Also available from Bloomsbury

The Bloomsbury Companion to Analytic Philosophy, edited by Barry Dainton and Howard Robinson

The Bloomsbury Companion to Aristotle, edited by Claudia Baracchi

The Bloomsbury Companion to Berkeley, edited by Richard Brook and Bertil Belfrage

The Bloomsbury Companion to Hume, edited by Alan Bailey and Dan O'Brien

The Bloomsbury Companion to Kant, edited by Gary Banham, Dennis Schulting and Nigel Hems

The Bloomsbury Companion to Locke, edited by S.-J. Savonius-Wroth, Paul Schuurman and Jonathan Walmsley

The Bloomsbury Companion to Philosophy of Science, edited by Steven French and Juha Saatsi

The Bloomsbury Companion to Plato, edited by Gerald A. Press

an historical introduction to the philosophy of mathematics

a reader

Edited by Russell Marcus and Mark McEvoy

Bloomsbury Academic
An imprint of Bloomsbury Publishing Plc

B L O O M S B U R Y
LONDON • OXFORD • NEW YORK • NEW DELHI • SYDNEY

there will be no such thing as (e.g.) the axiom of choice; there will only be the axiom of choice for a given universe.[5] Standard models are standard only because they are the intended interpretations. This, coupled with Balaguer's remark that standard semantics is "false but useful"[6] might be taken to give up standard semantics, rather than capture it.

We close by noting that Balaguer's overall view is actually more complicated than we have made it out to be. Though he argues that the main objections to platonism can be overcome by embracing FBP, he also argues that the main objections to fictionalism can be overcome. Finally, since he sees FBP and fictionalism as the only defensible versions of platonism and anti-platonism respectively, he concludes both that we can have no good reason for believing or disbelieving in mathematical objects, and that there is no fact of the matter as to whether such entities exist.

Themes to explore

1 According to Balaguer, what is the standard argument for mathematical platonism?

2 Balaguer's article presents Full-Blooded Platonism as the platonist's best means of solving the Benacerraf knowledge problem. Explain in outline how FBP solves this problem.

3 Here is one misguided objection to FBP:

> FBP implies that certain mathematical propositions are both true and not true. Balaguer himself notes that the continuum hypothesis is one such proposition. But contradictions can never be true, hence we must reject FBP.

Explain how this objection is misguided.

Suggested readings

Balaguer, Mark. *Platonism and Anti-Platonism in Mathematics*. New York: Oxford University Press, 1998. Balaguer's monograph defends plenitudinous platonism in the first part and fictionalism in the second. He concludes that there is no fact of the matter which will allow us to decide on one or other view.

Field, Hartry. *Realism, Mathematics and Modality*. Oxford: Blackwell, 1989. Field's fictionalism, an important influence on Balaguer's view, was first presented in *Science without Numbers* and is extended here, with responses to critics.

[5]Mark Balaguer, *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998), 59.

[6]Mark Balaguer, *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998), 89.

READINGS

Balaguer, "Full-Blooded Platonism"[7]

In this paper, I will argue that if we're going to endorse mathematical platonism, then we should endorse a very specific version of this theory, namely, what I have called *plenitudinous platonism*, or *full-blooded platonism*, or FBP.[8] I think there are a number of good arguments for this claim, but in this paper, I will focus on just one of them—an argument that's based on the idea that FBP is the only version of platonism that allows us to provide an acceptable account of mathematical *knowledge*.

In section 1, I will explain what mathematical platonism is. In section 2, I will lay out what's known as the epistemological objection to platonism; this objection is based on the claim that platonism is incompatible with the fact that human beings have mathematical knowledge. In section 3, I will explain how platonists can respond to this objection if they endorse FBP. And in section 4, I will respond to some objections to my view.[9]

I should note that this paper is intended to be somewhat introductory. For a more in-depth and less introductory treatment of these issues, see my (1998) and (2009).

Mathematical platonism

Let's start by defining platonism. After that, we can define *mathematical* platonism. *Platonism* (as I will use the term here, and as it is widely used in contemporary debates in metaphysics and the philosophy of mathematics) is the view that there are abstract objects. In other words, it's the view that there really do exist such things as abstract objects. But what does this mean? What's an abstract object?

An *abstract object* is an object that's non-physical, non-mental, and non-spatiotemporal. Thus, abstract objects don't exist in space and time, and they aren't made of physical stuff. They also don't exist inside the heads of human beings; that is, they aren't just ideas in our heads. On the contrary, if there are any abstract objects, then they exist independently of us. They are every bit as objective as planets and tables and cats. But, again, unlike planets and tables and cats, they aren't made of physical stuff.

[7]I would like to thank Russell Marcus and Mark McEvoy for helpful comments on an earlier draft of this paper.

[8]I first introduced this theory in "A Platonist Epistemology," *Synthese* 103 (1995): 303–25. The fullest development of the view is in *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998). For some more recent thoughts, see "A Theory of Mathematical Correctness and Mathematical Truth," *Pacific Philosophical Quarterly* 82 (2001): 87–114; and "Fictionalism, Theft, and the Story of Mathematics," *Philosophia Mathematica* 17 (2009): 131–62. Also, I should say that a similar view is endorsed by Edward Zalta and Bernard Linsky, "Naturalized Platonism vs. Platonized Naturalism," *Journal of Philosophy* 92 (1995): 525–55.

[9]This paper is intended to be somewhat introductory. For a more in-depth and less introductory treatment of these issues, see *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998) and "Fictionalism, Theft, and the Story of Mathematics," *Philosophia Mathematica* 17 (2009): 131–62.

Let me try to clarify this with an example. Consider *redness*—or as we might also put it, *the color red*. Let's suppose that there are one million red things in the universe. There's Mars, there's the Golden Gate Bridge, there's the tomato in my refrigerator, and so on. Moreover, in addition to these things, there are numerous mental representations of redness that exist in our heads: there's *my* idea of redness, *your* idea of redness, *Peyton Manning's* idea of redness, and so on. But you might think that none of these things is redness itself. You might think that in addition to all the things just mentioned—in addition, that is, to the various red objects that exist in the physical world and the various ideas of redness that exist in our heads—there is also *redness*, i.e., the color itself. And if you thought this, then you would probably be inclined to say that redness—the color itself—is an abstract object. That is, you would probably be inclined to endorse a platonist view of redness.

Now, this, of course, is just an example. Platonism, in general, is the view that there are at least *some* abstract objects. The view just described—the view that colors are abstract objects—can be called *color platonism*.

Another example of a platonist view is *mathematical platonism*. This is the view that mathematical objects—things like *numbers*—are abstract objects. Consider, e.g., the number 3. Platonists think that 3 is a non-physical, non-mental, non-spatiotemporal object. Now, of course, platonists admit that there are piles of three objects that exist in the physical world; e.g., there are three pennies in my pocket and three eggs in my fridge. And they also admit that there are 3-ideas in our heads. But they think that none of these things is the number itself. The number itself—3—is, according to platonists, an abstract object.

This is obviously a controversial view. Lots of people think there are no such things as abstract objects. Call these people *anti-platonists*. They think that in the physical world, there are lots of red objects (Mars, the Golden Gate Bridge, and so on), and there are lots of piles of three things (the three eggs in my fridge, the three pennies in my pocket, and so on). And they also think there are lots of ideas in our heads—there's my idea of redness, your idea of 3, and so on. But unlike platonists, anti-platonists think that *that's all there is*. There are no weird abstract objects that exist over and above the various physical and mental objects that we're familiar with.

So platonism is definitely controversial. Indeed, at first glance, it can seem a bit crazy. Believing in abstract objects can seem like believing in ghosts or fairies. Why would a level-headed, scientifically minded person believe in such things?

It turns out that there is an extremely powerful argument for mathematical platonism. I am not going to develop this argument in detail in this paper, but I would like to say a few words about it.[10]

[10]The argument I'll give here is deeply similar to arguments developed by Frege; see, e.g., *The Foundations of Arithmetic*, translated by J. L. Austin, 2nd ed. (Evanston: Northwestern University Press, (1884) 1980) and "The Thought: A Logical Inquiry," translated by A. M. and M. Quinton, in *Essays on Frege*, edited by E. Klemke (Urbana, IL: University of Illinois Press, (1919) 1968), 507–35. Platonism has also been famously endorsed by Bertrand Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, (1912) 1959); Kurt Gödel, "What is Cantor's Continuum Problem?" in **BP**, 470–85; and of course, Plato (see, e.g., Meno and Phaedo, translated by G. M. A. Grube, in *Five Dialogues*. Indianapolis, IN: Hackett Publishing, 1981).

The first point to note here is that the discipline of mathematics seems to be in the business of giving us *theories*. In other words, like other sciences, it attempts to *tell us things*, or to *state facts*. For instance, Euclid proved the following theorem:

(A) There are infinitely many prime numbers.

And more mundanely, mathematics tells us things like this:

(B) 3 is prime.

Sentence (B) seems obviously true, and it seems to be a claim about the number 3, and so it seems that the number 3 must exist. To appreciate this, let's change the example. Consider the following sentence:

(C) Mars is round.

This sentence seems obviously true, and moreover, it seems clear that the truth of (C) requires the existence of Mars. If Mars doesn't exist at all, then it couldn't be round. And the same goes for (B): if the number 3 doesn't exist at all, then it couldn't be prime, and so if sentence (B) is true, then it follows that the number 3 exists.

Given that the number 3 exists, the next thing that platonists argue is that it couldn't be a physical object or a mental object. There are several compelling arguments for this. One of them is that questions about the existence of physical and mental objects—or about how many physical and mental objects there are—seem entirely irrelevant to claims about numbers. Consider, again, Euclid's theorem—i.e., sentence (A) above. This sentence says that there are infinitely many prime numbers. If this sentence were about physical or mental objects, then its truth would depend on the controversial idea that there actually exist infinitely many physical or mental objects. But this just seems obviously wrong; it seems that the truth of Euclid's theorem *doesn't* depend on this controversial claim. Imagine, a mathematics professor proving Euclid's theorem for her students (the proof is not very hard), and imagine a student raising his hand with the following objection:

> I understand the proof, Professor, but something must be wrong with it, because my physics professor told me that there are only finitely many physical objects in the entire universe, and we also know that there are only finitely many mental objects (after all, there are only finitely many people, and each person has finitely many thoughts). So there couldn't be infinitely many prime numbers for the simple reason that there are only finitely many physical and mental objects in the universe.

How should the professor respond to this objection? It seems that she should respond by saying something like this:

> None of that is relevant to Euclid's proof. Because Euclid wasn't talking about physical or mental objects. He was talking about numbers. The proof goes through no matter how many physical and mental objects there are. In fact, even if there were no physical objects, and even if there were no people, it would still be true that there are infinitely many prime numbers. Because Euclid's proof doesn't rely on any claims about how many physical or mental objects there are.

This seems right. And if it is, then the only reasonable conclusion to draw is that when we talk about numbers, we're not talking about physical or mental objects. And, the argument continues, the only other thing we could be talking about is abstract objects. But if sentences about numbers—sentences like (A) and (B) above—are about abstract objects, then it would seem that we have to endorse the existence of abstract objects. For, again, these sentences seem true, and their truth seems to require the existence of the objects that they're about. Thus, if the objects that they're about (namely, numbers) are abstract objects, then it seems that we're committed to endorsing the existence of abstract objects.

That's the standard argument for mathematical platonism. There are, of course, responses that one might make to this argument, but I won't get into this here.[11]

The epistemological argument against platonism

In the remainder of this paper, I will be concerned only with mathematical platonism and not with platonism in general. But I will sometimes drop the qualifier and use the term "platonism" to refer to mathematical platonism.

We saw in the last section that there is a strong argument for mathematical platonism. But there are also reasons to disbelieve that view. Probably the most powerful and important argument against platonism is what's known as the *epistemological argument*.[12] This argument can be formulated as follows:

(1) Human beings exist entirely within space and time.

(2) If abstract objects exist, then they do not exist in space and time.

(3) If (1) and (2) are both true, then even if abstract objects exist, human beings could not acquire knowledge of them. Therefore,

(4) If abstract objects exist, then human beings could not acquire knowledge of them. But

(5) If mathematical platonism is correct, then abstract objects do exist and human beings *can* acquire knowledge of them. Therefore,

(6) Mathematical platonism is not correct.

The important part of this argument is the part in statements (1)–(4). Premise (5) seems pretty obvious (after all, platonists think that our mathematical knowledge just *is* knowledge of abstract objects), and the inference from (4) and (5) to (6) is clearly valid. Thus, it seems that if (4) is true, then (6) is also true, and so platonists need to

[11]I have argued elsewhere that the only viable way to respond to the argument for platonism is to endorse *fictionalism*—i.e., to admit that our mathematical sentences and theories are supposed to be about abstract objects and to say that since there are in fact no such things as abstract objects, our mathematical sentences and theories aren't strictly speaking true. They are rather useful fictions. For more on this view, see, e.g., Hartry Field, *Science Without Numbers* (Princeton: Princeton University Press, 1980); Hartry Field, *Realism, Mathematics and Modality* (Oxford: Blackwell, 1989); Mark Balaguer, *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998); and Mark Balaguer, "Fictionalism, Theft, and the Story of Mathematics," *Philosophia Mathematica* 17 (2009): 131–62.

[12]This argument goes all the way back to Plato, but it has received renewed interest since the publication of Benacerraf's "Mathematical Truth," Chapter 15, above.

block the argument for (4). Moreover, since the argument for (4) is valid, it seems that platonists have no choice but to reject either (1), (2), or (3).

The strategy of rejecting (3) looks pretty difficult. If platonists accept (1) and (2), then they have to admit that there's no way for us to *gather information* about abstract objects like numbers. Or to put the point differently, there's no way for information to flow from abstract objects to human beings. Thus, if platonists pursue the strategy of rejecting (3), what they'll need to do is explain how we could acquire knowledge of abstract objects—e.g., numbers—despite the fact that we can't have any information- gathering *contact* with such objects. And, again, it might seem that this would be pretty hard to do.

The problem, though, is that the other two strategies—the strategies of rejecting (1) and rejecting (2)—seem completely untenable. If platonists pursue either of these two strategies, they can claim that human beings are capable of having some sort of information-gathering contact with numbers. For instance, if they pursue the first strategy—i.e., if they reject (1)—then the claim will presumably be that humans are capable of somehow "leaving" the physical, spatiotemporal world and "accessing" the platonic realm and gathering information about what abstract objects are like. Most people who work in this area would say that this view is pretty implausible. Indeed, if you endorse a naturalistic, scientific view of the world (and of human beings), then the view probably seems extremely implausible.[13] It's worth noting, however, that this view has been endorsed by some very influential thinkers—namely, Plato and Gödel, on at least some interpretations of their work.[14]

The strategy of rejecting (2) might seem even more untenable than that of rejecting (1). This is because platonists are liable to think that (2) is true by *definition*. Platonists think that abstract objects are non-spatiotemporal—i.e., that while they exist, they don't exist in space and time. But that's all that (2) *says*—that abstract objects don't exist in space and time. So it seems that platonists have to accept (2). There's one philosopher, though—namely, Penelope Maddy[15]—who has tried to respond to the epistemological argument by developing a non-standard version of platonism on which (2) is false. Roughly, the view is that we can take mathematical objects like numbers to be *sets* (i.e., sets of objects, like the set of eggs in my fridge), and we can say that (a) sets exist in space (e.g., the set of eggs is located *in my fridge*, right where the eggs themselves are located) and (b) we can acquire information about sets in the same way that we acquire information about ordinary physical objects—namely, via sensory perception.

I think there are compelling arguments against the strategy of rejecting (1) and the strategy of rejecting (2). I have developed these arguments elsewhere,[16] but I do not have the space to run through the arguments here. Instead, I just want to explain how

[13]Actually, I have argued elsewhere, in Mark Balaguer, *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998), that even if you endorse the Cartesian view that every human being has a non-physical soul, you still shouldn't claim that we can acquire knowledge of abstract objects by, so to speak, "accessing platonic heaven."

[14]See Kurt Gödel, "What is Cantor's Continuum Problem?" in **BP**, 470–85; and Plato's *Meno and Phaedo*, translated by G. M. A. Grube, in *Five Dialogues*. Indianapolis, IN: Hackett Publishing, 1981.

[15]Penelope Maddy, *Realism in Mathematics* (Oxford: Clarendon Press, 1990).

[16]In my *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998).

I think platonists *can* solve the epistemological problem. I think they can do this by rejecting (3). More precisely, I think they can accept that (1) and (2) are true—and, hence, that we don't have any sort of information-gathering contact with abstract objects like numbers—and I think they can explain how we can nevertheless acquire knowledge of what these objects are like. I said above that, at first blush, this looks like a hard thing to do. But despite this, I think it can be done, and I now want to explain *how* it can be done.

The FBP-ist response to the epistemological argument

I think that platonists can adequately respond to the epistemological argument if they endorse a specific version of platonism that I have elsewhere called *plenitudinous platonism*, or *full-blooded platonism*, or FBP. FBP can be expressed, somewhat roughly, as the view that all possible abstract mathematical objects exist, or the view that all the abstract mathematical objects that *could* exist actually *do* exist, or the view that there are as many abstract mathematical objects as there could be.

FBP can be contrasted with what might be called *sparse platonism*. This is the view that (a) there do exist abstract mathematical objects, but (b) there are certain *possible kinds* of mathematical objects that have no instances. We can clarify this by noting that the physical world is obviously sparse. In other words, it seems obvious that there are certain possible kinds of physical objects that have no instances. For example, there are no 400-story buildings, and there are no talking donkeys. Plenitudinous realism about the physical world (or full-blooded realism) is the view that there actually exist physical objects of all of these kinds—i.e., it's the view that there actually exist things like 400-story building and talking donkeys and so on. This, of course, is a crazy view. But full-blooded *platonism* is not so crazy. Indeed, I want to argue that this is the best kind of platonism there is—i.e., that if we're going to be platonists at all, we should be FBP-ists.

I think there are numerous arguments for the claim that FBP is the best version of mathematical platonism, but in this paper, I will focus on just one of them, namely, the one that's based on the claim that FBP gives platonists a way—a very *plausible* way—of responding to the epistemological argument against their view. In particular, FBP gives platonists a way of explaining how human beings can acquire knowledge of abstract objects like numbers, despite the fact that they have no information-gathering contact with such objects. In a nutshell, the explanation proceeds as follows:

> Since FBP says that there are abstract mathematical objects of all possible kinds, it follows that if FBP is true, then every purely mathematical theory that could be true—i.e., that's internally consistent—accurately describes some collection of actually existing abstract objects. Thus, it follows from FBP that in order to acquire knowledge of abstract objects, all we have to do is come up with an internally consistent purely mathematical theory (and know that it's internally consistent). This is because, again, according to FBP, *every* consistent purely mathematical theory accurately describes a collection of actually existing abstract objects. But if all we need to do in order to acquire knowledge of abstract objects is come up with a consistent mathematical theory (and know that it's consistent), then it seems that we *can* acquire such knowledge. This is because it seems clear that (a) we *are* capable of formulating internally consistent mathematical theories (and of knowing that they're internally consistent), and (b) being able to do this does not require us to have any sort of information-gathering contact with the abstract objects that the theories in question are about. Thus, if all of this is right, then the epistemological problem with platonism has been solved.

We can better understand how this account of mathematical knowledge is supposed to work by showing how we could adopt an analogous account of ordinary knowledge of the physical world, *if we endorsed full-blooded realism about the physical world.* Recall what full-blooded realism about the physical world says: it says that there actually exist physical objects of all possible kinds. Thus, according to this view, there are things like talking donkeys, and talking donkeys who live in Cleveland, and talking donkeys who live in Cleveland and have purple eyes, and so on. Now, this view is obviously false, but if it were true—if the physical world were plenitudinous in this way—then every story about the physical world that was possible, or internally consistent, would accurately describe some bunch of objects. Thus, we could acquire knowledge of ordinary physical objects by simply coming up with an internally consistent story (and knowing that it was consistent). Thus, since we are capable of coming up with such stories (and of knowing that they're consistent), and since doing this doesn't require any information-gathering contact with the objects in question, it follows that if full-blooded realism about the physical world were true, then we could acquire knowledge of the physical world in this way.

We can make this more clear by considering a concrete example. Consider the following story:

> *The story of Zoton:* On the opposite side of the galaxy, there is a planet called Zoton, and on this planet there is a community of flying, talking donkeys who have purple eyes.

If full-blooded realism about the physical world were true, then I could read this story, and say, "Hmm, I just learned something. I now know that there are flying, talking, purple-eyed donkeys." If you asked me how I know this, I would say that since the story of Zoton is clearly possible, it follows that since full-blooded realism about the physical world is true, this story accurately describes some actually existing bunch of objects. And if you asked me how I know that the story of Zoton is possible, I would scratch my head and say, "What do you mean? That's just simple *logical* knowledge. I don't need to go to Zoton—or have any information-gathering contact with the creatures on Zoton—to know that this story is possible, or consistent. I just need to know that there is no incompatibility between, e.g., being able to fly, and being able to talk, and being purple-eyed, and so on.

Now, of course, full-blooded realism about the physical world is not true, and so we can't acquire knowledge of the physical world in this way—by just dreaming up possible stories. But full-blooded *platonism* is not so obviously false. Indeed, there seems to be some independent reason to think that it's the most obvious version of platonism to endorse. We can appreciate this by returning to the platonistic view of *colors*. Wouldn't it seem crazy to believe in some colors but not others—e.g., to say that while *redness* is a real abstract object, *blueness* isn't? If we're going to believe that some colors are abstract objects, we should say that they all are. And the same goes for numbers. If we believe that even numbers exist (and that they're abstract objects), then we should say the same thing about odd numbers as well. And likewise, it would seem, for all mathematical objects, like sets and functions and geometrical shapes.

In any event, if FBP is true, then we can acquire knowledge of mathematical objects in the above way—i.e., by simply dreaming up a possible story about mathematical objects. Consider, e.g., the following story:

> *The story of natural numbers*: In the mathematical realm—or as we might more metaphorically put it, *platonic heaven*—there is a sequence of things called natural numbers. The first one is 0. The successor of 0 is 1; the successor of 1 is 2; and so on—the sequence keeps going. Or more precisely, every natural number has a unique successor.

If FBP is true, then I can read this story, and say, "Hmm, I just learned something. I now know that 0 is the smallest natural number and that 1 is its successor." If you asked me how I know this, I would say that since the story of the natural numbers is pretty clearly possible,[17] it follows that since FBP is true, this story accurately describes some actually existing collection of objects. And if you asked me how I know that the story of natural numbers is possible, I would scratch my head and say, "What do you mean? That's just *logical* knowledge. I don't need to go to platonic heaven—or have any information-gathering contact with numbers—in order to know that this story is possible, or consistent. I just need to know that the story isn't internally inconsistent, or that it doesn't contradict itself.

So FBP gives us a sort of recipe for acquiring knowledge of abstract objects. The recipe goes like this:

1. Dream up a mathematical story. Or more precisely: come up with a kind of mathematical object and articulate a theory about those objects. For instance, the system of axioms known as *Peano Arithmetic* (*PA*) is a theory about the natural numbers; and *Zermelo-Frankel set theory* (*ZF*) is a theory about sets; and so on.
2. Make sure—via logic—that the theory is internally consistent.

[17]Some people might doubt that the story of natural numbers is possible. I don't have the space to address this worry here, but in a nutshell, I think FBP-ists can respond to it by saying something like the following: (a) *prima facie*, there doesn't seem to be anything impossible about the story of numbers—there doesn't seem to be any more reason to doubt that this story is possible than that the story of Zoton is possible—and (b) since claims of possibility are so weak, this already gives us good reason (although, surely, a defeasible reason) to think that the story of numbers is possible.

3. Conclude that the theory accurately describes a collection of abstract objects. E.g., PA accurately describes the natural numbers; ZF accurately describes a hierarchy of sets; and so on.
4. Finally, if you like, you can use logic to deduce (or prove) further facts about the objects in question from the theory that you've articulated. These facts can be surprising, or non-obvious. For instance, from PA we can prove (and Euclid did prove) that there are infinitely many prime numbers.

It's important to note that this recipe fits quite well with what we know about actual mathematical practice—i.e., with the methodology that mathematicians actually use to acquire mathematical knowledge. The methodology I'm referring to is that of axioms and proofs. We start by laying down basic axioms, and then we prove theorems from the axioms. The laying down of axioms is just step 1 in the above recipe—i.e., it's the articulation of a (hopefully consistent) theory about a certain part of the mathematical realm that we have in mind—or that we've "dreamed up."

One way to think about all of this (and as we'll see in a moment, what I'm about to say will require some qualification) is to say that on the FBP-ist picture, what we're doing when we lay down axioms is more or less *stipulating which abstract objects we're talking about*—or which part of the mathematical realm we're talking about. To bring this out, let's consider an example. Suppose that we're doing arithmetic—or studying the sequence of natural numbers—and suppose that we list the following sentence as an axiom: "Every number has a successor." Now, suppose that someone objects to this by saying the following:

> How do you *know* that every number has a successor? After all, since the natural numbers are non-spatiotemporal abstract objects, you don't have any information-gathering contact with them. So how do you know what they're like? For instance, how do you know that the number 16 has a successor? For all you know, there might be no such thing as the number 17. There might be a hole in the natural-number sequence where 17 is supposed to be.

On the FBP-ist picture, we can respond to this objection as follows:

> You don't get it. I'm talking about the sequence of objects that doesn't have a hole where 17 is supposed to be. Of course, since FBP is true, there's another sequence that *does* have a hole there. But that's not the one I'm talking about. I'm talking about the one that doesn't have a hole there. More precisely, I'm talking about the sequence of abstract objects that's described by the standard axioms of arithmetic, e.g., by "0 is a number," "Every number has a successor," and so on.

This, I think, is exactly right. But I now want to point out that it's not always strictly speaking true to say that the axioms are *stipulations*. To say that they're stipulations is to say that they're *definitions*. But while it's *sometimes* true that our axioms are definitional (and while it's almost always *close* to true), it isn't always exactly accurate to say that they're definitions. For (a) in some contexts (or in some branches of mathematics), we have pre-theoretic conceptions of the mathematical structures that we're talking about (e.g., in arithmetic, we have a pre-theoretic conception of

the sequence of natural numbers); and (b) the objects (or the structure, or the part of the mathematical realm) that we're talking about in a given context (or a given branch of mathematics) can be thought of as being determined not just by the axioms in question but by our *full conception* of the relevant objects (e.g., in arithmetic, we can be thought of as talking about the objects that correspond not just to the standard axioms of arithmetic, but to our full conception of the natural numbers); and (c) in situations like this, our axioms can be thought of as claims about *these* objects—i.e., the objects that are picked out by our full conception of the relevant objects—and so the axioms aren't really definitions; they could in principle turn out to be false; for they could be false claims about the objects that are picked out by our full conception of the relevant objects.

Let me make three more points about this before moving on. First, it's important to note that there are *some* cases in which our axioms *are* definitional. For sometimes mathematicians just "play around" with axiom systems to see what follows from them, and in cases like this, the axioms really are just stipulations about which objects we're talking about. In other words, the structures we're talking about in cases like this just are the structures that are characterized by the given axioms. Second, even when an axiom isn't strictly definitional, if it's just an obvious, core part of our full conception of the relevant objects, then it can be, so to speak, "close to definitional." For instance, in arithmetic, the axiom "Every number has a successor" can be thought of as a sort of "virtual definition" because it's so central to our conception of the natural numbers that it seems that any structure in which this claim isn't true is *thereby* not the sequence of natural numbers. Third and most importantly, the fact that our axioms aren't always strictly definitional doesn't undermine the FBP-ist response to the epistemological argument against platonism. For on the FBP-ist picture, even when the axioms aren't strictly definitional, it's still true that the part of the mathematical realm that we're talking about is being determined by *our* thoughts; in particular, it's determined by our full conception of the relevant objects. So in these contexts, we still get the result that we're performing step 1 of the above recipe; in other words, we're still using a methodology in which we start by (a) thinking up a kind of mathematical object and (b) articulating a (hopefully consistent) theory of those objects. When (a) and (b) are fused into one mental episode, the axioms are strictly definitional; when (a) and (b) aren't fused—when the articulation of the axioms comes *after* the "thinking up of the objects"—the axioms aren't strictly definitional; but in these cases, the "thinking up" part still plays the role of a stipulation (even if it's not a conscious stipulation) because it literally *determines* which objects (or which part of the mathematical realm) the given theory is about.

Objections and responses

In this section, I will respond to two objections to FBP and the FBP-ist theory of mathematical knowledge laid out above. The first objection can be put like this:

> Objection 1: In order for human beings to acquire knowledge of abstract objects in the above way, they would first need to know that FBP is true. Think about it.

> Without having any information-gathering contact with abstract objects, we can know that *if FBP is true*, then there's a sequence of objects that's characterized by the standard axioms of arithmetic, and so we can know that in this sequence, the number 16 has a successor (namely, 17), and there are infinitely primes, and so on. But we can't in this way come to know that there really *are* infinitely many prime numbers, because we haven't been given any way to know that FBP is true. Indeed, we haven't been given a way to know that there are any abstract objects at all.

Response to objection 1

This objection involves a misunderstanding of the epistemological challenge to platonism—i.e., the challenge that's generated by the epistemological argument. The purpose of that argument is to generate a genuine wonderment about how human beings could have *any idea* what the mathematical realm is like. The purpose is *not* to raise a Descartes-style skeptical worry. Descartes famously raised a worry about whether we can have any knowledge of an external (physical) world. That's all fine and good. But no one seriously wonders how human beings could acquire accurate beliefs about an external world, because it's entirely *obvious* how they could do this—they could point their eyes at physical objects, and photons could bounce off of those objects and into their eyes, carrying information about what those objects are like. So it's entirely obvious how human beings could acquire accurate beliefs about physical objects. But despite this, you might still have a Descartes-style worry about whether our sense perceptions really give us knowledge; you might wonder whether we could really know that there even *is* an external world.

The purpose of the FBP-ist response to the epistemological argument is to explain how we humans could acquire accurate beliefs about abstract mathematical objects, using our *ordinary* mathematical methodology. We could do this by (a) laying down axioms that amount to stipulations (or something like stipulations) about which mathematical objects we're talking about—or which part of the mathematical realm we're talking about—and (b) proving theorems from these axioms. If this is right, then the epistemological worry laid out in section 2 has been answered. Now, you might still have a Descartes-style *skeptical* worry about mathematics; you might still wonder whether this methodology really gives us knowledge; you might think that we could never really know that there even *is* a mathematical realm. But that's a different worry. And it's a worry that applies to our ordinary knowledge of physical objects as well to our knowledge of abstract objects.

There might be no adequate answer to this skeptical worry. That wouldn't be so bad. After all, it's plausible to suppose that there's no adequate answer to the corresponding skeptical worry about knowledge of the physical world. But even if there's an unanswerable skeptical worry about how we could have real knowledge of an external world, we still have a plausible story to tell about how our sense perceptions could give us accurate beliefs about physical objects. The epistemological argument against platonism was supposed to show that in the case of abstract objects, we don't even have this. But the above FBP-ist theory of mathematical knowledge

shows that this worry is misguided. It shows that even if there's an unanswerable skeptical worry about how we could have real knowledge of an external, platonistic mathematical realm, we have a plausible story to tell about how human beings could acquire accurate beliefs about abstract mathematical objects via our normal mathematical methodology—i.e., by using the method of laying down axioms and proving theorems from those axioms.

The second objection that I want to consider in this section is more complicated than the first, and in responding to it, I will further develop the FBP-ist view. The objection can be put in the following way:

> Objection 2: FBP entails that all consistent purely mathematical theories accurately describe actually existing collections of abstract objects. So it seems to entail that all such theories are *true*. But this leads to a contradiction. For there are pairs of consistent mathematical theories that contradict each other. Consider, for instance, ZF+CH and ZF+~CH. The former is standard Zermelo-Fraenkel set theory (ZF) plus the continuum hypothesis (CH); and the latter is ZF plus the *negation* of the continuum hypothesis (~CH). (CH is an open question in set theory. It doesn't matter what exactly it says, but for the curious, it says that the set containing all the real numbers is the second smallest kind of infinity.) We know that both of these theories—ZF+CH and ZF+~CH—are internally consistent.[18] Thus, FBP seems to entail that they're both true. But if they're both true, then CH and ~CH are both true, and that's a direct contradiction.

Response to objection 2

The problem with this objection is that it assumes that if a mathematical theory accurately describes a collection of abstract objects, then it's automatically true. In other words, it assumes the following:

> *Silly theory of mathematical truth* (or for short, *ST*): A mathematical sentence or theory is *true* if and only if it accurately characterizes some collection of mathematical objects—i.e., if and only if it's true of *some* mathematical structure, or some part of the mathematical realm.

FBP-ists should say that this theory is false. One reason is that ST leads to a contradiction—namely, the one mentioned in objection 2. But this isn't the only reason to reject ST. Another reason is that it fails to capture the intuitive notion of truth. What it captures is the notion *satisfiability*, or *truth in a structure*. But there's a difference between being true in a structure—*any* structure—and being *true*. The difference is that ordinary truth—or truth *simpliciter*—has to do with being true in the *intended* structure. For instance, if you're talking about the sequence of natural numbers and you say, "The number 16 doesn't have a successor," then what you said isn't true. Now, of course, if FBP is true, then this sentence is true in *some* structure;

[18] Another way to put this is as follows: we know that the theory of ZF doesn't settle the question of whether CH is true. In other words, it's impossible to prove CH or ~CH from ZF—because ZF is consistent with both CH and ~CH.

in particular, it's true in non-standard structures in which 16 doesn't have a successor. But despite this, it's not *true*, because it's not true in the *intended* structure—i.e., it's not true of the sequence of natural numbers, which is what you were talking about when you made your claim.

Given these remarks, it might seem that FBP-ists should endorse the following theory, instead of ST:

> *Better theory of mathematical truth* (or for short, *BT*): A mathematical sentence or theory is true if and only if it's true in the intended structure, or the intended part of the mathematical realm—i.e., the part of the mathematical realm that we have in mind in the given branch of mathematics. Thus, e.g., an arithmetical sentence is true just in case it's true *of the sequence of natural numbers*; and a set-theoretic sentence is true just in case it's true *of the universe of sets*—i.e., the universe of things that corresponds to our intentions with the word 'set', or to our full communal conception of the universe of sets, or some such thing.

This is a big improvement over ST, and it enables us to avoid the contradiction alluded to in objection 2. If FBP is true, then ZF+CH and ZF+~CH are both *true in a structure*—i.e., they're both true of some collection of abstract objects—but it doesn't follow that they're both *true*, because it doesn't follow that they're true in the intended structure for set theory.

But there's a problem with BT: it assumes that there will always be a unique intended structure in every branch of mathematics. But this assumption is unjustified. It may be that in some branches of mathematics (or some mathematical conversations), our intentions have some imprecision in them. In other words, it may be that our *full* conception of the objects in question (or the objects being studied) isn't strong enough, or precise enough, to zero in on a unique structure (or more precisely, a unique structure up to isomorphism). Indeed, this *might* be the case in set theory. For instance, it may be that there's a pair of structures, call them H1 and H2, that satisfy the following three conditions:

(i) ZFC+CH is true in H1;

(ii) ZFC+~CH is true in H2;

(iii) H1 and H2 both count as intended structures for set theory because they're both perfectly consistent with *all* of our set-theoretic intentions—or with our full communal conception of the universe of sets.

Thus, it seems that there could be some mathematical sentences (and it *may* be that CH is such a sentence) that are true in some intended structures, or some intended parts of the mathematical realm, and false in others.

So I don't think BT is right. I think that FBP-ists should endorse the following theory instead:

> *Intention-based theory of mathematical truth* (or for short, *IBT*): A pure mathematical sentence S is *true* if and only if it's true in *all* the parts of the mathematical realm that count as intended in the given branch of mathematics (and there is at least one such part of the mathematical realm); and S is *false* if

> and only if it's false in all such parts of the mathematical realm (or there is no such part of the mathematical realm); and if S is true in some intended parts of the mathematical realm and false in others, then there is no fact of the matter as to whether it is true or false.

This gives us the result that if CH is true in all of the intended structures for set theory, then it's true; and if it's false in all such structures, then it's false; and if it's true in some of these structures and false in others, then there's no right answer to the question of whether CH is true.[19]

It's important to remember in this context that sometimes, there is no pre-theoretic conception of the relevant objects—because, again, sometimes mathematicians just "play around" with axiom systems to see what follows from them. In settings like this, any structure that satisfies the given axioms thereby counts as *intended*. Moreover, this could happen with CH—or better, it could happen with certain *utterances* of CH. Suppose I was exploring the axiom system ZF+CH without worrying about whether I was capturing our intuitive conception of set; and suppose you were doing the same thing with ZF+~CH; and finally, suppose that in this setting, I said that CH was true and you said that ~CH was true. Both of our utterances would—*in this context*—be true. But this would not be a contradiction, because we would be talking about different objects. Here's an analogy: suppose that while thinking of Aristotle Onassis, I say, "Aristotle married Jackie Kennedy"; and suppose that while thinking of Aristotle the philosopher, you say, "Aristotle didn't marry Jackie Kennedy." Both of these utterances are true, but this is not a contradiction, because when we made these claims, we were talking about different people. And the same goes for our claims about CH and ~CH: in the above context, both of our utterances are true, but they don't contradict one another, because they're about different objects.

(Of course, this won't be the case in *all* contexts. If you and I both intended to be talking about the intended structures for set theory—i.e., the structures that correspond to our full communal concept of set—and if I said that CH was true in these structures and you said it wasn't, then at most one of our claims could be true.)

There are many more objections that could be raised against FBP and IBT, and there is much more that could be said in defense of these views than I have been able to say here.[20]

[19]You might think it's problematic that IBT entails that there could be mathematical sentences that are neither true nor false. I can't get into this here, but I argue in "Fictionalism, Theft, and the Story of Mathematics," *Philosophia Mathematica* 17 (2009), that this is in fact *not* a problem and, indeed, that it's a welcome result that gives us positive reason to endorse IBT.

[20]For a fuller defense of these views, see my *Platonism and Anti-Platonism in Mathematics* (New York: Oxford University Press, 1998) and "Fictionalism, Theft, and the Story of Mathematics," *Philosophia Mathematica* 17 (2009): 131–62.