

NOTICE

This material may be
protected by copyright
law (Title 17 U.S. Code.)

A New Platonist Epistemology

1. Introduction

In chapter 2, I argued that in order to salvage their view, mathematical platonists have to explain how human beings can acquire knowledge of abstract mathematical objects, given that they are not capable of coming into any sort of *contact* with such objects, that is, receiving any *information* from such objects. In this chapter, I will explain how platonists can do this. In effect, I will be providing an epistemology of abstract objects, although I should say that it will not be a *complete* epistemology. I will only say enough to (a) motivate the claim that spatiotemporal creatures like ourselves *can* acquire knowledge of non-spatiotemporal mathematical objects and (b) provide a rough sketch of how this works. In short, I will provide enough of an epistemology to block the Benacerrafian epistemological objection to platonism, but I will not do any more than this; I will not go into any of the details of the epistemology that are irrelevant to the question of whether Benacerraf's objection can be answered.

2. Skeleton of the Refutation of the Epistemological Argument

What I want to argue is that if we adopt plenitudinous platonism—that is, the view I have been calling FBP¹—we can very easily explain how human beings could acquire knowledge of abstract mathematical objects. My argument here is going to be rather long and complicated, but the intuitive line of thought is quite simple. If FBP is correct, then all consistent purely mathematical theories truly describe some collection of abstract mathematical objects. Thus, to acquire knowledge of mathematical objects, all we need to do is acquire knowledge that some purely mathematical theory is *consistent*. (It doesn't matter how we come up with the theory; some creative mathematician might simply "dream it up".) But knowledge of the consistency of a mathematical theory—or any *other* kind of theory, for

that matter—does not require any sort of contact with, or access to, the objects that the theory is about. Thus, the Benacerrafian objection has been answered: we can acquire knowledge of abstract mathematical objects *without* the aid of any sort of contact with such objects.

I can formulate this as a direct response to an argument of Field's. He once wrote as follows:

But special 'reliability relations' between the mathematical realm and the belief states of mathematicians seem altogether too much to swallow. It is rather as if someone claimed that his or her belief states about the daily happenings in a remote village in Nepal were nearly all disquotationally true, despite the absence of any mechanism to explain the correlation between those belief states and the happenings in the village.²

Now, I admit that I could not have knowledge of a Nepalese village without any access to it. But if all possible Nepalese villages existed, then I *could* have knowledge of these villages, even without any access to them. To attain such knowledge, I would merely have to dream up a possible Nepalese village. For on the assumption that all possible Nepalese villages exist, it would follow that the village I have imagined exists and that my beliefs about this village correspond to the facts about it. Now, of course, it is not the case that all possible Nepalese villages exist, and so we cannot attain knowledge of them in this way. But according to FBP, all possible mathematical objects do exist. Therefore, if we adopt FBP, we can also adopt this sort of epistemology for mathematical objects.

Now, of course, in order to motivate this line of thought, I'm going to have to provide a good deal of supporting argument and also clarify a few key terms, most notably, 'consistent' and 'logically possible'. But before I do any of this, I want to address an objection that may have already occurred to the reader. One might put the worry like this. "Your account of how we could acquire knowledge of mathematical objects seems to assume that we are capable of *thinking about* mathematical objects, or *dreaming up stories about* such objects, or *formulating theories about* them. But it is simply not clear how we could do these things. After all, platonists need to explain not just how we could acquire *knowledge* of mathematical objects, but also how we could do things like have *beliefs* about mathematical objects and *refer* to mathematical objects."

To appreciate my response to this worry, we need to recall the distinction I made in chapter 2 between metaphysically *thin* senses of 'about' and metaphysically *thick* senses of that term. To have a belief that is *thickly* about an object *x*, one must be "connected" to *x* in some appropriate way, whereas to have a belief that is *thinly* about *x*, one needn't be "connected" to it in any non-trivial way. Indeed, on the thin sense of 'about', there needn't even be any such thing as *x*; for instance, we can say that a little girl's belief that Santa Claus is fat is (thinly) about Santa Claus, despite the fact that there is really no such person as Santa Claus. Now, if there are any worries about how human beings could have beliefs about mathematical objects, or how they could dream up stories about such objects, then these worries are surely based upon the thick sense of 'about'. No one doubts that we could formulate beliefs and theories that are thinly about mathe-

BALAGUER, MARK
"PLATONISM AND ANTI-PLATONISM IN MATHEMATICS"
NEW YORK: OXFORD U.P. 1998

mathematical objects. But my account of knowledge of mathematical objects is going to be based solely upon the claim that we can formulate beliefs and theories that are thinly about mathematical objects. I am going to argue that if FBP is true, then we can acquire knowledge of mathematical objects by merely formulating consistent beliefs and theories that are thinly about such objects, because, according to FBP, *all* consistent purely mathematical theories truly describe some collection of mathematical objects.

One might object as follows. "You may be right that if FBP is true, then all consistent purely mathematical theories truly describe *some* collection of mathematical objects, or *some* part of the mathematical realm. But *which* part? How do we know that it will be true of the part of the mathematical realm that its authors intended to characterize? Indeed, it seems mistaken to think that such theories will characterize *unique* parts of the mathematical realm at all. This, of course, is just the point of Benacerraf's other important paper, 'What Numbers Could Not Be': since all consistent purely mathematical theories (including those that are categorical*) have multiple models, it seems that platonists are committed to the thesis that such theories fail to pick out unique collections of mathematical objects."

I am going to discuss this non-uniqueness objection at length in chapter 4. But I want to say just a few words about this here in order to ease the worry that considerations of this sort can be used to block the response that I want to give to the epistemological argument. What I am going to argue in chapter 4 is that non-uniqueness is simply not a problem for platonists; that is, I am going to *embrace* non-uniqueness. Thus, in connection with the epistemological problem, the point is this: if I know that some theory truly describes part of the mathematical realm, then I have knowledge of that realm, regardless of whether it describes a unique part of that realm, and regardless of whether it is "about" some collection of mathematical objects in a metaphysically thick sense of the term. And as for the intentions of the authors, there are two things I want to say. First of all, they too are irrelevant: if a mathematical theory truly describes part of the mathematical realm, then I can attain knowledge of the mathematical realm by studying that theory, regardless of what its authors had in mind. Second of all, I simply deny that the "intentions" that we have in mathematical contexts are anything like the intentions that we have in empirical contexts. In particular, I do not think there are any *unique* collections of objects that correspond to what we have in mind when we formulate our mathematical beliefs and theories; in other words, I do not think that any of our beliefs or theories are "about" any mathematical objects in any metaphysically thick sense of the term.⁵ And I should say here that I do not endorse this view merely because it provides a way of avoiding various objections to platonism; on the contrary, I think this view dovetails with the actual facts about mathematical belief and mathematical knowledge. All of this will be discussed in more detail—and motivated—in chapter 4, where I discuss the non-uniqueness objection to platonism.

In any event, let me now provide a clear and precise statement of the central argument of this chapter. I will do this in Fieldian terms. He writes that the challenge to platonists is to account for the fact that *if mathematicians accept p*,

then p.⁶ I think this is right: what is at issue is not whether we *have* mathematical knowledge, but whether FBP-ists can *account* for this knowledge, that is, whether they can account for how the mathematical knowledge that we *do* have could be knowledge of an inaccessible mathematical realm. Now, of course, platonists do not have to account for there being a *perfect* correlation between our mathematical beliefs and the mathematical facts; this is simply because there isn't any such correlation to account for, because we're human, that is, we make mistakes. What needs to be explained is the fact that our mathematical beliefs are *reliable*, that is, the fact that *usually* (i.e., as a general rule) if mathematicians accept a purely mathematical sentence *p*, then *p* truly describes part of the mathematical realm. Thus, I will simply try to show that FBP-ists *can* account for this fact.

(Actually, I'm going to speak in terms of *theories* rather than sentences. That is, I'm going to argue that FBP-ists can account for the fact that, as a general rule, if mathematicians accept a purely mathematical theory *T*, then *T* truly describes part of the mathematical realm. This just makes the argument more general, because on my usage, a theory is just a collection of sentences, and so a sentence is really just a very simple theory. In addition to saying what a theory is, I should probably also say what a *purely mathematical* theory is. This is just a theory that speaks of nothing but the mathematical realm, that is, does nothing but predicate mathematical properties and mathematical relations of mathematical objects. In contrast to pure mathematical theories, there are also *impure* theories and *mixed* theories, which speak of both mathematical and physical objects.⁷ Now, I am going to concentrate here on pure theories, but of course, to give a *complete* platonist epistemology, one would have to account for our knowledge of impure and mixed theories as well. But in order to solve the Benacerrafian epistemological problem with platonism, we needn't give a complete epistemology; we need only explain how human beings could acquire *some* knowledge of the mathematical realm, and so we can concentrate on pure theories. It is worth noting, however, that my epistemology *can* be generalized to cover impure and mixed theories.⁸)

In any event, my argument proceeds as follows:

- (i) FBP-ists can account for the fact that human beings can—without coming into contact with the mathematical realm—formulate purely mathematical theories.
- (ii) FBP-ists can account for the fact that human beings can—without coming into contact with the mathematical realm—know of many of these purely mathematical theories that they are consistent.⁹
- (iii) If (ii) is true, then FBP-ists can account for the fact that (as a general rule) if mathematicians accept a purely mathematical theory *T*, then *T* is consistent.

Therefore,

- (iv) FBP-ists can account for the fact that (as a general rule) if mathematicians accept a purely mathematical theory *T*, then *T* is consistent.
- (v) If FBP is true, then every consistent purely mathematical theory truly

describes part of the mathematical realm, that is, truly describes some collection of mathematical objects.

Therefore,

(vi) FBP-ists can account for the fact that (as a general rule) if mathematicians accept a purely mathematical theory *T*, then *T* truly describes part of the mathematical realm.

The argument for (i) has already been given: this premise is trivial, because it is not making any strong claim to the effect that our purely mathematical theories have unique domains of mathematical objects that they are “about” in some metaphysically thick sense of the term.

The argument for (ii) will be given in section 5.

Premise (iii) is entirely trivial: (ii) tells us that FBP-ists can account for the fact that we have some skill at distinguishing consistent theories from inconsistent ones; all we have to add to this in order to get (iii) is that mathematicians use this skill in deciding what pure mathematical theories to accept—that is, that mathematicians will accept a theory only if they believe it is consistent—and that FBP-ists can account for this fact. I do not think anyone would question this. (I should note here that I am *not* saying that if mathematicians believe a theory is consistent, then they will automatically accept it. This is certainly not true: mathematicians generally require more than mere consistency before they will accept a theory. All I’m saying is that if mathematicians accept a theory, then it’s (probably) consistent; in other words, the skill alluded to in (ii) makes our acceptance of a theory a somewhat reliable indicator of the consistency of the theory.)

Premise (iv) follows from (ii) and (iii) by *modus ponens*.

Premise (v) is just as trivial as (i) and (iii). FBP says that all the mathematical objects that logically possibly could exist actually do exist. But this means that every consistent purely mathematical theory truly describes some collection of mathematical objects. For if there were some such theory that *didn’t* do this—that is, that spoke of a collection of objects that do *not* exist—we would have a violation of the assumption that all the mathematical objects that logically possibly could exist actually do exist. Now, of course, this argument relies upon the assumption that the notion of *logical possibility* that appears in the definition of FBP is at least as broad as the notion of *consistency* that appears in (v). But FBP-ists can obtain this result by stipulation. After all, FBP is *their* theory, and they can define it however they like. In section 5, I will reveal the exact line I want to take on the terms ‘consistent’ and ‘logically possible’, but for now, let me just say that I am going to take them to be *synonyms*, for this is already enough to deliver the result I need here.¹⁰

Finally, (vi) follows trivially from (iv) and (v). In the present context—that is, in the context of accounting for mathematical knowledge—FBP-ists are allowed to assume that FBP is true. But this, together with (v), gives us that all consistent purely mathematical theories truly describe part of the mathematical realm; but combining this with (iv) gives us (vi). Now, one might wonder *why* FBP-ists are allowed to assume here that FBP is true. There are at least two reasons why this

assumption is legitimate. The first is this: in the present context, FBP-ists are not trying to establish their theory; they are merely trying to account for a certain fact (namely, the fact that we have mathematical knowledge) from *within* their theory; but in general, when one is trying to show that a theory *T* can account for a fact *F*, one can assume that *T* is true and make use of all of its resources.¹¹ The second reason is that (a) in the present context, FBP-ists are merely trying to respond to the epistemological argument against their view, and (b) that argument assumes FBP.¹²

As I have set things up, (vi) is precisely what I need. The Benacerrafian worry is that platonists cannot account for the reliability of our mathematical beliefs, and (vi) simply asserts that FBP-ists *can* account for it. Now, the only real gap I have left in the argument for (vi) is (ii). I will close this gap in section 5. But before I do that, I would like to address two sorts of worries. In section 3, I will address the worry that I haven’t done enough, that is, that (vi) does not really eliminate the epistemological problem with platonism. And in section 4, I will respond to some objections to FBP (and I will argue that there are independent reasons for thinking that FBP is the best version of platonism there is). Now, we have just seen that this is not really necessary, that in the present context, I can legitimately *assume* FBP. But I want to quell the worry that I have solved the epistemological problem with platonism only by adopting an untenable version of platonism.

3. Internalist vs. Externalist Explanations

Consider the following objection to my line of argument. “All you’ve really explained is how it is that human beings could *stumble onto* theories that truly describe the mathematical realm. On the picture you’ve given us, the mathematical community accepts a mathematical theory *T* for a list of reasons, one of them being that *T* is consistent (or more precisely, that mathematicians believe that *T* is consistent). Then, since FBP is true, it turns out that *T* truly describes part of the mathematical realm. But since mathematicians have no conception of FBP, they do not know *why* *T* truly describes part of the mathematical realm, and so the fact that it does is, in some sense, *lucky*. This point can also be put as follows. Let *T* be a purely mathematical theory that we know (or reliably believe) is consistent. (That there *are* such theories is established by (ii).) Then the objection to your epistemology is that you have only an FBP-ist account of

(M₁) our ability to know that *if* FBP is true, *then* *T* truly describes part of the mathematical realm.¹³

You do not have an FBP-ist account of

(M₂) our ability to know that *T* truly describes part of the mathematical realm,

because you have said nothing to account for

(M₃) our ability to know that FBP is true."

The problem with this objection to my epistemology is that (a) it demands an *internalist* account of the reliability of our mathematical beliefs, but (b) in order to meet the Benacerrafian epistemological challenge, platonists need only provide an *externalist* account of the reliability of our mathematical beliefs. To give an externalist account of the reliability of S's beliefs, one merely has to explain why S's methods of belief acquisition are, *in fact*, reliable; but to give an internalist account of the reliability of S's beliefs, one must do more: one must also explain how S knows (or reliably believes) that her methods of belief acquisition are reliable.

My FBP-ist account of the reliability of our mathematical beliefs is externalist: I explain this reliability by pointing out that (a) we use our knowledge of the consistency of purely mathematical theories in fixing our purely mathematical beliefs, and (b) on the assumption that FBP is true, any method of fixing purely mathematical belief that is so constrained by knowledge of consistency is, *in fact*, reliable (that is, any system of purely mathematical beliefs that is consistent will, in fact, truly describe part of the mathematical realm). I do not claim that actual mathematical knowers can justify FBP or even that they have any conception of FBP. Thus, what I need to argue, in order to block the above objection, is that I don't *need* mathematical knowers to have any conception of FBP, that is, that I don't need an internalist account of mathematical knowledge in order to refute the Benacerrafian objection to platonism.

It seems obvious to me that platonists need only an externalist account of mathematical knowledge. We can appreciate this by reflecting on the sort of epistemological challenge that Benacerraf is trying to present and by locating the empirical analog of Benacerraf's challenge, that is, the analogous challenge to our ability to acquire empirical knowledge about ordinary physical objects. According to both Field and Benacerraf—and I think they are right about this—it is *easy* to solve the empirical analog of Benacerraf's challenge: we can do so by merely appealing to sense perception. But this means that Field and Benacerraf are merely demanding an externalist account of mathematical knowledge. For an appeal to sense perception can provide only an externalist account of our empirical knowledge; it cannot provide an internalist account. To see this, let R be some simple theory about the physical world that we could verify via sense perception (e.g., the theory that snow is white). In internalist terms, all we can account for by appealing to sense perception is

(E₁) our ability to know that *if* there is an external world of the sort that gives rise to accurate sense perceptions, *then* R is true.

An appeal to sense perception does *not* yield an internalist account of

(E₂) our ability to know that R is true of the physical world,

because it does nothing to explain

(E₃) our ability to know that there is an external world of the sort referred to in (E₁).

On the other hand, an appeal to sense perception is sufficient for an *externalist* account of (E₂). EWA-ists—that is, those who believe there is an external world of the sort referred to in (E₁)—can give an externalist account of our empirical knowledge of physical objects by merely pointing out that (a) we use sense perception as a means of fixing our beliefs about the physical world, and (b) on the assumption that EWA is true, any method of fixing empirical belief that is so constrained by sense perception is, in fact, reliable. Since this is an externalist account, EWA-ists do not need to claim that actual empirical knowers can justify EWA or even that such knowers have any conception of EWA.

So the FBP-ist's situation with respect to knowledge of mathematical objects seems to be exactly analogous to the EWA-ist's situation with respect to empirical knowledge of physical objects. The FBP-ist can provide an externalist account of our mathematical knowledge that is exactly analogous to the EWA-ist's externalist account of our empirical knowledge: where the EWA-ist appeals to sense perception, the FBP-ist appeals to our ability to separate consistent theories from inconsistent theories; and where the EWA-ist appeals to EWA, the FBP-ist appeals to FBP. Moreover, the FBP-ist's attempt to provide an internalist account of mathematical knowledge and the EWA-ist's attempt to provide an internalist account of empirical knowledge break down at exactly analogous points: the former breaks down in the attempt to account for knowledge that FBP is true, and the latter breaks down in the attempt to account for knowledge that EWA is true.

It seems to me that anti-platonists can block my argument only by finding some sort of relevant disanalogy between the FBP-ist's epistemological situation and the EWA-ist's epistemological situation. They cannot allow the two situations to be analogous, because the whole point of the Benacerrafian objection is to raise a *special* problem for abstract objects, that is, a problem that is easily solvable for physical objects. Now, it *may* be that there is some *other* epistemological problem—for example, one motivated by Cartesian-style skeptical arguments—that applies to both EWA and FBP; but I am not concerned with any such problem here; I am concerned only with the Benacerrafian worry that there is a special epistemological problem with abstract objects.¹⁴

The upshot of all this is that Benacerraf's argument has to be interpreted as demanding an externalist account of our knowledge of mathematical objects. The anti-platonist's claim has to be that while such an account cannot be given, an externalist account of our knowledge of *physical* objects *can* be given. We cannot interpret Benacerraf as demanding an internalist account of our knowledge of mathematical objects, because this is no easier to provide for our knowledge of physical objects. (I take it that this is all entirely obvious and precisely why Benacerraf and Field formulate the demand as a demand for an externalist account of mathematical knowledge.)

The question we need to consider, then, is whether there is any relevant disanalogy between the FBP-ist's externalist account of mathematical knowledge and the EWA-ist's externalist account of empirical knowledge. I will consider two ways in which anti-platonists might try to establish such a disanalogy. The first proceeds as follows. "While it is true that most people who know things about the physical world never cognize EWA, and while it is true that even if they did, they

could not justify their assumption that EWA is true, it seems that, at *some* level, people do accept EWA. But the situation with respect to FBP is entirely different: people just do not assume—at *any* level—that FBP is true.”

First of all, I am not sure that either of the two central claims here is right. I am not sure that people assume—at some level—that EWA is true; and if we decide to say that they do, then I do not see why we shouldn't *also* say that they assume—at some level—that FBP is true. To assume (at some level) that FBP is true is just to assume that our mathematical singular terms refer; but it seems fairly plausible to claim that this assumption is inherent (in some sense and at some level) in mathematical practice. If a mathematician comes up with a radically new pure mathematical theory, she can be criticized on the grounds that the theory is inconsistent or uninteresting or useless, but she cannot be criticized—legitimately, anyway—on the grounds that the objects of the theory do not exist. Now, criticisms of this sort *have* emerged in the history of mathematics—for instance, in connection with imaginary numbers—but, ultimately, they have never had any real effect; that is, they have never blocked the acceptance of an otherwise acceptable theory. I think it is fair to say that at this point in time, it is not a legitimate or interesting mathematical criticism to claim that the objects of a consistent purely mathematical theory do not exist.

But the real problem with the first attempt to establish a disanalogy between FBP and EWA is that it is irrelevant. Given that we need only an externalist account of our knowledge of mathematical objects, it simply doesn't matter whether anyone assumes (at *any* level) that FBP is true.¹⁵ My claim is that people can acquire knowledge of the mathematical realm—even if they do not assume (at *any* level) that FBP is true—by simply having a method of mathematical belief acquisition that (as a general rule) leads them to believe purely mathematical sentences and theories only if they are consistent. This is exactly analogous to the claim that people can acquire knowledge of the physical world—even if they do not assume (at *any* level) that EWA is true—simply by looking at it with a visual apparatus that (as a general rule) depicts the world accurately. And, of course, the *reason* we can acquire knowledge in these ways is that these methods of belief acquisition are, in fact, reliable.

A second way in which anti-platonists might try to establish a disanalogy between the externalist epistemologies of FBP-ists and EWA-ists proceeds as follows. “The FBP-ist is not on all fours with the EWA-ist, because FBP is not analogous to the bare claim that there exists an external physical world. FBP states not just that there *is* an external mathematical world, but that there is a very particular *kind* of mathematical world, namely, a *plenitudinous* one. Because of this, your explanation of knowledge of the mathematical realm is trivial. To see why, consider an analogous explanation. Let ZFP be a version of platonism that takes Zermelo-Fraenkel set theory to be true of part of the mathematical realm. Then ZFP-ists can give an externalist explanation of our knowledge that ZF is true because, on their view, any method of belief acquisition that leads to the acceptance of ZF will be, in fact, reliable.”

The problem with this argument is that it does not establish a disanalogy between FBP and EWA, because EWA is *not* the bare claim that there exists an

external world. It is the claim that there exists an external world *of the sort referred to in (E1)*, that is, the sort that gives rise to accurate sense perceptions, for instance, one containing photons, photon-reflecting objects, eyes, and so on. It seems to me that if anything, this is *farther* from the bare claim that there exists an external world than FBP is from the bare claim that there exists a mathematical realm. Moreover, there is also an empirical analog to the bit about ZFP. Let QMR be a version of realism that takes quantum mechanics to be true. Then QMR-ists can give an externalist explanation of our knowledge that QM is true because, on their view, any method of belief acquisition that leads to the acceptance of QM will be, in fact, reliable.

The externalist epistemologies of the EWA-ist and the FBP-ist are not trivial in the way that the externalist epistemologies of the ZFP-ist and the QMR-ist are. There are at least two reasons for this. I will state these reasons in terms of ZFP and FBP, but exactly analogous points could be made in terms of QMR and EWA. The first reason that the above ZFP-ist epistemology is trivial is that it does not describe a method of mathematical belief acquisition that both leads us to believe ZF *and* is reliable in general. My FBP-ist epistemology, on the other hand, does describe a method of mathematical belief acquisition that is reliable in general. Indeed, it describes a *class* of such methods, namely, the class of methods that forbid the acceptance of inconsistent purely mathematical theories.¹⁶ The second reason that the above ZFP-ist epistemology is trivial, while my FBP-ist epistemology is not, is that ZFP is a mathematical theory, whereas FBP is an ontological theory. (Unlike ZFP, which is essentially equivalent to ZF, FBP makes no claims about any *particular* mathematical objects; it merely asserts a *general* criterion for when we ought to countenance mathematical objects.) The upshot of this is that by adopting FBP, we *explain* our ability to acquire mathematical knowledge, whereas by adopting ZFP, we do no such thing, because here, mathematical knowledge is smuggled in from the start.

I can think of no other way of trying to draw a disanalogy between the epistemologies for FBP and EWA. Thus, I conclude that the two epistemologies are on all fours and, therefore, that my externalist FBP-ist epistemology is sufficient to refute Benacerraf's argument.

Before going on, I want to guard against a possible misunderstanding. My intention in this section is *not* to provide a self-contained refutation of Benacerraf's argument; that is, my point is not that platonists do not need an epistemology for mathematical objects, because we do not have an epistemology for physical objects. On the contrary, I think we do have an epistemology for physical objects, namely, a perception-based externalist epistemology. This epistemology might not do everything we would like it to do, but it surely does a lot. My purpose in this section has, rather, been to argue that the FBP-based externalist epistemology I sketched in section 2 is on equal footing with this perception-based epistemology. It doesn't do everything we would like an epistemology of mathematics to do, but it does do a lot. Indeed, it does just as much as the perception-based epistemology does in the empirical case.

If you doubt that my explanation does a lot, consider that the Benacerrafian argument is supposed to inspire an absolute befuddlement about our ability to

acquire knowledge of the mathematical realm. We find ourselves asking, "How in the world could we have *any clue* about the nature of such an inaccessible realm? How could we even begin to make a *guess* in this connection?" This is decidedly different from what skeptical arguments do to us. It is entirely obvious how people could make correct guesses about the physical world, that is, how they could stumble onto true hypotheses about it: they could do this by merely *looking* at it. All that one might wonder about is our ability to *know* things (in the skeptic's sense) about the physical world. But look what FBP does for us. It explains how rational people can formulate hypotheses that, in fact, truly describe parts of the mathematical realm: they can do this by merely constructing consistent purely mathematical theories. Of course, one might still wonder about our ability to *know* things (in the skeptic's sense) about the mathematical realm, but in the present context, this is irrelevant. For (a) all I am trying to establish here is that my FBP-based epistemology does everything in the mathematical case that the perception-based epistemology does in the empirical case, and (b) the latter does no better against skepticism than the former does.

I began this section with the worry that FBP *only* explains how our mathematical beliefs could turn out to be, in fact, reliable. The response, in a nutshell, is that this is exactly what *needs* to be explained, because the whole force of Benacerraf's argument lies in the fact that it makes us wonder how, if platonism were true, our mathematical beliefs could even be, in fact, reliable.

The only remaining hole in my argument is (ii). I will motivate this premise in section 5. But before I do that, I would like to provide a (partial) defense of FBP. Now, we have already seen that I don't *need* to do this, but I want to say a few words in this connection in order to block the objection that I have solved the epistemological problem with platonism only by adopting an untenable version of platonism.

4. Defending and Motivating FBP

I begin by fending off several different objections to FBP. Then at the end of this section, I argue briefly that FBP is actually the best version of platonism there is, that is, that non-full-blooded (or non-plenitudinous) versions of platonism are untenable.

Objection 1: FBP seems to lead to a *contradiction*. It entails that all consistent purely mathematical theories truly describe part of the mathematical realm, but there are numerous cases in which consistent purely mathematical theories contradict one another. An example is ZFC and $ZF + \sim C$ (that is, Zermelo-Fraenkel set theory with and without the axiom of choice). These theories are both consistent (assuming that ZF is consistent), and so FBP entails that they both truly describe part of the mathematical realm. Thus, FBP seems to lead to the contradictory result that C and $\sim C$ are both true.

Reply: This is not a genuine contradiction. According to FBP, both ZFC and $ZF + \sim C$ truly describe parts of the mathematical realm, but there is nothing wrong with this, because they describe *different* parts of that realm. In other words,

they describe different *kinds* of sets, or different *universes* of sets. Thus, while it does follow from FBP that both C and $\sim C$ truly describe parts of the mathematical realm, we can obtain this result only by interpreting C in two different ways in the two different cases, that is, by assigning different sorts of entities to the expressions of C in the two different cases. Therefore, insofar as ' C and $\sim C$ ' truly describes the mathematical realm, it is no more a genuine contradiction than is the sentence 'Aristotle married Jackie Kennedy and Aristotle did not marry Jackie Kennedy'. (And note that since, in mathematics, we never allow a term to shift meaning within a theory, ' C and $\sim C$ ' will not be a theorem of any of our mathematical theories, except for those that contain an unrelated contradiction.¹⁷)

We might express the idea that ZFC and $ZF + \sim C$ describe different universes of sets by saying that ZFC describes universes of sets_{*c*}, whereas $ZF + \sim C$ describes universes of sets_{*~c*}. Now, it is important to note that according to FBP-ists, ZFC does not describe a *unique* universe of sets_{*c*}; it describes many different universes of sets_{*c*}. For example, it describes some universes in which the continuum hypothesis (CH) is true and others in which it is not true. This is simply because $ZFC + CH$ and $ZFC + \sim CH$ are both consistent and, hence, both truly describe parts of the mathematical realm. In general, the point here is that if FBP is true, then there are as many different kinds of sets as there are consistent set theories.

I should note, however, that the phrase 'different kinds of sets' can be a bit misleading. One way to generate a set theory and a corresponding universe of sets is to relativize the quantifiers of another set theory. In such cases, we will have two different set theories describing different universes of sets, but it seems a bit misleading to suggest that we have two different *kinds* of sets here, because there will be entities that are members of both universes. But I think we can say that we do have two different kinds of sets here and then merely note that one of the kinds is *nested* in the other, or *less inclusive* than the other.

In contrast to the picture of nested universes of sets, there is also a picture of, so to speak, *side-by-side* universes of sets. This might be the best way to visualize the situation with respect to CH. Technically speaking, there are $ZF + CH$ universes nested within $ZF + \sim CH$ universes and $ZF + \sim CH$ universes nested within $ZF + CH$ universes, but it is perhaps best to think of these universes as existing side by side—or more generally, to think of the universes that are characterized by the theories

$ZF + CH$

$ZF +$ 'the size of the continuum is $\aleph_2$ '

$ZF +$ 'the size of the continuum is $\aleph_3$ '

and so on.

as existing side by side.

Objection 2: FBP entails that all consistent purely mathematical theories truly describe parts of the mathematical realm. Thus, it entails that among purely mathematical theories, consistency is sufficient for truth. But this seems to represent a shift in the meaning of the word 'true', as it is used by mathematicians, and so FBP seems to fly in the face of mathematical practice.

Reply: FBP does not entail that among purely mathematical theories, consistency is sufficient for truth. (It does entail that all purely mathematical theories truly describe parts of the mathematical realm, but as we will see, it does not follow from this that all such theories are *true*.) More importantly, FBP does not bring with it a shift in the meaning of the word 'true', as it is used by mathematicians. What mathematicians standardly mean when they say that a sentence is true is that it is true in the *standard model*, or the *intended structure*—or as we'll see, the class of intended structures—for the given branch of mathematics. It seems to me that this fits perfectly well with FBP. Indeed, it seems to me that FBP-ists can explain *why* mathematicians use 'true' to mean 'true in the standard model'.

Let me begin my argument for these claims by pointing out that talk of *truth in a model* dovetails with FBP. Models are just parts of the mathematical realm, so to say that a sentence *S* is true in a model *M* is just to say that *S* is true of some particular part of the mathematical realm. But if talk of truth in a model dovetails with FBP, then talk of truth in the standard model dovetails with FBP as well. The only thing that FBP-ists will want to emphasize in this connection is that there is nothing *metaphysically special* about standard models. Now, this is not to say that there can be no good reason for singling out a model (or class of models) as standard; it simply means that such models do not enjoy any privileged ontological status. The claim that a model (or class of models) is standard is a claim about *us* rather than the model; what is being claimed is that this is the model (or class of models) that is *intended*, that is, that we *have in mind* with respect to the given theory. Thus, for instance, a model of set theory is standard if and only if it jibes with *our notion of set*;¹⁸ and a model of arithmetic is standard if and only if it jibes with *our conception of the natural numbers*; and so on. (There are also cases, I think, in which we want to say that a model is standard because it is *inclusive*, but I think that whenever this is true, the inclusive model is also what we *have in mind*.)

These remarks show that FBP-ists can account for talk of truth in the standard model. But FBP-ists can also account for why 'true in the standard model (or models)' is more or less synonymous with 'true'. This might seem surprising, for FBP-ists maintain that all consistent purely mathematical theories truly describe some collection of mathematical objects, but they do not claim that all such theories are true in a standard model. But FBP-ists can make sense of this by maintaining that not all theories that truly describe parts of the mathematical realm are true. And they can do this by appealing to a certain, fairly standard, way of thinking about truth that distinguishes the notion of *truth*, or *truth simpliciter*, from the notion of *truth in a language L* (where a language is an abstract object that, at the very least, maps sentence types onto truth conditions). The fact that every consistent purely mathematical theory truly describes a collection of mathematical objects shows that every such theory is true in some language *L*. (Actually, it shows more than this; it shows that every such theory is true in a language *L* that interprets the given theory in a "natural way", that is, a way that takes the given theory to be about the objects that, intuitively, it is about.) But none of this shows that all consistent purely mathematical theories are true simpliciter. Indeed, insofar as many of these theories have never been tokened, the notion of truth

simpliciter does not even make sense in connection with them. For on the present view, the notion of truth simpliciter is defined only for sentence tokens (and collections of sentence tokens) and not for sentence types. In particular, and very roughly, a sentence token is *true simpliciter*, on this view, if and only if it is true in the *intended language* (or rather, if and only if it is a token of a type *t* that is true in the intended language). This is obviously very rough, but given the present view of what a language is, it is clear that we have to define 'true simpliciter' in some such way, because every sentence type is true in some languages and false in others. For instance, the sentence type 'Snow is white' is true in English but false in all languages that map it onto the truth condition of grass being orange.

In any event, it should be clear that if this view of the notion of truth simpliciter is at least roughly correct, then the notion of truth in the standard model is a notion of truth simpliciter. For as we've seen, to say that a model is standard is just to say that it's *intended*. But to say that a sentence is true in the intended model is essentially equivalent to saying that it is true in the intended *interpretation*, and so it's also more or less equivalent to saying that it is true in the intended *language*. Thus, I conclude that by appealing to the above view of the notion of truth simpliciter, FBP-ists can account for why mathematicians use 'true' to mean 'true in the standard model'.

It is important to note here that, according to FBP, none of this provides any *metaphysical* distinction to the mathematical theories that happen to be true simpliciter. Take any consistent purely mathematical theory *T*. If mathematicians became interested in the structure that *T* describes and formulated *T* in an effort to describe that structure, then according to FBP, *T* would be true simpliciter. But it doesn't follow from this that *T* is true simpliciter *right now*. Thus, whether a mathematical theory is true simpliciter depends partially upon facts about us. But this is just what we want. For in general, whether our utterances are true depends partially upon our intentions, upon what we intend these utterances to mean. For instance, part of the reason that our utterances of 'Snow is white' are true is that when we utter this sentence, we intend to be saying that snow is white, as opposed to, say, that grass is orange. (Of course, this isn't the *whole* reason that our utterances of 'Snow is white' are true; part of the reason is that snow is white.) In any event, my point here is that, according to FBP, mathematical truth works in an essentially the same way as ordinary truth.

(One might object here as follows: "Given the above remarks, it's not clear that FBP succeeds in marking any important difference between mathematics and empirical science; after all, every consistent physical theory is true in some language *L*." My response is that by maintaining that the mathematical realm is plenitudinous, FBP-ists obtain the result that every consistent purely mathematical theory is true in a language *L* that *interprets the theory in a natural way*. This guarantees that whenever mathematicians think of a mathematical structure and formulate a theory that characterizes that structure, the theory will be true. But of course, the corresponding claim about physical reality is false. It is simply not the case that whenever somebody dreams up a physical situation and formulates a theory that characterizes it, the theory is true. If this were the case, then all logically consistent novels would be literally true stories.)

Before going on, I should point out that some of the remarks in the last few paragraphs might be a bit misleading, because they might lead one to think that FBP-ists are committed to the claim that, for each branch of mathematics, there is a *unique* intended model. As we will soon see, this is not true. To take the case of set theory, FBP-ists allow that it *may* be that there are multiple models of set theory that are not isomorphic to one another and that are all perfectly consistent with all of our set-theoretic intentions and, indeed, with the totality of all of our thoughts about sets. Thus, to leave room for this possibility, FBP-ists maintain that a mathematical sentence is *true simpliciter*, or *correct*, if and only if it is true in *all* of the standard models for the given branch of mathematics; and it is *incorrect* if and only if it is false in all of these models; and if it is true in some of these models and false in others, then it is neither correct nor incorrect. I will discuss this in more detail in my reply to objection 3. We will see there that this fits very well with mathematical practice.

(In response to the arguments of this section, one might wonder how human beings could succeed in mentally picking out a particular model, or class of models, as standard. I will address this worry in my reply to objection 4.)

Objection 3: FBP seems to sacrifice the *objectivity* of mathematics. Now, it does entail that mathematical theories are objectively true in the sense that they are true of an objective mathematical realm and, hence, true independently of us and our mathematical theorizing. Nonetheless, it seems that FBP-ists cannot salvage the objectivity of certain open questions. For instance, FBP seems to entail that undecidable sentences like CH do not have determinate truth values. Once it has been established that CH and $\sim$ CH are both consistent with the set theories that we currently accept, all we can say is that CH is true of some kinds of sets and false of others; we cannot maintain that there is any interesting mathematical question left to answer. But this flies in the face of mathematical practice, because there are a good many set theorists who think that CH does have a determinate truth value, that is, that there is an objectively and uniquely correct answer to the question 'How big is the continuum?' (Moreover, if one were inclined to agree with Kreisel that it is the *objectivity* of mathematics, and not its *ontology*, that is the really important issue, one might be inclined to conclude that FBP is actually not a very platonistic view.)^{19,20}

Reply: The claim that FBP-ists cannot salvage the objectivity of undecidable open questions is simply false. Most mathematical disputes can be interpreted as disputes about what is true in the standard model (or models). Consider, for instance, arguments over the truth or falsity of CH. When people argue about whether some axiom candidate that's supposed to settle the CH question is true, what they are really arguing about is whether the given axiom candidate is inherent in *our notion of set*. In other words, they're arguing about whether the axiom candidate is true in the standard model (or class of models) of set theory. Thus, FBP-ists claim that CH and $\sim$ CH are both true in various set-theoretic hierarchies and that arguments over the truth of CH are arguments over its truth value in the intended hierarchy (or hierarchies).

Now, it *may* be that *our notion of set* is non-categorical, that is, that there are numerous models of set theory that are not isomorphic to one another but are,

nonetheless, standard—or *as* standard as any other model. In other words, it may be that the totality of our set thoughts fails to pick out a unique model (or more precisely, a unique class of mutually isomorphic models). If this is the case, then for *some* open set-theoretic questions, there is no objectively correct answer. FBP-ists can easily account for this: to say that there is no objectively correct answer to, for instance, the CH question, is just to say that our notion of set isn't strong enough to settle that question, that is, that neither CH nor $\sim$ CH is inherent in our notion of set, that is, that CH is true in some standard models of set theory and false in others. It seems to me that this is an extremely important point, because traditional platonists *cannot* account for the existence of open questions without correct answers. If there is only *one* universe of sets, then CH is either true or false. But this is a problem for traditional platonism, because one of the dominant opinions about the CH question among contemporary set theorists—if not *the* dominant opinion—is that it doesn't have an objectively correct answer. The problem is that traditional platonists cannot account for how this could be so.

Thus, far from providing an *objection* to FBP, considerations involving objectivity and methodology actually provide us with a reason to *favor* FBP over traditional platonism. For FBP-ists can account for *more* of mathematical practice in this connection than traditional platonists can. In particular, they can account for the existence of undecidable open questions with objectively and uniquely correct answers *and* undecidable open questions *without* objectively correct answers. Most philosophies of mathematics *dictate* that we take one stance or the other here with respect to *all* open questions. But FBP allows mathematicians to say whatever they *want* to say in this connection with respect to each different open question. This, I think, is an extremely appealing feature of FBP. I say this for two reasons. First, a good philosophy of mathematics should not dictate things like this to mathematicians; the point of the philosophy of mathematics is to *interpret* mathematical practice, not to place metaphysically based *restrictions* on it. And second, the FBP-ist stance here just seems intuitively pleasing. It just seems right to say that it *may* be that some open mathematical questions have objectively correct answers whereas others do not. Moreover, it seems very plausible to suppose that what determines whether a given open question has an objectively correct answer is whether it is independent not just of the currently accepted theory in the given area of mathematics, but also of what we *have in mind* with respect to this theory, that is, our notion of set, or our conception of the natural numbers, or whatever. If multiple answers to an open question are consistent with our intentions and concepts and intuitions, then different answers to the question will be true in different standard models, and so the question will not have a unique, objectively correct answer. But if one answer to an open question is, in some sense, already contained in our intentions or concepts or intuitions, then the question does have a unique, objectively correct answer.²¹

Now, actually, this picture of things is a bit oversimplified. In the first place, there are surely going to be cases in which there is no clear fact of the matter as to whether a given answer to some open question is "inherent in our concepts". For insofar as our concepts can be vague and fuzzy, there can be answers to open

questions that are “borderline inherent in these concepts”. And in the second place, even if we assume that our concepts are always precise and well-defined, it is not the case that the *only* way an answer to an open question can be correct is if it is, in some sense, inherent in these concepts. For even if no answer to a given open question is inherent in our concepts, it may be that good reasons—for instance, aesthetic or pragmatic reasons—can be given for *revising* (or perhaps, *refining*) our concepts in a way that would settle the question. And in some cases of this sort, we might still want to say that the question had a correct answer. (Cases like this suggest that there is a relationship of back-and-forth influence between theory and intuition. It is obvious that our theories are influenced by our intuitive pre-theoretic concepts, but cases of this sort suggest that our concepts are also influenced by our theorizing. And this, in turn, suggests that which models count as standard is influenced by our theorizing as well.)

One might wonder what my FBP-ists would say about Euclid’s fifth postulate, for we have here an undecidable proposition that seems a bit different from CH. What I would say is this: we know in exactly which kind of geometrical space this postulate is true and in which kinds of spaces it is false; moreover, we know which of these spaces corresponds to our pre-theoretic intuitions, and so there is simply no mystery here at all. That is, there is no mathematical question left to answer. The only interesting question is whether *physical* space is Euclidean, and that is a non-mathematical, empirical question.

As for open *arithmetical* questions, I think we can safely say that *all* of these have unique, objectively correct answers. If *Q* is a question about the natural numbers, then even if it were shown that all answers to *Q* were undecidable in a theory as strong as ZFC, we would still maintain that *Q* has an objectively correct answer, because we’re convinced that *our conception of the natural numbers* is categorical. That is, we’re convinced that the totality of our natural-number thoughts picks out a unique model, or at worst, a unique class of models that are all isomorphic to one another. This, at any rate, is what *mathematicians* would say. But there is at least one philosopher—namely, Putnam²²—who has argued that our concept of number *isn’t* categorical. There is no reason to go into this here, though, because even if Putnam is right about this, it is not a problem for FBP because, again, that view is compatible with the claim that our mathematical notions and conceptions—such as our notion of set and our conception of the natural numbers—are non-categorical. (I will say a bit more about Putnam’s argument, and the point I’m making here, at the end of chapter 4.)

I want to make one more point before going on. I have claimed that it may be that our notion of set is non-categorical. But I want to emphasize that FBP-ists are not *committed* to this. If it turns out that our notion of set picks out a unique universe of sets (or at least does this up to isomorphism), that will not be a problem for FBP, for as the above remarks already make clear, FBP is neutral as to whether or not this is the case.

Objection 4: The appeal to standard models in the reply to objections 2 and 3 seems to give rise to an epistemic problem for FBP-ists about how human beings could acquire *knowledge* of what the various standard models are like.

Reply: In fact, there is no epistemic problem here at all. This is simply because standard models aren’t metaphysically special. They’re only *sociologically* special,

or *psychologically* special. To ask whether some proposition is true in, for example, the standard model (or class of models) of set theory is just to ask whether it is inherent in *our* notion of set. Thus, since our notion of set is clearly accessible to us, questions about what is true in the standard model (or models) of set theory are clearly within our epistemic reach.

Thus, the answer to the question ‘How do we know what the various standard models of mathematics are like?’ is just this: we formulate axioms that are intuitively pleasing (that is, that jibe with our notion of set, or number, or whatever) and then we prove theorems. This, of course, is exactly true to mathematical practice. Mathematicians try to settle open questions by constructing proofs that rely only upon currently accepted propositions, that is, propositions that we already believe hold in the standard model. And if a question *cannot* be answered in this way—that is, if the propositions that answer the question are undecidable in the current theory—then mathematicians seek new axioms that are (a) powerful enough to entail an answer to the question and (b) intuitively pleasing. (Or if they can’t find any axioms that are intuitively pleasing, they might use an axiom that is pragmatically appealing; I mentioned this above, and I’ll say more about it below.) This is exactly what set theorists have tried to do in connection with CH, or at any rate, it’s what has been attempted by those set theorists who think that CH has a determinate truth value.

The worry behind objection 4 can also be put in this way: “How could human beings mentally pick out a unique model (or class of models) to call standard?” My response to this way of putting the point is similar to my response to the other way of putting it. We do not do anything special here that involves some sort of *connection* between our heads and standard models. We just have our intuitions and notions and conceptions, and we slowly build theories out of them. Since these theories are consistent, it follows (on the assumption that FBP is true) that these theories truly describe parts of the mathematical realm. But some of our theories aren’t categorical, that is, they have multiple models that aren’t isomorphic to one another. Moreover, in some such cases, we’re inclined to say that one of the models of the given theory (or one class of these models) is *standard*, whereas the others are not. But in such cases, standardness is determined not by any sort of *contact* that we have with the mathematical realm, but by the simple fact that our intuitions, notions, and conceptions happen to jibe more with one of the models (or class of models) than any of the others. Our intuitions, notions, and conceptions “pick out” one of the models (or class of models), but they do this only in a *thin* way; they are not “about” the standard model (or class of models) in any *thick* way. (I will elaborate on this picture, and justify it more thoroughly, in chapter 4.)

Objection 5: FBP seems to forbid us to speak of *all* sets. For it seems that according to FBP, every set theory is about a restricted universe of sets. But there doesn’t seem to be any good reason for this; we ought to be able to develop a theory of *all* sets and say whether CH is true in this theory.

Reply: In fact, FBP *doesn’t* forbid the development of such a universal set theory. If somebody came up with a set theory—call it ST—such that (a) each of the axioms of ST seemed intuitively plausible, that is, seemed to jibe with our notion of set, and (b) ST settled all important open questions of set theory and,

indeed, picked out a *unique* universe of sets that seemed intuitively to correspond to our notion of set, then it would be reasonable to maintain that ST was a theory about *all* sets and that it told us once and for all whether CH is true. Now, there might still be other consistent theories, incompatible with ST, that purported to be about sets, and of course, FBP would entail that these other theories truly described parts of the mathematical realm; but we could simply maintain that these other theories weren't about *the universe of sets*, that is, that all the models of these other theories either contained some things that weren't really sets, or didn't contain some things that *were* sets, or both.

Now, at this point, one might be inclined to raise an objection that is, in a sense, the *opposite* of objection 5. In particular, one might argue as follows. "FBP-ists speak of various universes of sets. But we can surely *amalgamate* all these universes to form a single universe of sets. But if this is true, then the various open questions of set theory ought to be taken as being uniquely about this particular universe. Thus, it seems that FBP-ists are committed to the claim that open set-theoretic questions, such as the CH question, have unique, objectively correct answers. But this is problematic, for as we've seen, FBP-ists want to claim that their theory allows that it *might* be that some such questions don't have objectively correct answers."

The problem with this objection is that any FBP-ist who denies that CH has a determinate truth value (i.e., that CH has the same truth value in all of the standard models of set theory) will also deny that there is a unique amalgamated universe that clearly contains all of the things that legitimately count as sets and none of the things that don't. Moreover, if there is a unique amalgamated universe of this sort, then this shows not just that FBP-ists ought to say that CH has a determinate truth value, but that *everyone* ought to say this. In other words, it shows that CH does have a determinate truth value. The important point to note here is that by *itself*, FBP is *neutral* with respect to the question of whether there is a unique amalgamated universe that contains all and only things that legitimately count as sets; that is, it's neutral as to whether there is a unique universe that corresponds to our notion of set. Therefore, it's also neutral as to whether all of the open questions of set theory have unique, objectively correct answers.

So neither objection 5 nor its "opposite" succeeds: FBP doesn't entail that there *isn't* a theory of all and only sets (or a unique *universe* of all and only sets), and it doesn't entail that there is. It remains neutral here, so that mathematicians can settle this question however they want to.

Objection 6: Let's take objection 5 one step farther. FBP seems to prohibit us from making claims about the *entire mathematical realm*.

Reply: Again, this is just false. The sentence 'All mathematical objects are abstract objects' is about the entire mathematical realm, and it is presumably true. Now, it may be that there is nothing *mathematically interesting* to say about the entire mathematical realm, but that's just because the mathematical realm is so vast.

Objection 7: This reply misses the point of objection 6. The problem is that there are *some* things that we might want to say about the entire mathematical realm such that FBP prohibits us from saying *them*. For instance, suppose that

Scottie, a mathematical lunatic, believes that there is no such thing as the number 7, that the natural-number sequence goes straight from 6 to 8. If he utters the sentence 'There is no number 7', then since this is *consistent*, FBP entails that what Scottie said is true of part of the mathematical realm. But this seems to get things wrong. Scottie was talking about the *entire* mathematical realm, and so what he said is surely false.

Reply: FBP tells us that the sentence 'There is no number 7' truly describes *part* of the mathematical realm, but it *doesn't* tell us that it truly describes the *entire* mathematical realm. Thus, there is simply no problem here: if Scottie was talking about the entire mathematical realm when he uttered this sentence, then he was simply wrong, and FBP-ists can say that he was wrong. Why can't they? They never said that all consistent purely mathematical sentences and theories truly describe the *entire* mathematical realm.

Objection 8: Let's think a bit more about sentences like 'There is no number 7', or ' $2 + 2 = 5$ '. FBP entails that such sentences are true of part of the mathematical realm. But this seems wrong; it seems that sentences like these are just false. Now, of course, you might claim that what we *mean* when we say that sentences like ' $2 + 2 = 5$ ' are false is that they are false in the *standard model*. But it's not clear that this is acceptable. After all, we seem to think that these sentences are false in some *absolute* sense.

Reply: FBP-ists can account for the intuition we have that sentences like ' $2 + 2 = 5$ ' are false in some absolute sense. We could construct a consistent purely mathematical theory in which ' $2 + 2 = 5$ ' was a theorem, but to do this, we would have to use at least one of the terms in this sentence in a non-standard way. For instance, we might simply use the symbol '5' to denote the number 4, or we might use '+' to express some unusual operation. But if we interpret the terms of this sentence in a non-standard way, then it would not really say that $2 + 2 = 5$. As long as we interpret ' $2 + 2 = 5$ ' in the standard way, that is, according to *English*, it will be false. Now, of course, this is just to say that ' $2 + 2 = 5$ ' is false in the standard model, but I think this way of putting the point explains why the fact that ' $2 + 2 = 5$ ' is false in the standard model leads to the intuition that it is false absolutely.

Similar remarks can be made with respect to the sentence 'There is no number 7'. I could construct a consistent purely mathematical theory that had this sentence as a theorem. Indeed, I could construct *many* such theories, and according to FBP, they would all truly describe parts of the mathematical realm. For example, one such theory would characterize a part of the mathematical realm that is just like the natural-number sequence except that it has a sort of *hole* in it where 7 is in the natural-number sequence. But, of course, this weird quasi-sequence²³ is not the natural-number sequence; thus, while we can use the sentence 'There is no number 7' to say something true about this quasi-sequence, in doing this, we would be using the sentence in a non-standard way, and so we would not really be saying that there is no number 7. So long as 'There is no number 7' is taken to mean that there is no number 7, it is, according to FBP-ists, false.

Objection 9: We normally think that in order for a person S to know that p,

there has to be a counterfactual relationship between S's belief that *p* and the fact that *p*, so that if things would have been different (that is, if it wouldn't have been the case that *p*), then S would have believed differently (that is, he or she wouldn't have believed that *p*). But the FBP-ist epistemology described here does not salvage this.

Reply: This objection is based on a worry that can be lumped together with the original Benacerrafian worry that human beings could not acquire knowledge of abstract mathematical objects, because they do not have any information-gathering *contact* with such objects. Both worries arise from taking epistemic principles that seem applicable (to at least a limited extent) in empirical contexts and applying them in mathematical contexts. Now, *prima facie*, it might seem that these principles *are* applicable in mathematical contexts—that is, it might seem that we couldn't have knowledge of mathematical objects without having any contact with such objects and without there being a counterfactual relationship here of the above sort—but one of the main points of the present chapter is precisely that these principles are *not* applicable in mathematical contexts. That is, the arguments that I'm developing in this chapter suggest that human beings could have knowledge of mathematical objects even if they don't have any contact with such objects and even if there isn't a counterfactual relationship here of the above sort. Thus, since we never had any good *argument* for thinking that these principles *are* applicable in mathematical contexts,²⁴ we can conclude that the present chapter provides a good response to our two worries here, that is, the worry behind objection 9 and the original Benacerrafian worry.

Objection 10: FBP-ists claim that our mathematical singular terms refer to mathematical objects. But *which* objects do they refer to? It seems pretty clear that if FBP is true, then our mathematical theories do not describe *unique* parts of the mathematical realm. Thus, it seems that singular terms like '3' and 'the null set' do not pick out unique objects. But on the standard view of reference, this is just to say that they suffer a kind of reference *failure*.

Reply: This objection is deeply related to the Benacerrafian non-uniqueness objection to platonism. Thus, since chapter 4 is going to be entirely devoted to providing an FBP-ist reply to the non-uniqueness objection, I want to put the present objection on hold for now and respond to it in chapter 4. To tip my hand a bit, though, my strategy will be to *embrace* the thesis that our mathematical singular terms do not uniquely refer. I should say, however, that the non-uniqueness I embrace will be very *limited*. It might seem that FBP-ists could be forced into the result that all mathematical singular terms refer (non-uniquely) to all mathematical objects, but we'll see that, in fact, they cannot be forced into this or, indeed, anything like it.

In addition to fending off the objections to FBP, I would also like to argue in its *favor*; that is, I would like to argue that it is the best version of platonism there is. The most important argument here is the one I have been developing over the last two chapters: FBP survives the Benacerrafian epistemological attack, whereas non-full-blooded, or non-plenitudinous, versions of platonism do not. If non-plenitudinous platonism were correct, it would be a mystery how we could ever know what the universe of sets was like. Since, in FBP-ist terms, there are infinitely many different universes of sets, traditional platonists have to allow that

the universe of sets could correspond to any of these FBP-ist universes. But given all these possibilities, it's totally unclear how spatiotemporal creatures like ourselves could discover the nature of *the* universe of sets.²⁵

But there are also *independent* reasons for favoring FBP over other sorts of platonism. Indeed, I already gave one such reason in connection with objection 3: FBP-ists can account for the existence of open questions with correct answers *and* open questions without correct answers, whereas traditional platonists can account only for the former.

A second independent argument for FBP is that it reconciles the objectivity of mathematics (to which all platonists are committed) with the legitimacy in mathematics of pragmatic modes of justification, that is, with the fact that the adoption of a new axiom for a mathematical theory can be justified pragmatically, for instance, because it solves certain open questions or simplifies the theory. We saw in chapter 2 (sub-section 5.3) that non-full-blooded platonists cannot account for the legitimacy of pragmatic modes of justification: since, according to their view, an axiom candidate (that has been shown to be independent of the other axioms of the given theory) could very easily turn out to be false of the given domain of objects, it is unclear why pragmatic modes of justification should be legitimate. Why should the fruitfulness of a claim about *the* universe of sets have anything to do with its truth?²⁶ FBP-ists, on the other hand, can easily explain the legitimacy of pragmatic modes of justification. Suppose that a sentence *A* is an axiom candidate for a purely mathematical theory *T* and that *A* is independent of the other axioms of *T*. If we assume that FBP is true, then even if *A* can be justified only pragmatically—that is, even if we do not have any intuition as to whether or not *A* is true—there is nothing wrong with adopting *A*, or *T*+*A*, because according to FBP, there do exist objects for which *T*+*A* holds. Of course, there are also objects for which *T*+~*A* holds, but that is irrelevant. The point is that the decision to adopt *T*+*A* can simply be seen as a decision to study a certain kind of object (or perhaps, as a decision to refine our concepts in a certain way).

A third (related) advantage of FBP is that it reconciles the objectivity of mathematics with the extreme *freedom* that mathematicians have.²⁷ As I have already pointed out, mathematicians cannot be legitimately criticized on the grounds that the objects of their (consistent and pure) theories do not exist. Indeed, just the opposite seems true: one way for a mathematician to become famous is to develop an interesting theory about a kind of mathematical entity or structure of which no one has yet conceived. (Now, of course, a physicist could also become famous in this way, but before we would accept the new physical theory, we would demand independent evidence that the objects in question exist.)

5. Consistency

It remains only to justify premise (ii) of my argument. To this end, I need to argue that FBP-ists can account for the fact that human beings can—without coming into contact with the mathematical realm—know of certain purely mathematical

theories that they are consistent. Let me begin my argument here by stating a reason one might give for being *skeptical* of premise (ii). One might reason as follows.

"Look, there are two different notions of consistency: a theory *T* is *semantically consistent* (or *satisfiable*) iff it has a model, and it is *syntactically consistent* iff there is no derivation of a contradiction from *T* in any logically sound derivation system. But insofar as models and derivations are abstract objects, these are both platonistic notions. Thus, knowledge of consistency is going to be knowledge of abstract objects, namely, models and derivations. Indeed, Gödel has shown that knowledge of syntactic consistency is essentially equivalent to *arithmetical* knowledge. Thus, FBP-ists have not accomplished *anything* by reducing the question of how we could know that our mathematical theories are true to the question of how we could know that they are consistent."

FBP-ists can avoid this worry by merely claiming that the notion of consistency at work in (ii)—and, more generally, in (i)–(vi)—is an anti-platonist notion of consistency. One anti-platonist notion of consistency that FBP-ists can use here is suggested in the work of Kreisel and discussed recently by Field.²⁸ The main idea here is that 'consistent' is simply a *primitive* term. More precisely, the claim is that in addition to the syntactic and semantic notions of consistency, there is also a primitive or intuitive notion of consistency that is not defined in any platonistic way. Now, the standard view here is that the semantic notion of consistency can be thought of as a definition (or perhaps a reductive analysis) of our intuitive notion of consistency, but according to the Kreisel-Field view, this is wrong. On this view, the intuitive notion is related to the two formal notions in analogous ways: neither of the formal notions provides us with a *definition* of the primitive notion, but they both provide us with information about the *extension* of the primitive notion. More specifically, it follows from the definitions of the two formal notions, and from our intuitive understanding of the primitive notion, that (a) if a theory *T* is semantically consistent, then it is intuitively consistent; and (b) if *T* is syntactically inconsistent, then it is intuitively inconsistent. Moreover, if we combine these two points with the completeness theorem—or more precisely, with the Henkin theorem that (among first-order theories) syntactic consistency implies semantic consistency—we arrive at the result that (among first-order theories) the intuitive notion of consistency is coextensive with both formal notions of consistency.

Advocates of the Kreisel-Field view might want to claim that the primitive notion of consistency is equivalent to a primitive notion of *possibility*. Now, there are, of course, many different *kinds* of possibility, but this fits perfectly well with the present view. To see this, consider the fact that for each different kind of possibility, we can define formal notions of syntactic and semantic consistency. Here are two examples:

A theory *T* is *semantically conceptually consistent* iff the union *T*+*C* of *T* and the set *C* of all conceptual truths has a model; and *T* is *syntactically conceptually consistent* iff there is no derivation of a contradiction from *T*+*C* in any logically sound derivation system.

A theory *T* is *semantically physically consistent* iff the union *T*+*P* of *T* and the set *P* of all physical laws has a model; and *T* is *syntactically physically consistent* iff there is no derivation of a contradiction from *T*+*P* in any logically sound derivation system.

Now, given this, we can say that there is an intuitive notion of possibility, or consistency, corresponding to each such pair of formal notions. Thus, the Kreisel-Field intuitive notion is simply the *broadest* of these notions; it is a notion of *logical possibility*.²⁹ All of the other intuitive notions of possibility can be defined in terms of the Kreisel-Field intuitive notion. For example, a theory *T* is *intuitively physically possible* iff *T*+*P* is intuitively possible, or consistent, in the Kreisel-Field sense.

At any rate, in order to use all of this to motivate premise (ii), I need to argue two points. First, I need to argue that it is acceptable for me to use the Kreisel-Field primitive notion of consistency here, that is, that this notion is a genuinely *anti-platonist* notion. And second, I need to argue that if we insert this notion of consistency into premise (ii), then that premise is true; that is, I need to argue that FBP-ists can account for the fact that human beings can—without coming into contact with the mathematical realm—know of certain purely mathematical theories that they are intuitively consistent, in the sense developed by Kreisel and Field.

Now, actually, the first thesis here is stronger than what I really need to establish. All I really need here is that there is *some* legitimate anti-platonist account of consistency. Whether it is the account provided by the Kreisel-Field view is irrelevant; for so long as there is some legitimate anti-platonist account of consistency, I will be able to run the argument in (i)–(vi) and simply understand the word 'consistent' there in the given anti-platonist way. But it seems to me that the claim that there is *some* legitimate anti-platonist account of consistency is extremely weak. If there were *no* legitimate anti-platonist account of consistency, then anti-platonists wouldn't even be able to account for the simple fact that some of our theories are consistent, and so their view would be totally unacceptable. Thus, all of the arguments that I am considering in this book would be moot, because we would know that platonism was correct.³⁰

But in any event, it is not hard to motivate the first thesis of the paragraph before last, that is, the thesis that the Kreisel-Field primitive notion of consistency is a genuinely anti-platonist notion. For since that notion is a *primitive* notion, it is entirely obvious that it isn't defined in terms of abstract objects, because it doesn't have any definition at all.

Now, I suppose that one might try to argue here that we ought not to think of our intuitive notion of consistency as a primitive notion, that we ought to take the semantic notion as providing a definition (or at least a reduction) of the intuitive notion. But I have no idea how one might proceed in trying to argue this point. Indeed, it seems to me that there is no reason to think that the intuitive notion is even *coextensive* with the semantic notion. Henkin's proof shows that they are coextensive in connection with first-order theories, but that argument doesn't extend to higher-order theories.³¹ Moreover, even if we assume, for the

Way too quick

sake of argument, that these two notions are coextensive, there are good reasons for thinking that the semantic notion doesn't provide a definition (or reduction) of the intuitive notion. Field argues this point very convincingly by showing that the semantic notion doesn't capture the "essence" of the intuitive notion. He does this by pointing out that there are certain theories for which it is *obvious* that they are intuitively consistent but *not* obvious that they are semantically consistent. For instance, the theory *S* consisting of all the truths about sets that are storable in the language of set theory is obviously consistent in the intuitive sense, but it is not at all obvious that *S* is semantically consistent, that is, that it has a model. For intuitively, it seems that a model of *S* would have the set of all sets as its universe, but we know that there is no such thing as the set of all sets.³² Now, if the language of *S* is *first-order*, then by Henkin's theorem, *S* does have a model. But (a) the model produced by this proof is extremely unnatural; and (b) this result doesn't extend to cases where the language of *S* is higher-order; and most important, (c) the mere fact that this result is non-trivial, that it has to be *proven*, shows that the semantic notion doesn't capture the "essence" of the intuitive notion.

One might object as follows. "Even if the semantic notion doesn't capture the 'essence' of the intuitive notion, and indeed, even if it's not coextensive with the intuitive notion, there is still something illegitimate about FBP-ists claiming that we use platonistic notions to help us get a grip on the extension of the intuitive notion of consistency. For according to FBP-ists, mathematical knowledge is supposed to arise out of knowledge that certain sentences and theories are intuitively consistent. But given this, how can they also claim that knowledge of the extension of the intuitive notion of consistency is obtained, or partially obtained, by means of an appeal to the platonistic notions of syntactic and semantic consistency? Isn't this circular?" This objection would be cogent only if it were impossible to know that a sentence or theory was intuitively consistent without having some knowledge involving the syntactic or semantic notion of consistency. But this is clearly false: anyone who has taught an introductory logic course can attest that students can be pretty reliable judges of whether a set of sentences is consistent, even if they have no conception whatsoever of syntactic or semantic consistency. Thus, the idea here is that before we developed the notions of syntactic and semantic consistency, our knowledge of intuitive consistency was good enough to give rise to some mathematical knowledge. Then, once the "ball of mathematical knowledge was rolling", so to speak, we developed the formal notions of consistency, acquired some knowledge of them, and in this way, increased our knowledge of the extension of the intuitive notion of consistency. (It should be noted that this picture of things doesn't just refute the above objection—it also jibes with the historical facts.)

In any event, I now want to move on. Assuming that there is a legitimate anti-platonist notion of consistency, what I need to argue in order to motivate premise (ii) is this: FBP-ists can account for the fact that human beings can—without coming into contact with the mathematical realm—know of certain purely mathematical theories that they are consistent, where 'consistent' is understood anti-platonistically. It seems to me, however, that as soon as we realize that we are

working with an anti-platonist notion of consistency here, this premise becomes pretty trivial. If the fact that a sentence or theory is consistent is an anti-platonistic fact, then that fact isn't about any abstract objects, and so it would seem that we don't need any contact with any abstract objects in order to know that the fact obtains, that is, that the given sentence or theory is consistent.

Now, one might object to this on the grounds that sentences and theories are abstract objects. But FBP-ists can simply restrict their attention here to concrete *tokens* of sentences and theories (where it is understood that a "token of a mathematical theory" includes only tokens of the *axioms* of the theory—or more precisely, the axiom schemata and the axioms that aren't instances of axiom schemata). If FBP-ists can explain how human beings could know that concrete tokens of our mathematical theories are consistent, that is all they need to do, because they can claim that these tokens truly describe parts of the mathematical realm in the same way that types do.

A second objection that one might raise here is this. "If *T* is a purely mathematical theory, then although the claim that concrete tokens of *T* are consistent isn't about abstract objects, *T* itself is about abstract objects; thus, it may be that we need contact with abstract objects in order to know that *T* is consistent, because it may be that we need contact with *T*'s own ontology in order to know that *T* is consistent." But this worry is misguided. For in general, knowledge of the consistency of a set of sentences—whether the sentences are purely mathematical, purely physical, mixed, or whatever—does *not* require any sort of epistemic access to, or contact with, the objects that the sentences are about. For instance, I do not need any access to the seventh child born in 1991 in order to know that the sentences asserting it to be female and Italian are consistent with each other; likewise, I don't need any access to this child to know that the sentences asserting it to be male and not male are *inconsistent* with each other. And the same is true of mathematical sentences: I do not need any access to the number 4 in order to know that '4 is even' and '4 is positive' are consistent with each other, or that '4 is odd' and '4 is not odd' are inconsistent with each other.

I take it that this point is entirely obvious, but it is worth mentioning because it provides a clear picture of the intuitive idea behind my epistemology. If FBP is true, then knowledge of the mathematical realm falls straight out of knowledge of the consistency of mathematical theories. But knowledge of the consistency of a theory does not require any contact with the objects of that theory, and so the Benacerrafian lack-of-contact worry about platonism has vanished completely.

Now, I suppose that one might press me here to actually explain how human beings arrive at knowledge of consistency. But it would be entirely inappropriate to dive into this here, because this issue is irrelevant to the question of whether platonism is true. There may be a deep and important question about how we come to know that various sentences and theories are consistent, but it is no more pressing for FBP-ists than it is for anyone else. Everyone has to account for knowledge of consistency, and FBP-ists can accept any explanation here that anyone else can accept. In particular, they can accept any explanation that anti-platonists can accept, because consistency—or rather, the sort of consistency we're discussing here—is an anti-platonist notion. Moreover, the question of how we acquire

knowledge of consistency is no more pressing in connection with our mathematical theories than with our empirical theories. But this means that platonists do not have to address this issue in order to respond to the Benacerrafian epistemological argument, for as I've already pointed out, if that argument is to succeed, it cannot work equally well against physical and mathematical objects. It can succeed only if it shows that there is a *special* epistemological problem with platonism that arises as a result of the inaccessibility of mathematical objects.

In the early sections of this chapter, I tried to reduce the question of how we could acquire knowledge of abstract mathematical objects to the question of how we could acquire knowledge that some of our mathematical sentences and theories are consistent. In this section, I have been trying to argue that this latter sort of knowledge is not problematic. That is, I have been trying to argue that platonists will not encounter any special problems in accounting for our knowledge of mathematical consistency. I think I can provide additional support for this claim by pointing out that knowledge of consistency is *logical* knowledge.³³ We wouldn't expect platonists to have any more trouble than anti-platonists in accounting for our logical knowledge because, intuitively, logical knowledge doesn't seem to be platonistic knowledge at all. This is perhaps the central point of this chapter: if FBP is true, then mathematical knowledge can arise directly out of logical knowledge. Platonists who do not endorse FBP cannot make this claim, because they have to account for how people could know which of our consistent purely mathematical theories truly describe mathematical objects and which do not, and this could not be logical knowledge. But FBP-ists do not have to account for this sort of knowledge, because according to them, *all* of our consistent purely mathematical theories truly describe mathematical objects.³⁴

All of this should be reminiscent of Field's view. He argues that anti-platonists can take mathematical knowledge to be logical knowledge.³⁵ If I'm right, then FBP-ists can do the same thing, although I think the point is better put by saying that mathematical knowledge *can arise directly out of* logical knowledge. And it is worth noting that this doesn't commit FBP-ists to logicism any more than Field's view commits *him* to logicism. Mathematical truth is not logical truth, because the existence claims of mathematics are not logically true. More precisely, if T is a purely mathematical theory implying the existence of various mathematical objects, then (a) if FBP is true, then knowledge of the logical fact that T is consistent can give rise to knowledge of the mathematical fact that T truly describes part of the mathematical realm; but (b) 'T is consistent' is not *equivalent* to 'T truly describes part of the mathematical realm', because whereas the latter can be true only if there are mathematical objects, the former can be true even if there are no mathematical objects. What we can say, however, is that 'T truly describes part of the mathematical realm' follows from the conjunction of FBP and 'T is consistent'.

I end by tying up a loose end. In chapter 1, I said that FBP is (roughly) the view that all the mathematical objects that logically possibly could exist actually do exist. But I noted there that more needed to be said about exactly what is meant here by 'logically possible'. In this section, I have cleared this up: 'logically possible' just means 'consistent, in the primitive, intuitive sense'. (Or at any rate,

this is the stance that FBP-ists should tentatively endorse. If it turns out that there is an anti-platonist account of consistency that is superior to the view that 'consistent' is a primitive term, then FBP-ists might want to change what they say here.) Of course, this means that 'logically possible' can be applied only to whole sentences (or collections of sentences, or sentential components of sentences), and so the above definition of FBP will have to be reworded; but as I showed in chapter 1, this is not a serious problem.

tuition. But this is wrong; as we will see, the construction of representations of numbers no more requires contact with numbers than does the formation of beliefs about numbers.

46. The exception to this is Parsons. He seems wholly unconcerned with the worry about reliability, or at any rate, he seems to think that the worry is somehow trivial: at one point (1994, p. 142) he asserts without argument that “intuition of an object involves intuitive knowledge of that object”. There is a sense in which this makes Parsons the purest advocate of the NCTI strategy, because he does not fall back upon any other strategy of responding to the Benacerrafian challenge. But it also makes his view the most puzzling of the NCTI-platonist views, because it is totally unclear how an appeal to NCTI can help if it is not conjoined with an explanation of reliability.

47. It is obvious, I think, that people do use the word ‘about’ in this thin way. We say, for instance, “The novel *Oliver Twist* is about an orphan who . . .”, but in saying this, we certainly don’t mean to commit to the existence of such an orphan. Now, I suppose that one might wonder what the *truth conditions* of sentences like

(THIN) The sentence *S* is thinly about the object *n*

are, but I don’t think there’s any real mystery here. I haven’t taken the time to think of all the scenarios in which (THIN) would be true, but the most obvious is that *S* contains a singular term that’s coreferential with, or that’s “supposed to be” coreferential with, the singular term ‘*n*’ that appears in (THIN). (The point of the ‘supposed to be’ clause here is to allow for cases in which the two singular terms—that is, the one in *S* and the one in (THIN)—are *vacuous*. Thus, for instance, ‘Santa Claus’ and ‘Kris Kringle’ are supposed to be coreferential, but ‘Santa Claus’ and ‘Oliver Twist’ are not supposed to be coreferential.)

48. I will explain how platonists can solve the problem of belief in chapter 4.

49. There are two ways in which NCTI could be made stronger and, hence, more controversial. First, if NCTI were understood as the view that we are capable of arriving at no-contact intuitions and beliefs that are *thickly* about mathematical objects, then it would be quite controversial (and indeed, as I’ll argue in chapter 4, false). And second, if one were to go beyond the *broad* formulation of NCTI in the text and endorse a specific view of the inner workings of the faculty of no-contact intuition, then the resulting theory might be very controversial.

50. See, e.g., Quine (1951), section 6; Steiner (1975), chapter 4, especially section IV; and Resnik (1997), chapter 7.

51. One might worry here that mathematical theories are often developed *in connection* with their applications and not *before* them. But all my argument requires is that at least *once* in the history of mathematics, people had knowledge of an unapplied mathematical theory. It is this that I am taking to be obvious.

52. All of this accords very well with the practice of mathematics and empirical science. Mathematicians do not think that their theories are confirmed by applications, and empirical scientists do not take themselves to be confirming the mathematical theories that they use. In fact, the opposite is true: when empirical scientists use a mathematical theory, they think that we already know it’s true. This, of course, is why they feel entitled to use it.

53. Katz (1981), p. 207. I should note that according to Katz, the appeal to necessity is not the only way in which mathematical claims can be justified. Once our intuitions have been systematized into a theory, we can justify the theory by using all of the standard criteria for theory evaluation; see Katz (1981), pp. 14–16, 212–214. But while Katz may well be right here, this is not relevant to Benacerraf’s challenge. The reason is identical to the

reason I cited in connection with the Quine-Steiner-Resnik epistemology: what we have here is an *after-the-fact* justification, whereas what platonists need in order to respond to Benacerraf is an account of how the intuitions that engendered the theory could have been accurate in the first place.

54. Lewis (1986), p. 111.

55. Field (1989), p. 238.

56. I should point out here that in opposition to this, Hale and Wright (1992) have argued that the existence of mathematical objects is conceptually necessary. But Field (1993) has argued convincingly that their argument is flawed. The bottom line here is that it just seems wrongheaded to suggest that existence is built into the very *concept* of a mathematical object. Of course, there is much more to be said in this connection, but there is no need to pursue it here.

57. See Resnik (1997), chapter 11, section 3; and Shapiro (1997), chapter 4, section 7.

58. Resnik (1982, p. 101) says: “a pure [mathematical] theory can be falsified by showing that it fails to characterize any pattern at all, that is, that it is inconsistent.” It seems to me that the ‘that is’ here suggests that he thinks that the *only* purely mathematical theories that *don’t* characterize patterns are the inconsistent ones. Shapiro doesn’t come as close as Resnik does to explicitly stating the thesis, but it seems to me that the thesis is at least “suggested” by the following passage:

It is conceivable, barely, that arithmetic is incoherent, in which case *no* structure is characterized. . . . But it is nonsense to claim that the theory of arithmetic does successfully refer to a single, fixed structure . . . but says hardly anything true about it. . . . On our view, the language characterizes or determines a structure . . . if it characterizes anything at all.

(This passage comes from Shapiro (1997), p. 131, but I should note that Shapiro says the same thing—almost word for word—on pp. 167–168 of his (1989).)

59. I should say, however, that I do not think that the appeal to pattern recognition is philosophically important anyway. This was the point of sub-section 6.2: nothing philosophically important can be accomplished by providing an account of the psychological mechanisms responsible for the genesis of our no-contact mathematical beliefs and intuitions. We might clear up some *psychological* mystery in this way, but we cannot clear up any *philosophical* mystery concerning the inaccessibility of mathematical objects.

Chapter 3

1. FBP was defined in chapter 1, sub-section 2.1. Roughly, it is the view that all the mathematical objects that logically *possibly could* exist actually *do* exist.

2. Field (1989), pp. 26–27.

3. See Benacerraf (1965).

4. A theory is *categorical* if and only if all of its models are isomorphic to one another.

5. Here’s another way to put this point: on my view, human beings do not have *de re* mathematical beliefs or *de re* mathematical knowledge. All mathematical knowledge is knowledge by description, and what’s more, it is almost certain that none of our mathematical theories yields a uniquely satisfied description.

6. See Field (1989), p. 26.

7. The difference between impure theories and mixed theories is that the former are *mathematical* theories that refer to physical objects, whereas the latter are *empirical* theories that refer to mathematical objects. Examples of the two sorts of theories are, respectively, impure set theory and quantum mechanics. One might also speak of an *applied* mathe-

mathematical theory. This is a pure theory together with a claim that the theory in question is satisfied by some group of physical objects. An example here would be Riemannian geometry taken as a theory about *physical* space.

8. This might seem surprising, because such theories can be consistent but false. But all I need to do in order to accommodate such theories in my epistemology is shift from talk of consistent theories to talk of theories that do not imply any falsehoods about the physical world. I thank Hartry Field for pointing this out to me. (Note that this shift is irrelevant in the pure case because, here, the theories that imply no falsehoods about the physical world just *are* the consistent theories. This is because inconsistent pure theories imply everything about the physical world, whereas consistent pure theories imply nothing about the physical world; that is, consistent pure theories are *conservative*, in Field's (1989, p. 58) sense of the term.)

9. Note the scope of the existential quantifier here. I am not merely saying that FBP-ists can account for our knowledge that there exist many consistent purely mathematical theories. Rather, I am saying that there exist many mathematical theories such that FBP-ists can account for the fact that we can know that *these theories* are consistent.

10. It might seem that we could provide additional motivation for (v) by appealing to the Henkin theorem that all syntactically consistent first-order theories have models. But I think this is wrong. For whereas the Henkin theorem provides very *unnatural* models for many of our mathematical theories, FBP entails that all consistent purely mathematical theories truly describe parts of the mathematical realm *that they correspond to in a very natural way*, that is, that they are intuitively and straightforwardly (albeit thinly) *about*. Moreover, it is also worth pointing out here that whereas the Henkin theorem applies only to first-order theories, (v) is concerned with higher-order theories as well as first-order theories. (In connection with this last point, one might wonder what I would say about higher-order theories that are syntactically consistent but not semantically consistent. I will discuss this in section 5 when I describe the notion of consistency at work in (v). As we will see, this notion is distinct from both the syntactic and the semantic notions of consistency, although among first-order theories, it is coextensive with both of them.)

11. A bit less generally, the point is that in constructing an epistemology for a theory T, it is legitimate to assume that T is true.

12. The epistemological argument is a *reductio ad absurdum*: the claim here is that *if platonism were true*, then mathematical knowledge would be impossible. Thus, the version of the argument that's aimed at FBP holds that *if FBP were true*, then mathematical knowledge would be impossible. Therefore, to respond to this argument, FBP-ists have to show that *if FBP is true*, then mathematical knowledge is *not* impossible. (Moreover, I don't see how anti-platonists could alter the argument so that it no longer assumes platonism. For the whole point of the epistemological argument is to talk about the inaccessibility of mathematical objects. Thus, to talk about mathematical objects in this way, anti-platonists need to assume, for the sake of argument, that there *are* such objects and then discharge this assumption at the end of the argument.)

13. The FBP-ist account of (M₁) is simple: we can learn what FBP says and recognize that if FBP is true, then *any* theory like T (i.e., any consistent purely mathematical theory) truly describes part of the mathematical realm.

14. Katz has made a similar point. He says (1981, p. 212) that "Empirical knowledge ... has no advantage over a priori knowledge in encounters with the skeptic."

15. Actually, even if we were trying to provide an *internalist* account, it still wouldn't matter whether any actual knowers assume (at any level) that FBP is true. For in order to solve the Benacerrafian problem, platonists don't need to explain how human beings *actually* acquire knowledge of the mathematical realm. They only need to explain how

human beings *could* do this (or how they *plausibly* could do this—see chapter 2, note 37).

16. One such method tells us to believe *all* consistent purely mathematical theories. Others demand that a theory satisfy other conditions, in addition to consistency. But if FBP is true, then *all* of these methods will be reliable.

17. This last point is important, because it enables FBP-ists to block the worry that 'C and not-C' truly describes the mathematical realm but isn't satisfiable. If we are working within a formal mathematical theory, 'C' and 'not-C' will contradict one another, and so 'C and not-C' will be false. The only way that 'C and not-C' can be taken as truly describing the mathematical realm is if we are in an informal extra-mathematical setting in which the two occurrences of 'C' are read differently. But in this informal setting, there is nothing wrong with saying that 'C and not-C' is informally satisfiable, that is, that it can be given an informal interpretation on which it comes out true, namely, one that interprets the two occurrences of 'C' in different ways. Since we are only speaking of interpretation and satisfiability only *informally* here, there is nothing wrong with this. After all, the sentence 'Aristotle married Jackie Kennedy and Aristotle didn't marry Jackie Kennedy' can also be given an informal interpretation on which it comes out true.

18. It *may* be that, *strictly speaking*, set theory doesn't have any standard *models*. For (a) it may be that in order for a model to count as a standard model of set theory, it would need to have the set of all sets as its domain, and (b) there is no such thing as the set of all sets. This is not a *problem*, however, because all of the claims that I am going to make here about standard *models* of set theory could be reformulated and turned into claims about intended set-theoretic *hierarchies*, or intended *parts of the mathematical realm*. But for rhetorical reasons, I am not going to do this; I am going to be speaking simultaneously of all the standard models for the various different branches of mathematics, and it would simply be too cumbersome to keep adding parenthetical provisos to allow for the above complication about set theory. Thus, since there isn't any real problem here, I am simply going to ignore this and write in terms of standard models. I should note, however, that I will return to this point in section 5, for it will become important there, albeit very briefly.

19. I should note here that Field (forthcoming) has made a point similar to the one in the text. He claims that while FBP is ontologically similar to traditional platonism, it is more like fictionalism in connection with objectivity and methodology. But Field does not see this as an *objection* to FBP.

20. A related objection that one might raise here is this: "FBP isn't a genuinely platonistic view, because (a) it tells us, more or less, that in connection with pure mathematics, existence and truth amount to nothing more than consistency, and (b) this is a very anti-platonistic thing to say." But this version of the objection can be answered very easily, because FBP simply doesn't say that "existence and truth amount to nothing more than consistency". Rather, it says that all the mathematical objects that logically possibly could exist actually do exist, and then it *follows* from this that all consistent purely mathematical theories truly describe some collection of actually existing mathematical objects. Moreover, as I pointed out in my reply to objection 2, it doesn't follow from this that all consistent purely mathematical theories are *true*.

21. One might try to argue here that FBP-ists are already committed to the thesis that CH has a determinate truth value—and, hence, that they cannot adopt a neutral stance here of the sort that I am suggesting—because (a) we can amalgamate all the universes of sets to form a *single* universe, and (b) FBP-ists ought to understand CH as being about this particular universe. I will respond to this worry below, in connection with objection 5.

22. Putnam (1980).

23. It isn't a *real* sequence, because real sequences don't have holes in them. If it were

a real sequence, then we would simply be using '8' to denote 7, '9' to denote 8, and so on. In the theory I'm imagining, we use '8' to denote 8—or rather, quasi-8, that is, the analog of 8 in this structure. I say this because sentences like ' $4 + 4 = 8$ ' are true in this structure. What's weird about this structure—or one of the things that's weird about it—is that its numbers (or rather, quasi-numbers) are not closed under successor or addition or subtraction. Terms like 'the successor of 6' and ' $4 + 3$ ' and ' $19 - 12$ ' are undefined.

One might respond as follows: "You do have an ordinary sequence of objects here; you've just got some very odd relations and operations defined on it." I would reply: I admit that I've got a denumerable domain, but the weird relations and operations give me a non-sequential, or quasi-sequential, structure.

24. This is a point that *all* platonists—or at any rate, all platonists who maintain that human beings do not have any information-gathering contact with mathematical objects—would make in response to objection 9. For all such platonists would want to respond to this objection by rejecting all principles that suggest mathematical knowledge requires a counterfactual relationship of the above sort.

25. I argued this point in much more detail in chapter 2 by explaining how the various platonist epistemologies not based upon FBP—for instance, those of Gödel, Maddy, and Quine—fail.

26. One might think that non-full-blooded platonists could object to this argument on the grounds that pragmatic considerations are relevant to empirical theory construction as well as to mathematical theory construction. But I responded to this objection in chapter 2.

27. Cantor (1883, p. 564) wrote that "The very essence of mathematics is its freedom."

28. See Kreisel (1967) and Field (1991). I should note that Kreisel's discussion is actually couched in terms of the notion of logical validity and not consistency; but everything he says applies, *mutatis mutandis*, to consistency. Moreover, Kreisel was not concerned with the fact that his notion of validity was anti-platonistic; indeed, he used platonistic language (in particular, talk of *structures*) in describing the notion. But as is clear from Field's discussion and my discussion here, the view can be developed without using any platonistic language to describe the notion of consistency (or the notion of validity).

29. That this is really the broadest intuitive notion of possibility can be appreciated by noting that the formal notions of syntactic and semantic consistency to which this primitive notion corresponds cannot be made any broader. Or more precisely, the point is that if we try to broaden these two notions, they will no longer be notions of *consistency*.

30. Of course, there would still be a mystery about how human beings could acquire knowledge of abstract mathematical objects—the epistemology that I am proposing here would, I suppose, have to be abandoned along with anti-platonism—but we would know that that mystery was solvable.

31. Let Σ be the set of all higher-order theories that are syntactically consistent but semantically inconsistent. Then the question of whether the intuitive notion of consistency is coextensive with the semantic notion reduces to the question of whether every theory in Σ is intuitively inconsistent. Now, *some* of the theories in Σ do seem to be intuitively inconsistent. For instance, if P is a standard second-order axiomatic theory of arithmetic and G is the Gödel sentence for P , then it seems natural to suppose that $P + \neg G$ is intuitively inconsistent, because it seems natural to suppose that it *couldn't possibly be true*. But there is no reason to think that *all* the members of Σ are intuitively inconsistent. Indeed, as the remarks in the text will show, it is quite possible that there are theories in Σ that are *obviously* consistent in the intuitive sense. These theories are in Σ because of "technical difficulties", because there are certain limitations to our model-theoretic apparatus that

prevent these theories from having models, despite the fact that they are intuitively consistent, that is, despite the fact that, intuitively, they *could be true*.

32. It might seem that this creates a problem for me, for I have been assuming here that set theory has models and, indeed, standard models. But I addressed this worry in note 18.

33. One might complain in this connection that our reasons for believing that, say, ZFC is consistent are not logical reasons. But I didn't mean to suggest that whenever we believe that a mathematical theory is consistent, our *reasons* are always logical; I simply meant that in such cases, the *fact* about which we have knowledge—that is, the fact that the given theory is consistent—is a logical fact. (But in any event, there are *some* mathematical theories that are so simple that we *can* acquire knowledge of their consistency via direct logical intuition and reasoning. And if the arguments of this chapter are cogent, then these theories can give us knowledge of abstract mathematical objects.)

34. One might think that FBP-ists have to account for how people could *know* that all of our consistent purely mathematical theories truly describe mathematical objects. But I responded to this worry in section 3.

35. In particular, Field (1989, essay 3) thinks that we have two different kinds of mathematical knowledge, namely, knowledge that certain axiom systems are *consistent* and knowledge that certain theorems *follow* from certain axiom systems. Thus, my platonist epistemology is very similar to his anti-platonist epistemology.

Chapter 4

1. The thesis that Benacerraf was explicitly attacking was the thesis that arithmetic is about objects. I won't go into this, but the present chapter provides a defense of this thesis, as well as a defense of platonism, because the view defended against Benacerrafian considerations here is a version of object-platonism.

2. See Resnik (1980), p. 231.

3. One might worry that a similar argument could be constructed in an *empirical* context. But such an argument would not go through. For we have *perceptual contact* with physical objects, and this serves to ground our concrete singular terms to *particular* objects in a way that our abstract singular terms are *not* tied to particular objects. Now, I am not suggesting here that a mere appeal to perceptibility eliminates all worries about uniqueness in connection with physical-object talk. To name just one problem here, the appeal to perceptibility won't do anything to block Quine-style inscrutability-of-reference worries. My point is simply that the appeal to perceptibility blocks *the worry in the text* from applying in physical-object cases. To see that this is so, consider an attempt to apply this worry in such a case. One might begin here by, for example, telling us to imagine a Twin-Earth version of Bill Clinton who is just like our Bill Clinton in all ways that we've ever considered but different in some way that we've never thought about. Now, I think it's dubious that such a creature could really be just like our Bill Clinton in *all* ways that we've ever considered, but even if we assume that this is possible, it should be clear that our term 'Bill Clinton' doesn't refer to any such creature, because no such creature is appropriately "connected" to us.

4. I am going to defend a version of platonism that embraces non-uniqueness, and so it might seem as though my response is *not* consistent with the falsity of (2) and (3). But my real point is going to be that platonists can *allow* non-uniqueness, not that they have to *commit* to it.

5. See, e.g., Resnik (1981) and (1997) and Shapiro (1989) and (1997). I discuss platonistic structuralism in chapter 1, sub-section 2.1.