathematical fact is explained by observing that every axiom of Σ is a mathematical truth and that the rules of Σ preserve mathematical truth.

In fact, this explanation establishes slightly more than (1P). It shows that every sentence mathematicians are prepared to accept when suitably prompted follows from Σ. That is, it establishes

$$(1P') \quad \forall S(\text{mathematicians are disposed to accept } S \rightarrow S \text{ follows from } \Sigma).$$

No counterpart to (1P′) can be established in the case of the Lucky Fool. For the Lucky Fool is not guided by any mathematical theory. It is pure chance that his coin tosses always give the right answer. And we know that luck in the past does not raise the probability of luck in the future. Every time the Lucky Fool encounters a new mathematical sentence, there is a fifty percent chance he will get it wrong. So the sort of explanation that can be given of the Lucky Fool's success will not support counterfactual claims about what would have happened had the Lucky Fool been confronted with some novel mathematical sentence. This means that the Boring Explanation identifies at least one feature that distinguishes competent mathematicians from the Lucky Fool.

However, there is a much more important distinction that the Boring Explanation fails to identify. To see this, consider *the Swamp Calculator*. The Swamp Calculator is a molecule by molecule copy of our best device for proving mathematical theorems. But the Swamp Calculator has an unusual origin: One day it assembled spontaneously out of a swamp due to random physical processes. Despite this unusual origin, we can give a boring explanation[15] of the Swamp Calculator's success that completely parallels the Boring Explanation of (1_d). The Swamp Calculator instantiates a sophisticated program. This program disposes it to "accept" a mathematical sentences S only if S follows from a particular theory. And because of certain mathematical features of this theory, all of its theorems are true.

Unlike the case of the Swamp Calculator, however, it is not just by chance that competent mathematicians operate with a theory Σ that is true. They hold their mathematical theory Σ *for a reason*. And this reason is somehow connected with *the truth* of this theory. Had they not taken this theory to be true, they would not have held it. So an adequate explanation of the mathematical reliability claim will have to explain why the same theory Σ figures both in the psychological fact and in the mathematical. But the Boring Explanation explains these two facts separately. It says nothing about the relation between these two occurrences of Σ. Thus it fails to distinguish competent mathematicians from the Swamp Calculator.

At this point, I think the best strategy for the defenders of the Boring Explanation is to admit that their explanation is boring and to argue instead that we have no right to insist on a more interesting explanation. Some correlations admit nothing but boring explanations. Perhaps (1_d) is one of them. I will now examine three arguments to that effect.

The first argument attempts to show that (1_d) is not strictly speaking a correlation at all.[16] It does this by analyzing (1_d) in a possible worlds framework. (1_d) is supposed to be a correlation between certain psychological facts, recorded by (1P), and certain mathematical facts, recorded by ($1M_d$). The psychological side of the correlation is straightforward: (1P) is true in all possible worlds in which mathematicians accept those mathematical sentences that our mathematicians actually accept. But the mathematical side offers a surprise: ($1M_d$) is true in all possible worlds whatsoever. To see this, recall that we are operating with a deflationary truth-predicate and that 'S is true$_d$' for this reason is cognitively synonymous with S itself. Hence, when S is a mathematical truth, so is 'S is true$_d$'. Moreover, since facts about logical consequence are mathematical, so will be the instances of ($1M_d$); for they are of the form

$$\text{S follows from } \Sigma \rightarrow \text{S is true}_d.$$

Hence it follows that ($1M_d$) is a mathematical truth and thus true in all possible worlds. This means that the mathematical side of the alleged correlation drops out, leaving only the psychological fact (1P) to be explained. And this fact is adequately explained by the Boring Explanation.

This argument has an air of sophistry. I think the most natural response to it is that the apparent failure of (1_d) to express a genuine correlation is an artifact of the possible worlds framework that is brought to bear. Since this framework identifies all necessary propositions and hence all propositions of pure mathematics, it's no wonder it fails to register a correlation! But intuitively, we take propositions that this analysis identifies to say very different things. For instance, we take '$2+2=4$' and the Axiom of Choice to express entirely different propositions. In order to capture these intuitive distinctions, we will either have to leave the possible worlds framework altogether or adopt a more fine-grained space of possibilities. In either case, the correlation between psychological facts and mathematical facts will re-emerge.

The second argument in defense of the Boring Explanation suggests that, when the correlation to be explained has no counterfactual force, a boring explanation is all the explanation we need.[17] This suggestion receives support from our imaginary examples of the Lucky Fool and the Swamp Calculator: These correlations lack the relevant counterfactual force, and they are adequately explained by boring explanations. To gather more evidence, consider the following two real examples:

$$\forall x(x \text{ has been President of the US} \rightarrow x \text{ is not a child of Swedish immigrants}), \quad (3)$$

$$\forall x(x \text{ has been President of the US} \rightarrow x \text{ is not a woman}). \quad (4)$$

These examples point to correlations between the occupants of the American presidency and certain properties of these occupants. Although the first correlation (3), may have an interesting explanation, such as a lack of political involvement among Swedish immigrants, or prejudices in the American population, or what not, most likely its only explanation will be a boring one, consisting of 54 individual explanations of why the winner of each presidential election, who happened not to be the child of Swedish immigrants, succeeded in getting more electoral votes than any other candidate. By contrast, the second correlation (4), clearly needs a deeper explanation. Now, what is the relevant difference between (3) and (4)? The answer seems to be that the latter, unlike the former, has counterfactual force. Even if things had gone slightly differently in some presidential campaign, a woman still would not have been elected president. To be adequate, an explanation of (4) must explain this counterfactual claim.

If this suggestion is correct – that is, if a correlation needs an interesting explanation only when it has counterfactual force – then the Boring Explanation of (1_d) may be sufficient. For it is hard to see how (1P) could depend counterfactually on ($1M_d$). Since mathematical truths are necessary, we have no idea how to vary them in the way that would be required to make sense of this counterfactual.

However, this suggestion is *not* correct as it stands. We see this by reflecting on the completeness theorem for first-order logic. Although the correlation between consistency and satisfiability of first-order theories has no counterfactual force, it is not adequately explained by a boring explanation. It needs an explanation that connects the properties of consistency and satisfiability, such as the explanation afforded by the familiar proofs of the theorem.

The third argument in defense of the Boring Explanation suggests that, as long as a correlation is not very pervasive, no deeper explanation is needed.[18] Now, by itself this suggestion too is insufficient. For (4) is a non-pervasive correlation that still needs a deeper explanation. However, non-pervasiveness is promising as a patch on the second argument. So this is how I will interpret the third argument. Then, the suggestion is that a correlation requires nothing but a boring explanation just in case it neither has counterfactual force nor is particularly pervasive. This suggestion gives plausible diagnoses of our examples. (4) needs more than a boring explanation because it has counterfactual force. The completeness theorem needs more because it involves an infinite and thus very pervasive correlation. But the correlation (3) lacks both counterfactual force and pervasiveness and hence is adequately explained by a boring explanation.

If this suggestion is correct, it may allow us to defend the Boring Explanation. We have already seen that the mathematical reliability claim (1_d) does not have the right sort of counterfactual force. And arguably, (1_d) rests upon a non-pervasive correlation. To see this, observe that both of its components, (1P) and ($1M_d$), are generated by the axiomatic theory Σ. Since knowledge of logic is not currently at issue, the mathematical reliability claim (1_d) will thus be explained if we can only explain the correlation between mathematicians' acceptance of *the axioms* of Σ and these axioms' being true. This yields an enormous reduction of our problem. For if Σ is a theory of sets, it can be chosen so as to contain only a small number of axioms. In fact, by conjoining these axioms, the theory Σ can be reduced to just one non-logical axiom "super-axiom" A.[19] This reduces the mathematical reliability claim to the correlation between mathematicians' acceptance of A and A's being true. But this is just a conjunction of two sentences, not a pervasive correlation!

I have two problems with this defense of the Boring Explanation, one technical and one more intuitive. The technical problem is that it is debatable whether our total mathematical competence can be captured by any finite formal system. After all, Gödel's incompleteness theorems shows that even arithmetic cannot be finitely axiomatized.

The defender of the Boring Explanation can attempt to bypass this problem by invoking full second-order logic. First, he can argue that second-order logic is ontologically innocent and that it qualifies as pure logic. Then, he can let Σ be a finitely axiomatized second-order set theory and argue that the incompleteness in question resides in the logical part of the theory, not in the mathematical.[20] This would show the incompleteness to be a problem in the epistemology of logic, not in that of mathematics.

However, I have serious doubts about the first step of this argument. The most promising defense of the logicality and ontological innocence of monadic[21] second-order logic is due to George Boolos.[22] Boolos first proves that monadic second-order logic can be interpreted in a theory of plural quantification. This is an indisputable technical result. Then he gives a philosophical argument that the requisite theory of plural quantification is ontologically innocent and qualifies as pure logic. But this philosophical argument has been challenged elsewhere, where it is denied that the theory of plural quantification qualifies as pure logic and that it is ontologically innocent.[23]

My second problem with the above defense of the Boring Explanation is more intuitive and completely independent of the first. The defense in question attempts to read off the pervasiveness of a correlation from its syntactic representation. But there is no reason to think that the syntactic representation of a correlation should give any information about its pervasiveness. For instance, we can conjoin, on the one hand, 54 sentences listing who won the various presidential elections throughout the history of the United States, and on the other, 54 sentences to the effect that each winner was male. These two sentences, conjoined with a third sentence stating that these are all the U.S. presidential elections, yields a sentence logically equivalent to (4). But obviously, this does not remove the need for a deeper explanation of (4).

So perhaps the defender of the Boring Explanation will concede that we must "unpack" the super-axiom A to retrieve a finite number of ordinary set theoretic axioms. But even so, he could argue that since this yields less than a dozen axioms, mathematicians' reliability with respect to them is not much of a correlation. But this argument too would be unconvincing. For there is reason to think that the axioms of set theory contain highly compressed mathematical information. Indeed, they contain within them all of classical mathematics. It took mathematicians decades of hard work to find such a minimal basis for mathematics. So it is implausible to deny that there is a connection between these axioms' being accepted as true and their actually being true.

3. THE INTERNAL EXPLANATION

Let us concede that the Boring Explanation fails adequately to explain the connection between the mathematical fact and the psychological fact. An adequate explanation must bring out why it is not just an accident that mathematicians tend to accept as axioms only true sentences. We have already conceded that a causal explanation is out of the question. Are there any alternatives? I will now discuss a second minimal explanation of the mathematical reliability claim that purports to be an alternative. This explanation is slightly more informative than the Boring Explanation, and accordingly, it cannot as easily be dismissed.

According to this explanation, mathematicians' tendency to accept as axioms only true sentences is adequately explained by pointing out that the historical process that led to the acceptance of these axioms is a *justifiable* one according to the standards of justification implicit in the mathematical and scientific community.[24] Never mind that mathematicians may not be able to explicitly articulate these standards of justification. What matters is that such standards *exist*. To deny that such standards exist would be to put mathematics at an unfair disadvantage. Such a denial would be particularly unpalatable to a naturalistic philosopher, who seeks to respect successful science. I will call this explanation of the mathematical reliability claim the *Internal Explanation*. This name is appropriate because the explanation is given from the point of view of the science in question: Since this science has developed through a justifiable process, it is not just an accident that mathematicians' acceptance of axioms is a reliable indicator of their truth.

Of course, the Internal Explanation makes no mention of causal contact with a realm of mathematical objects. But so what? Unlike the standards of justification implicit in the empirical sciences, those implicit in mathematics do no require any such contact between the knower and the known. To insist on a requirement of causal contact would therefore be philosophical arrogance. After all, mathematicians know better than we philosophers do when a mathematical proposition is justified. Based on considerations of this sort, Burgess and Rosen argue that a naturalistic philosopher, who seeks to respect successful science, has no right to require more than the Internal Explanation. To require more, they say, would be to ask for some justification higher than what science is in the business of providing. But this would be to go "outside, above, and beyond" natural science, which to a naturalistic philosopher would be illegitimate.[25]

How can the challenger respond to this defense of the Internal Explanation? One option is suggested by Field. Field starts out by being puzzled how mathematicians can have justified belief in entities with which they don't causally interact. He therefore argues that, for mathematicians to be justified, it must be possible to explain the mathematical reliability claim. But the Internal Explanation simply *assumes* that mathematicians' beliefs are justified. So from Field's point of view, the Internal Explanation begs the question.

We have to agree with Field that internal explanations[26] are not very informative. To see this, assume the scientific status of some discipline is contested and that someone therefore demands an explanation of the associated reliability claim. An internal explanation of this reliability claim will do nothing to reassure us. For this explanation simply assumes that the practitioners of the discipline are justified in what they do and uses this to establish that it is not just an accident that the reliability claim is true. So it is built into the very structure of an internal explanation that it will find a connection between the practitioners' beliefs and the subject matter of these beliefs. Internal explanations are therefore deeply dissatisfying.

The problem is how one can require more. Field appears to think that a more informative explanation is possible only if we bracket the justifications that mathematicians themselves offer for their belief in the mathematical axioms. It is because he operates with this very strict restriction that he can claim, with at least some plausibility, that there is no alternative to a causal explanation. But this restriction is excessive. I do not think the reliability claim associated with *any* irreducible branch of science can be explained without employing the sorts of justification that are peculiar to this very science. For instance, the explanation of the reliability of physicists' beliefs sketched at the beginning of last section makes essential use of physical modes of justification. The mathematical platonist can therefore respond that Field places platonistic mathematics at a disadvantage already at the outset. By accepting at face value the justifications that physicists provide for their beliefs but refusing to extend this honor to the justifications that mathematicians provide for theirs, it is Field who begs the question.

So despite its unsatisfactory character, the Internal Explanation remains undefeated.

4. EXTERNAL EXPLANATIONS

I will now describe a kind of explanation that is more demanding than the Internal Explanation but that can still reasonably be required of the mathematical platonist. Assume again that some discipline is contested. As we saw, an internal explanation of the reliability claim associated with this discipline is unsatisfactory. All it does is point out that the relevant sentences are accepted as a result of a justifiable process. Although this may ensure *that* the process is reliable, it does nothing to explain *what makes it the case* that the process is reliable. An explanation that addresses this latter question as well would be much more illuminating. And the demand for such an explanation seems completely reasonable. The scientists in question employ certain methods, based on which they make various claims about reality. So we can carry out a scientific investigation of these scientists, their methods, and the claims they make. First, we describe and analyze their claims. Then, we describe and analyze their methods. And finally, we attempt to explain why these methods are conducive to finding out whether these claims are true. Let's call such an explanation *an external explanation*.

As an example, consider the case of perceptual knowledge. Ordinary people are very good at finding out about middle-size bodies in their immediate vicinity. Most of the claims they make about such bodies are true. A semantic analysis of these claims will show them to be about physical objects outside of people's sensory surfaces. And an investigation of people's methods for deciding whether to accept such claims will show that people mostly rely on the verdicts of their senses. *Prima facie*, it is puzzling that these two systems of facts should be so nicely correlated. However, contemporary science provides at least a good beginning of an explanation of why this is so, having to do with light being reflected from our physical surroundings and impinging on our retinas, and this information's being interpreted by our visual system.

As this example shows, external explanations too have to rely on claims from the contested discipline. To explain the reliability of our perceptual beliefs, we have to appeal to our knowledge of light and of the workings of our perceptual system. And this knowledge is ultimately based on the verdicts of our senses. So external explanations too are in a sense circular and will not, for this reason, satisfy a skeptic. However, our task is not to answer the skeptic but to provide scientific explanations of reliability claims. And for this task, the circularity involved in an external explanation is benign – or so I will now argue.

The problem with internal explanations is that there is so little distance between what is assumed (that some process is justifiable) and what is to be explained (that this process is reliable). To see this, consider an internal explanation of the reliability of our perceptual beliefs. According to this explanation, our perceptual beliefs are reliable because they have been formed through the justifiable process of carefully looking before making up one's mind. But of course, to someone who is puzzled why perception is reliable, this explanation is completely unhelpful. Whatever doubts this person had about perception in the first place will automatically be inherited by this internal explanation. By contrast, external explanations secure a much greater distance between what is assumed and what is explained. It assumes *that* certain claims and methods are reliable. But what it attempts to explain is *why* these claims are reliable, what makes it the case that the methods giving rise to these claims are conducive to determining the truth of such claims. So even if claims from the contested discipline figure among the assumptions, this is no guarantee of success. So this circularity is not trivializing in the way that the circularity of internal explanations is. In particular, when an external explanation is possible, this will increase our confidence in the contested discipline.

Another way of bringing out the benign nature of the circularity involved in external explanations is the following. Call the theory whose reliability is investigated *the object theory*. Call the theory employed in this investigation *the meta-theory*. In an external explanation, the meta-theory need not be identical to the object theory, whereas in an internal explanation, it must be. In particular, if we have two competing theories, say of physics, it is in principle possible to use one to give an external explanation of the reliability of the other. Clearly, this would not be possible with an internal explanation.[27]

Having argued that external explanations are non-trivial, I must now argue that the demand for such an explanation remains neutral. Recall that the problem with the traditional epistemological challenges is that they make unreasonable demands. Benacerraf feels a need to point to a direct causal connection between the knower and the known, and Field demands that the platonist give a causal explanation of the mathematical reliability claim. So both require that the mathematical reliability claim be explained in roughly the same way as empirical reliability claims are explained.[28] But given how different pure mathematics is from the empirical sciences, this requirement is unreasonable. (This is the Natural Response mentioned at the beginning of the paper.) But the demand for an external explanation is not unreasonable in this way. For here nothing is assumed about what resources an external explanation may use. In particular, the mathematical platonist is allowed to assume the reliability and justifiability of all the science he needs, including platonistic mathematics.

In their defense of the Internal Explanation, Burgess and Rosen argue that, in requiring more than an internal explanation, one would have to appeal to considerations "outside, above, and beyond" natural science, thus making this requirement unacceptable to naturalistic philosophers.[29] But this argument is mistaken. Just as it makes perfect naturalistic sense to require more than an internal explanation of perception, the same requirement makes sense of mathematics as well. This requirement is neither contrary to science nor outside its scope; rather, it flows from science itself. We want an account of how the central epistemological notions, such as evidence, reliability, and justification, fit into a scientific account of the epistemic agent and his natural surroundings. And in order to develop such an account, we need to engage in external explanations.

I conclude that external explanations are both substantive and neutral. Thus we have an improved epistemological challenge, based on requiring that it be possible to give an external explanation of mathematics. Although this challenge puts severe constraints on any acceptable philosophy of mathematics, I am hopeful it can be met. In Section 5, I will make some suggestions about how.

5. TOWARDS AN ANSWER TO THE IMPROVED EPISTEMOLOGICAL CHALLENGE?

In Section 1 I presented an argument that we cannot make sense of a counterfactual dependence of people's disposition to accept mathematical sentences upon these sentences' being true. I will now reconsider this argument and suggest that even in mathematics, a certain counterfactual dependence can be made sense of. Using this dependence, I will suggest a way in which the improved epistemological challenge may be met.

The argument from Section 1 went as follows. The mathematical reliability claim says that for each mathematical sentence S, if mathematicians accept S, then S is true. If this correlation is to have counterfactual force, it will have to take the form of a counterfactual dependence of mathematicians' disposition to accept sentences upon these sentences' being true. But in order to make sense of this counterfactual, we would have to make sense of how things would have been had a mathematical sentence S which in fact is true not been true. But since mathematical truths are necessary, we are unable to make sense of contrary-to-fact mathematical scenarios. Hence we are unable to make sense of the requisite counterfactual.

On a closer look, however, this argument turns out to depend on our using a *deflationary* truth-predicate in the formalization of the mathematical reliability. It is only because we use a deflationary truth-predicate that ‘S is true’ comes out cognitively synonymous with S itself. This is why ‘S is true’ must be regarded, *modulo* cognitive synonymy, as a purely mathematical statement. And this in turn is why we cannot make sense of S’s having a different truth-value than the one it actually has. If instead we formalize the mathematical reliability claim by means of a *semantic* truth-predicate, then ‘S is true’ will not be a purely mathematical statement. For on this construal, what ‘S is true’ says is that S has a truth-condition that is satisfied; and this claim is concerned also with the contingent facts in virtue of which S has the truth-condition it happens to have.[30] The truth of a sentence generally depends on two factors: not just on the fact that reality is such as to satisfy its truth-condition, but also on the meta-semantic fact that this sentence has the truth-condition it happens to have.[31] Accordingly, people’s acceptance of a sentence as true generally depends on both the non-semantic and the meta-semantic components of its truth. Had the relevant aspects of reality been different, people might not have accepted the sentence; but equally, had the sentence meant something different, people might not have accepted it.

This suggests that even in mathematics, where we cannot vary the fact that reality is such as to satisfy the relevant truth-conditions, we can still vary the meta-semantic facts involved in mathematical sentences’ having these truth-conditions in the first place and study what effect this would have on people’s disposition to accept sentences as true. So even in mathematics, we can make sense of at least one component of the usual two-dimensional dependence of people’s disposition to accept sentences as true upon the truth of these sentences. The trick is to employ a semantic truth-predicate, not just a deflationary one.

When we do this, it makes perfect sense to ask how things would have been, had the mathematical sentence S had a different truth-value than the one it actually has. For this to have been the case, S would have had to have different truth-conditions. For instance, had the sentence '$2+2=4$' not been true, this might have been because the numeral '4' was used to denote the number 5. Thus, when we employ a semantic truth-predicate, it makes sense to ask whether mathematicians would have been disposed to accept the same mathematical sentences, had these sentences not been true.

This means that we can make sense of a counterfactual dependence of the psychological fact

(1P) $\forall S$ (mathematicians accept S $\rightarrow$ S follows from Σ)

upon a semantic version of the mathematical fact, namely

($1M_s$) $\forall S$ (S follows from $\Sigma \rightarrow$ S is true_s)[32].

For ($1M_s$) serves as a partial truth-definition for the relevant mathematical language. If we replace the formal theory Σ with an alternative theory Σ', this induces a (not necessarily unique) change in the truth-theory for this language. So when we formalize the mathematical reliability claim as

(1_s) $\forall S$ (mathematicians accept S $\rightarrow$ S is true_s)

we can make sense of at least one dimension of counterfactual dependence. Henceforth we will therefore be concerned exclusively with the semantic notion of truth, not with the deflationary notion.

Before considering how far the above considerations take us, I want to address two objections. So far, I have simply assumed we have a semantic truth-predicate at our disposal. But this is denied by so-called deflationists about truth.[33] However, a semantic truth-predicate in my present sense does not require great philosophical resources. All that is required is that there be some sense of the word 'true' on which it makes sense to say: S is actually true, but had S been used differently, S might have been false. In particular, it is not required that this truth-predicate be sharply determined by the facts on which semantic facts supervene. In a situation of semantic indeterminacy, where the semantic truth-predicate is not sharply determined by these facts, the right response is not to reject the reliability claim (1_S) but to observe that the claim has less content than it appears to have because one of its constituent predicates is indeterminate. And the less content the reliability claim has, the smaller will be the associated explanatory burden.[34]

A second and more serious objection is the following. I may have succeeded in making sense of a dependence of mathematicians' acceptance of *sentences* upon the (semantic) truth of these sentences. But it is more important, and much harder, to make sense of the analogous dependence of mathematicians' *beliefs* upon the (semantic) truth of these beliefs.

This is *more important* because an explanation of the latter dependence could easily be extended to one of the former. Assume we had an explanation of why mathematicians believe only true mathematical propositions. Then their disposition to accept only true mathematical sentences could easily be explained. For they have a general desire to accept only sentences that are true. And for each mathematical sentence S, there is a mathematical proposition *p* such that they know: S is true if and only if *p*. Being rational, they will therefore accept S only if they believe *p*. And by assumption, we can explain why their belief that *p* is reliable.

It is *harder* to make sense of the dependence in question at the level of beliefs than at the level of sentences because the contents of belief – propositions – appear to have their truth-conditions by necessity. This prevents us from straightforwardly extending to beliefs the counterfactual dependence I found to obtain between (1P) and ($1M_S$).

With a bit of cleverness, however, I believe it *is* possible to extend this dependence to beliefs.[35] Let's consider a simple example of a belief in a necessary proposition. Assume that Smith sees Venus in the morning and names it 'Phosphorus', and that he sees Venus again in the evening but, not knowing that it is identical to what he saw in the morning, now names it 'Hesperus'. But assume that Smith later (perhaps after a calculation of orbits) comes to believe: Hesperus is identical to Phosphorus. (That is, he forms the belief that *he* would express in the words following the colon.) Although the *content* of this belief may have its truth-condition by necessity, it is contingent that the belief itself (as a psychological state) has this content. The reference of the terms 'Hesperus' and 'Phosphorus' is based on (but need not be exhausted by) Smith's causal interactions with certain spatiotemporal chunks of heavenly bodies. And it is only because these chunks belong to a common planet that these terms are co-referential. Had the chunks belonged to different heavenly bodies, the terms would have referred to different objects. The fact that the chunks belong to a common planet is therefore meta-semantically involved in Smith's belief's having the semantic content it happens to have. But since the fact that the chunks belong to a common planet is perfectly contingent, it is no mystery how Smith became aware of it and as a result came to believe: Hesperus is identical to Phosphorus. Indeed, had the spatiotemporal chunks with which Smith interacted not belonged to a common planet, he would not have formed this belief. So this explains how Smith's belief is responsive to its truth.

Let's now return to where we left off before considering the two objections. We had found that no sense can be made of a dependence of mathematical beliefs upon whatever facts may be non-semantically involved in their truth. However, we've now seen that there is good reason to believe that sense can be made of how such beliefs depend on the facts responsible for these beliefs' having the semantic content they happen to have. How far does this one dimension of counterfactual dependence take us? I would like to close this paper by making a bold proposal: perhaps all facts that are involved in the truth of mathematical beliefs are meta-semantically involved. If so, the one dimension of dependence which we have been able to make sense of will suffice to explain how mathematicians' beliefs are responsive to the truth of these beliefs.

In fact, this proposal isn't quite as radical as it may at first appear. Firstly, the class of facts that are meta-semantically involved in the truth of a belief may be quite large. For instance, in the above example the fact that two spatiotemporal chunks belong to the same planet was meta-semantically involved in Smith's belief's having the semantic content it happened to have, although there is obviously nothing intrinsically semantic about this fact. Secondly, although the belief in question depends entirely on facts that are meta-semantically operative, its content need not be analytic. For instance, the belief that Hesperus is identical to Phosporus is not analytic.

As it stands, my proposal is obviously purely programmatic. And a more substantive development will have to await another paper.[36] But even in its current form the proposal does establish something important, namely that there turns out to be logical space for a response to the improved challenge where no such space appeared to exist.[37]

NOTES

[1] Frege supported the argument by means of his logicism. But that is not obligatory. Later philosophers have supported the argument in other ways; Quine, for instance, by means of his holistic empiricism.

[2] In fact, the causal theories of knowledge were only intended for knowledge of contingent empirical truths; see e.g. Goldman (1967).

[3] See Burgess and Rosen (1997) and Lewis (1986).

[4] I here present the epistemological challenge as a challenge to the possibility of *justified belief* in platonistic mathematics rather than to the possibility of *knowledge* of such matters. As Burgess and Rosen (1997, pp. 36–7) point out, this yields a stronger challenge. Even if the mathematical platonist cannot claim to *know* truths about platonistic mathematics, as long as her mathematical beliefs are *justified*, she will be within her rights to propagate these beliefs. So in order to silence the platonist, the challenger must question the *justification* of her beliefs.

[5] See Field (1989, pp. 25–30 and 230–9). I have modified his argument slightly to incorporate the improvement mentioned in the previous footnote.

[6] Clearly, reliability claims can also be formalized by quantifying over propositions: For almost all p, if mathematicians accept p, then p. I prefer the formalization in terms of sentences since it avoids commitment to sharply individuated propositions, and since it suffices for present purposes. But I discuss the formalization in terms of propositions in Section 5.

[7] This characterization of deflationism follows Field (1994).

[8] This notion of a semantic truth-predicate corresponds roughly to Field's notion of an *inflationary* truth-predicate; other philosophers use the label "robust". I will have more to say about this notion in Section 5.

[9] Field (1989, pp. 228–30).

[10] Field (1989, p. 238).

[11] However, in Section 5 I will argue that a certain counterfactual dependence is possible after all.
[12] When formalizing reliability claims, I will always assume that the variable 'S' has been restricted to the relevant language. Moreover, by using the universal quantifier I idealize slightly. Since mathematicians make mistakes, I should strictly speaking use the quantifier 'for almost all'. Henceforth, these two qualifications will be suppressed.
[13] See Field (1989, p. 232). I have modified the example slightly for improved clarity.
[14] Of course, this is a *mathematical* claim. But recall that the epistemological challenge seeks to establish that *even if mathematical platonism is true*, the beliefs of the mathematical platonist would not be justified. So in the present dialectical situation, the mathematical platonist is allowed to appeal to this claim.
[15] I will speak of "boring explanations" generally and reserve the capitalized form "the Boring Explanation" for the boring explanation described above of the mathematical reliability claim.
[16] This argument is suggested in Lewis (1986, pp. 111–12).
[17] This argument too is suggested in Lewis (1986, pp. 111–12).
[18] See Burgess and Rosen (1997, I.A.2.c).
[19] See Burgess and Rosen (1997, p. 45).
[20] See McGee (1997).
[21] Monadic second-order logic is second-order logic whose only second-order quantifiers are monadic. But since set theory contains a pairing function, this restriction has no significance for the present application.
[22] See Boolos (1984, 1985).
[23] See Resnik (1988), Hazen (1993), and Linnebo (2002).
[24] See Burgess and Rosen (1997, pp. 46–9).
[25] See, in addition to Burgess and Rosen (1997, pp. 46–9), also Maddy (1997, part III).
[26] I will use "internal explanation" as a general term designating explanations of the form discussed, and "the Internal Explanation" for the internal explanation just sketched of the mathematical reliability claim.
[27] This situation is analogous to one that we find in the philosophy of logic. When we have a non-disquotational semantic theory, admitting a certain logical law in the meta-language does not automatically entail admitting it in the object-language. See Dummett (1991).
[28] Recall that this requirement is less absolute for Benacerraf than for Field. For as we saw at the beginning of the paper, all Benacerraf argues is that there is a conflict between the desideratum that our language be given a uniform semantics and the desideratum that this semantics be able to mesh with (what Benacerraf takes to be) a reasonable epistemology.
[29] See Burgess and Rosen (1997, pp. 46–9).

[30] This and following uses of the notion of fact are completely innocent and could, if desired, be replaced by talk about truths or by use of suitable gerunds.
[31] *Semantics proper* studies how the semantic value of a complex expression is determined by its semantic structure and the semantic values of its simple constituents, whereas *meta-semantics* attempts to explain in virtue of what expressions have the semantic structure and the semantic values that our semantic theory ascribes to them. (This distinction coincides with the distinction between "descriptive" and "foundational" semantics in Stalnaker (2001).)
[32] ($1M_s$) is like ($1M_d$) except that it uses a semantic truth-predicate, true_s, rather than the deflationary one, true_d.
[33] See e.g. Field (1994) and Horwich (1998).
[34] In [article name suppressed] I discuss in some detail how the epistemological challenge would be affected by semantic indeterminacy.
[35] For present purposes I need not take a stand on whether it is also *necessary* to study the dependence in question at the level of beliefs rather than at the level of sentences. Although in the remainder of this article I will talk about the dependence as it arises at the level of beliefs, everything I say can be reproduced at the level of sentences.
[36] See Linnebo (forthcoming).
[37] This article is draws heavily on Chapters 2 and 4 of my dissertation Linnebo (2002). For discussion and written comments I am extremely grateful to my dissertation advisors, Warren Goldfarb, Richard Heck, and Charles Parsons, as well as to Matti Eklund, Jim Pryor, Augustin Rayo, Michael Rescorla, an anonymous referee, and the participants in a Harvard-MIT discussion group.

REFERENCES

Benacerraf, P. (1973): 'Mathematical Truth', reprinted in Benacerraf and Putnam (1983).

Benacerraf, P. and Hilary P. (eds.) (1983): *Philosophy of Mathematics. Selected Readings* 2nd edn., Cambridge: Cambridge University Press.

Boolos, G. (1984): 'To Be Is To Be a Value of a Variable (or to Be Some Values of Some Variables)', *Journal of Philosophy* 81, 430–450.

Boolos, G. (1985): 'Nominalist Platonism', *Philosophical Review* 94, 327–344.

Burgess, J. and Gideon R. (1997): *A Subject with No Object. Strategies for Nominalistic Interpretation of Mathematics,* Oxford: Clarendon Press.

Dummett, M. (1991): *The Logical Basis of Metaphysics,* Cambridge, MA.: Harvard University Press.

Field, H. (1989): *Realism, Mathematics, and Modality,* Oxford: Basil Blackwell.

Field, H. (1994): 'Deflationist Views of Meaning and Content', *Mind* 103, 249–285.

Goldman, A. (1967): 'A Causal Theory of Knowing', *Journal of Philosophy* 64, 355–372. Hazen, A.P. (1993): 'Against Pluralism', *Australasian Journal of Philosophy* 81, 132–144.

Horwich, P. (1990/98): *Truth,* Oxford: Basil Blackwell.

Lewis, D.K. (1986): *On the Plurality of Worlds,* Oxford: Basil Blackwell.

Linnebo, Ø. (2002): Science with Numbers: A Naturalistic Defence of Platonism, *Ph.D. Dissertation*, Harvard University.

Linnebo, Ø. (2003): 'Plural Quantification Exposed', *Nous* 37, 71–92.

Linnebo, Ø. Forthcoming. 'Reference and Frege's Context Principle', In *Proceedings of Uppsala Conference on the Philosophy of Mathematics*.

Maddy, P. (1997): *Naturalism in Mathematics,* Oxford: Clarendon.

McGee, V. (1997): 'How We Learn Mathematical Language', *Philosophical Review* 106, 35–68.

Stalnaker, R. (2001): 'On Considering a Possible World as Actual', *Proceedings of the Aristotelian Society* Suppl. 65, 141–56.

Merton College
Oxford
UK
E-mail: oystein.linnebo@philosophy.oxford.ac.uk