9
Formal Tools and the Philosophy of Mathematics

Thomas Hofweber

1 How can we do better?

In this chapter, I won't try to defend a particular philosophical view about mathematics, but, in the spirit of *New Waves*, I would instead like to think about how we can hope to make real progress in the field. And to do that it seems best to think about what in the field as it is now is holding up progress. What in the field should change so that we can hope to do better than that has been done before? Having considered the question, I propose the following answer: The single biggest obstacle to real progress in the philosophy of mathematics is a lack of reflection on which questions should be addressed with what methods. And this is particularly apparent in the role that formal tools have in the discipline as it is today. Formal tools encompass paradigmatically formal, artificial languages, formal logic expressed with such languages, and mathematical proofs about such languages. These tools were developed during the rise of logic over 100 years ago, and they are ubiquitous in the philosophy of mathematics today. But even though formal tools have been used with great success in other parts of inquiry, in the philosophy of mathematics they have done a lot of harm, besides quite a bit of good. In my contribution to this volume, I would thus like to think a bit about what role formal tools should have in the philosophy of mathematics. I will argue that they should have merely a secondary role, unless one holds certain substantial views in the philosophy of mathematics, ones that very few people hold. The role of formal tools is thus tied to substantial questions in the philosophy of mathematics.

Much work done under the heading of "philosophy of mathematics" consists of proofs, precise mathematical proofs. Many talks at philosophy of mathematics conferences present new proofs, and often a proof is taken to be the surest sign of progress in the field. But this should be a bit puzzling. Proof is the method to establish results in mathematics. But it is rather unusual that the method to achieve results in the philosophy of X, some discipline, is the same as the method for achieving results in X. For example,

June 6, 2009 9:33 MAC/MATS Page-197 9780230_219434_11_cha09

physics achieves results via experimentation, amongst other methods. But the philosophy of physics does not. It might learn about the results of physics, achieved amongst others via experimentation. But it itself does not proceed via experimentation to do whatever it hopes to do. In general the method of the philosophy of X is distinct from the method of X, and this might be particularly compelling when X has a very distinct method, as does mathematics with that of precise proof. So, why should philosophers of mathematics hope to achieve results with the same method with which the discipline they hope to understand achieves its results?

The common justification for the role of proof in the philosophy of mathematics is this: metamathematics! Mathematical languages can be formulated with mathematical precision, and the notion of a proof can be characterized precisely as a consequence. Given this we can carry out proofs about proofs, including proofs about what can't be proven. And this justifies the central role of proofs and formal tools in the philosophy of mathematics, or so the story goes. This story is partly right, but also full of holes and sources of error. In this chapter, we will have a closer look at what role formal tools should have in the philosophy of mathematics. We will discuss both the use of formal languages and the use proof about them in metamathematics.

The role of formal tools depends crucially on what the answer is to two other questions: first, the question of the status of axioms in mathematics, and, second, the question whether or not all of mathematics is philosophically equal. We will discuss these two questions in more detail below. Thus these three questions are closely related:

1. What is the proper role of formal tools?
2. What is the status of axioms?
3. Is all of mathematics philosophically equal?

In the end, I will argue that much damage has been done in the philosophy of mathematics by relying on formal tools. Their role in the philosophy of mathematics is limited, unless one holds a certain position with respect to the other two questions. These answers to the other two questions are not widely accepted, but I do accept them. Even though I hold that the role of formal tools is limited, I think most philosophers should think that it is even more limited.

Much of the discussion in the following will have a special emphasis on formal tools in the philosophy of arithmetic, but this is merely our example. The general issues carry over to other parts of mathematics as well, and the conclusion I hope to reach is general.

2 Formal tools

Mathematics is first and foremost an activity carried out in ordinary natural language with the use of symbols. Mathematical conversations are carried out in natural language, mathematical papers are written in natural language, but with the additional use of symbols. When we wonder about the role of formal tools, what is not at issue is the role of symbols. To make this clear, let us distinguish three different uses of, for example, what is pronounced “four” in mathematics. First there is the *natural language expression* “four”. It is part of ordinary English, and it occurs in ordinary English sentences as

(1) You shouldn’t drink more than four beers.

Secondly, there is the symbolic use of “4”. It will occur in simple mathematics, as when an elementary school teacher writes on a blackboard

(2) What is 4+2?

And it also occurs in serious arithmetic, of course. How is the *symbolic use* of “4” related to the English word “four”? It is tempting to think that it is merely a symbolic abbreviation of English word. After all, for every number word in English there is a symbolic one that is pronounced just the same way. And symbolic uses of number words occur in ordinary written English, as in

(3) Over their life, the average Bavarian drinks 15432 beers.

Finally, there is the *formal use* of “4”, as it occurs in the artificial languages of formal logic.[1] This language is introduced and defined by logicians, and various expressions in it are given a meaning. Even if symbolic uses of number words are merely abbreviations of natural language expressions, there is a real question as to how formal uses of number words relate to the other ones.

This division of three different uses carries over to a variety of other expressions as well. For example, “for all” and symbolic and formal uses of “∀”. As a symbol it can merely be an abbreviation of a natural language expression. But as a formal expression, there is an issue how it relates to the other two uses.

When we are wondering what role formal tools should have in the philosophy of mathematics, we are not wondering what role symbols should have, but what role in particular artificial languages should have. Such artificial language are perfectly precise and thus subject to mathematical proofs about them. Much technical work in the philosophy of mathematics is closely related to mathematical proofs about such artificial languages. The question will be what role such proofs should have. But first, lets make clearer what is and what isn’t at issue.

3 Two projects in the philosophy of mathematics

When we ask what role formal tools should have in the philosophy of mathematics it of course depends on what one means by the philosophy of mathematics. The philosophy of mathematics is a diverse discipline and in a sense formal tools clearly play a central role in some parts. So, there is an uncontroversial case in favor of formal tools, and there is a controversial case about their status. Let us distinguish two large-scale projects, both usually carried out under the umbrella of the philosophy of mathematics. One of them is *foundational*. The foundational project is directly concerned with issues tied to formal mathematical theories. It tries to find out, say, which axioms of a formal theory are needed to prove what theorems. Or, whether a certain formal theory can be represented in another. Or, whether a certain sentence is independent of a certain set of axioms. Much great work is done here, and all of it is essentially tied to using formal languages, formal logic, and formal tools more generally. This is the uncontroversial part. Next, there is the *large-scale philosophical* project. This project concerns the large-scale philosophical questions related to such notions as truth, knowledge, fact of the matter as they apply to mathematics, and questions about the large-scale relationship between mathematics and other parts of inquiry, and so on. These are the questions that philosophy is traditionally concerned with, and here they are directed at mathematics. These questions are not directly about any formal languages, and so here there is a real issue whether formal tools play or should play a role in making progress in answering them. This is our first topic, and this is the more controversial part. This question is related to, but not exactly the same as, what role the foundational project should have in the large-scale philosophical project. The foundational project does not, of course, need to have any role in the large-scale philosophical one to be a good and legitimate project. There are other interesting things in the world than large-scale philosophical questions. But many people who engage in the foundational project have the ambition for their work to be of significance for the large-scale philosophical project as well. And there might well be a role for formal tools in the large-scale philosophical project that goes beyond the foundational project. What is at issue for us here is the role of formal tools in the large-scale philosophical project, which is related to, but not quite the same, as the role of the foundational project in the large-scale philosophical project. Many philosophers throughout the history of the philosophy of mathematics have given formal tools a central role in the project of finding answers to the large-scale philosophical questions. A clear example is Hilbert, who thought that the main missing part for the defense of a certain large-scale philosophical picture was a certain technical result: a proof of the consistency of a certain formal axiom system by finitistic means. Once that had been done, Hilbert thought, a certain large-scale picture of mathematics would be vindicated.[2] We will see more about this below.

Next, we should also distinguish a *substantial role* from an *insubstantial role*. Formal tools might be useful as tools of representation. We might use them to represent something or other about our mathematical activity or language, and such a representation might be useful in various ways. There is no question that formal tools can be useful this way, but this is insubstantial. But they might also be used to do something that can't be done without them, they might really make a substantial difference. They might play a crucial role in achieving results, not just representing them. Whether formal tools can have such a substantial role is what is at issue.

Finally, we should distinguish a *revisionary* from a *descriptive* role of formal tools. Someone might hold that formal tools should play a central role in the philosophy of mathematics since mathematics as it is presently done should be replaced with a discipline that relies exclusively on formal languages. But this revisionary approach will simply be sidelined here. What is at issue is to understand mathematics from a philosophical point of view, as it actually is. We are concerned with the descriptive project, not the revisionary one, in part because it is hard to see how the revisionary could be made plausible, given how well mathematics is doing. Revision in any discipline should be triggered by some problem in that discipline. Mathematics as it is today does not seem to need a large-scale revision.

Whether formal tools should have a substantial role in the descriptive, large-scale philosophical project is what is at issue. We will divide the discussion by considering various options where one might think formal tools might have such a role in the large-scale philosophical project.

4 Mathematical activity

Mathematics is, amongst other things, an activity. It is something people do, within a certain community, with results, winners and losers, and all that. To understand this activity, one might employ the social sciences. And the social sciences might rely on formal methods in their study of the social aspect of mathematics. Of course, such formal methods would not likely involve the formal tools we talked about above, namely artificial languages. They will more likely be formal methods from statistics to, for example, show that problems posed by mathematicians with tenure are given more weight than ones posed by those without tenure. But I take it that reliance on results from the social sciences that are achieved with the help of formal methods, like the ones from statistics, are only insubstantial for showing the role of formal tools in the large-scale philosophy of mathematics. Even if these tools are substantial parts of the social sciences, it remains to be seen why the results of the social sciences are a substantial part of the philosophy of mathematics. It is a big step from the sociology of any science to the philosophy of that science, and it seems no different in the philosophy of mathematics. In any case, though, this is not what anyone has in mind when they want to defend the role of formal tools in the philosophy of mathematics. What is more crucial for us is the role of the formal tools of formal logic, the use of artificial languages in the philosophy of mathematics.

What the activity of mathematics is like is a largely empirical question. It is to be studied like other questions about the activity of individuals or groups of individuals. This simple point seems to be somewhat under-appreciated, though. Some positions that are hotly debated in the philosophy of mathematics are nothing more than questions about mathematical activity, and thus largely empirical questions. However, in the philosophical debate they are hardly treated as such. Take, for example, the debate about fictionalism. Most philosophical defenses of fictionalism are not in the ballpark of what kind of considerations are required to defend it. At most, they can be taken to argue that fictionalism is coherent or that it is a route to nominalism or that it is compatible with this, that, or the other philosophical position. But to argue that it is true one has to argue that actual utterances involve a pretense, or that mathematical activity in fact involves a certain stance by its participants, or the like. These are empirical questions, ones that can be properly addressed in psychology or the social sciences. But defenses of fictionalism hardly even consider these.

I will suggest below that several questions that are traditionally considered part of the philosophy of mathematics are in fact solidly empirical questions. That doesn't mean that they should not be considered part of the philosophy of mathematics any more, but it should reflect on the methods with which they should be addressed.

5 Mathematical language

Traditionally, the role of formal tools comes through the analysis of mathematical language. The language of first-order arithmetic is taken to represent, or analyze, the language used in mathematical practice. And important features of the language of mathematical practice can be read off, and are made explicit by, the language of first-order arithmetic. This move, from the language used in mathematics to the language of first-order arithmetic, is so common that it is hardly made explicit. In the early days of our discipline, that is, after Frege, this was more explicitly discussed, but the transition has mostly moved to the background.

This is where the real damage is done in the use of formal tools in the philosophy of mathematics. Consider, for example, the language in which arithmetic is ordinarily carried out, and the language of first-order Peano Arithmetic, $\mathcal{L}_{PA}$. The $\mathcal{L}_{PA}$ here is an artificial language in standard first-order logic, with constant symbols, function symbols, quantifiers and variables, and so on.[3] The $\mathcal{L}_{PA}$ does about as badly as it could as a way to analyze the language of arithmetic. The $\mathcal{L}_{PA}$ gets almost everything wrong about the language in which arithmetic is ordinarily carried out. I will explain this in this section. But still, I will argue in the Section 6 that $\mathcal{L}_{PA}$ is supremely useful and good, even though it gets almost everything about the language of arithmetic wrong. What it is good at, and why it doesn't matter for this that it gets so much wrong, will be discussed after we see what it gets wrong.

5.1 Syntax

When we take $\mathcal{L}_{PA}$, the language of first order Peano Arithmetic, as an analysis of mathematical language, there is clearly a lot that it gets badly wrong. For example, in natural language, quantified noun phrases can be arguments of predicates, like transitive verbs:

(4) John kicked a man.

which has a syntactic argument structure like this

(5) kick(John, a man)

But in first-order logic only terms can be arguments of predicates, and quantifiers aren't terms. So, instead of the quantifier phrase "a man" we have to make a variable the argument, and put the quantifier somewhere else:

(6) $\exists x(kick(John, x) \wedge man(x))$

Similarly, the noun phrase (NP)–verb phrase (VP) structure of natural language sentences gets badly misrepresented in first-order languages. First-order languages do not split up sentences into NPs and VPs in the same way in which natural languages do. They both have a predicate–argument structure in general, but what is a predicate and what is an argument is quite different. More sophisticated formal theories of natural language have moved on, away from first-order languages, to ones more capable of representing NP–VP structure, predicate–argument structure, and so on.

This failing of first-order languages very much applies to first-order arithmetical theories. Consider a mathematical statements in natural language, for example,

(7) Every number is less than some number.

It has a a clear argument structure, splits up into a (quantified) noun phrase and a verb phrase, and so on. But the first-order statement in $\mathcal{L}_{PA}$

(8) $\forall x \exists y(x < y)$

completely misrepresents these aspects of the English sentence. Other formal languages do much better at getting that aspect of the syntax of English right, but such languages are in many other ways more complicated, and less suitable for what $\mathcal{L}_{PA}$ is good for. Even though $\mathcal{L}_{PA}$ does badly when it comes to the syntax of the language of mathematics, $\mathcal{L}_{PA}$ nonetheless is obviously a great formal tool. It gets something right, but not that.

To say this is not to dismiss formal language in general, in particular when it comes to understanding the syntax of natural languages. There is no question that there are formal languages that capture various syntactic aspects of natural languages well. But $\mathcal{L}_{PA}$ is not one of them. And, as we will see below, $\mathcal{L}_{PA}$ is really good for something, and those languages that capture the syntax of natural languages better aren't as good at that as $\mathcal{L}_{PA}$. The lesson so far simply is as follows: don't hope to learn about the syntax of ordinary talk about numbers from the syntax of $\mathcal{L}_{PA}$. The syntax of ordinary talk about numbers is studied in syntax, the sub-discipline of linguistics, and is then well or badly captured in various formal languages. But to find out what the syntax of ordinary talk about numbers is we have to look at linguistics. And to find out how well a certain formal language captures it, we have to first look at linguistics, and then compare what it tells us with the formal languages.

5.2 Number words

The $\mathcal{L}_{PA}$ treats number words as either constants (in the case of "0") or complex terms (in the case of all other number words). All terms in first-order logic pick out an object in any model, and so number words are treated as expressions that pick out objects, in other words as referring expressions. That number words are referring expressions is widely agreed upon among philosophers of mathematics. That is, that number words aim to refer is agreed upon by philosophers of mathematics. Whether or not they succeed in referring to anything is very controversial, at least in literal uses. But whether or not number words in natural language are referring expressions is far from clear. As Frege had already observed in his groundbreaking *Grundlagen* (1884), there is a puzzle about the uses of number words in natural language, which we can call *Frege's other Puzzle*,[4] unrelated to *Frege's Puzzle*, which is about about belief ascriptions and identity statements.[5] In natural language, number words can occur like adjectives, modifying nouns:

(9) Jupiter has four moons.

but apparently they can also occur like names or singular terms:

(10) The number of moons of Jupiter is four.

But, in general, the grammatical category of a singular term is different from that of an adjective. In general, a singular term cannot occur in the position of an adjective, and the other way round. How come number words can do this?

Although almost all philosophers of mathematics take number words to be referring expressions, like proper names, linguists usually focus on number words as they are used in completely different ways, for example, as parts of quantified noun phrases. And both groups have good reasons for their point of view. There is a real puzzle here.

I have argued in Hofweber (2005a, 2007b) that number words in natural language are not referring expressions. They are rather determiners that appear syntactically in the position of a singular terms for a variety of different reasons. Furthermore, number words in arithmetic are merely symbolic abbreviations of natural language number words, and thus they, too, are not referring expressions. In addition, arithmetic statements are literally true, although no reference is even attempted in them. Since there are different reasons why number words appear syntactically as singular terms, there are different cases to consider. The reason why number words appear as singular terms in examples like (10) is quite different from the reason why they appear as singular terms in arithmetical statements. I will briefly give those reasons in the following.

In Hofweber (2007b), I argued that the difference between (9) and (10) is one analogous to the difference between

(11) Otávio knows judo.

and

(12) Its Otávio who knows judo.

Both communicate the same information, but they do so with a different emphasis or focus. In addition, the focus effect is the result of the syntactic structure, not the result of special intonation as in

(13) OTÁVIO knows judo.

Similarly, the difference between (9) and (10) is one of a focus effect that is the result of the syntactic structure of (10). I have also argued that explanation for why there is a focus effect in (10) involves that the number word "four" in it is still a determiner, as it is in (9), but one that was displaced from its usual syntactic position. If this is correct, then "four" is not a referring expression in (10).

The more important case for the philosophy of mathematics is, however, a completely different one. I here disagree with Frege and neo-Fregeans, who take great inspiration from examples like (10) for the philosophy of arithmetic. The important case for arithmetic is the occurrence of number words in arithmetical statements, like

(14) $2+2=4$

In Hofweber (2005a), I have argued, on largely empirical grounds, that even there number words are not referential expressions, although such statements are literally true. One key to this argument is to notice that (14) when read out aloud is commonly read in two different kinds of ways, one corresponding to a plural reading, the other to a singular one:

(15) Two and two are four.
(16) Two and two is four.

and similarly using "plus", or "make". In Hofweber (2005a), I argued that (i) in the plural reading number words are "bare determiners", determiners without their noun argument present, and that (ii) the singular reading is the result of "cognitive type coercion", coercing the syntactic form of the plural statement into a different one for cognitive reason, thereby leaving the semantic function of the relevant expressions untouched. I won't repeat the details, but if it is correct, then even in arithmetical equations like (14) number words are still determiners, and not referential expressions. In addition, it follows that arithmetical statements are literally true, and their truth is independent of what there is, including how many things there are.[6]

My point for the rest of this chapter doesn't depend on that you believe any of this about the language of arithmetic. It is rather that there is a real issue whether or not $\mathcal{L}_{PA}$ correctly captures the semantic function of number words as picking out objects. And that this issue is one in the study of the language of mathematics, as it is actually carried out. This issue is not at all easily settled, and the rather simplistic pictures of natural language that were popular during the time when first-order languages ruled the study of natural language are long overthrown. What is important is to keep in mind that we need to settle this issue, in linguistics and the philosophy of language, and then, once we have settled it, we can see whether or not $\mathcal{L}_{PA}$ represents the semantic function of number words correctly. The hard work here is done somewhere else, in the study of natural language, or whatever the language of mathematics is. But unfortunately many philosophers are too quick to conclude from the facts that $\mathcal{L}_{PA}$ clearly gets something right (we will see shortly what) and that it gets many other things right as well, including what the semantic function of number words is. This is a great source of error in the philosophy of mathematics. The success of $\mathcal{L}_{PA}$ as the default formal language of arithmetic has more than once led to unjustified conclusions about features of number words in natural language from those in $\mathcal{L}_{PA}$. But what $\mathcal{L}_{PA}$ gets right does not indicate one way or another whether it gets the semantics of number words right. Before we can see what it gets right, lets look at one more thing that it gets wrong.

5.3 Quantifiers

The $\mathcal{L}_{PA}$ is a first order language, and as such it comes with a certain picture of quantification. On the one hand, quantification is given a certain syntactic treatment. Quantifiers are never directly arguments of predicates, for examples, contrary to ordinary English. We have already noted above that this, uncontroversially, misrepresents ordinary mathematical talk in its syntactic aspect. On the other hand, there is the semantics of quantifiers. Quantifiers in first-order logic are, standardly, given a model theoretic semantics, where they range over a domain of objects that are the objects that the terms in the language denote. This semantics of the quantifiers does fairly uncontroversially correspond to a reading that quantifiers have in ordinary natural languages, one where we make claims about whatever is out there in reality, our domain of quantification. But, I hold, this is not the only use that quantifiers have in natural language, and moreover, ordinary uses of quantification over natural numbers is not in accordance with this "domain" use of quantifiers. This can be argued for on the basis of considerations about natural language semantics and the particular uses of quantifiers over natural numbers, and it coheres with the picture of number words outlined above. I make my case for these claims in Hofweber (2005a, 2005b). Whether or not this is correct is not central in the following. But we should agree that whether or not $\mathcal{L}_{PA}$, as a first-order language, correctly captures certain aspects of ordinary quantification over numbers is an issue that is to be settled by first finding out how ordinary quantification over number works, then how it is represented in $\mathcal{L}_{PA}$, and finally to see whether $\mathcal{L}_{PA}$ represents it correctly. Quantification in natural language is much more complicated along a variety of dimensions than quantification in first or higher-order languages. It won't be too easy to figure this out, and we shouldn't draw any conclusions from how quantification works in $\mathcal{L}_{PA}$ to how it works in ordinary talk about numbers.

I hold that $\mathcal{L}_{PA}$ gets basically everything wrong about the syntax and semantics of talk about natural numbers. This, I take it, is uncontroversial for the question of the syntax, but controversial for the question of the semantics. But whether or not this is correct, the question what $\mathcal{L}_{PA}$ gets right and what it gets wrong has to be assessed via looking at the semantics of ordinary uses of quantifiers over natural numbers, and the semantic function of number words, and then to compare it to how these expressions are taken to be, semantically, in $\mathcal{L}_{PA}$. One should not draw any direct conclusions from the way they are represented in $\mathcal{L}_{PA}$ to how they are in ordinary talk about numbers.

Even though I think $\mathcal{L}_{PA}$ gets all this wrong, I do think that $\mathcal{L}_{PA}$ is great, and supremely useful, despite getting all this wrong. This is because $\mathcal{L}_{PA}$ gets something important right, and getting that right is independent of getting the syntax and semantics of number talk wrong. Here is what $\mathcal{L}_{PA}$ gets right.

6 Inferential relations

What we get right with the use of $\mathcal{L}_{PA}$ is capturing the inferential relations among the sentence of arithmetic. When we talk about numbers we use ordinary natural language. The sentences in our natural language stand in certain inferential relations to each other. Certain ones imply other ones, and there are discernible patterns of implication among them. We can properly capture those patterns of implication with the use of formal languages. First, we specify an artificial language, and some technical notion of consequence among sentences in this language. Then, we systematically assign sentences of our natural language that we use to talk about natural numbers to sentence in this artificial language such that inferential relations are exactly mirrored. Whenever an English sentence about natural numbers implies another one, in the ordinary sense of implication, then the corresponding sentence in the formal language implies, in the technical sense, the other corresponding sentence, and the other way round. This is what $\mathcal{L}_{PA}$ is good at. It manages to mirror the inferential relations among our ordinary sentences about natural numbers. And given the mathematical precision of the artificial language, it makes mathematical proofs about inferential relationships possible. Let me spell this out a little more.

Let E be a fragment of English, in particular the fragment consisting of the sentences employed in arithmetical reasoning. We take it that there is a notion of logical consequence that applies to sentences of E, call it "implies$_E$", and we take this notion as primitive for present purposes. Let L be an artificial language, say $\mathcal{L}_{PA}$. We assume that there is a defined notion of consequence that applies to the sentences of L, call it "implies$_L$". Let Φ be a systematic assignment of sentence of E onto sentences of L. We then say that Φ is *inferentially adequate* iff: for all sentences α and β of E, α implies$_E$ β iff $\Phi(\alpha)$ implies$_L$ $\Phi(\beta)$. Thus, an assignment of sentences in a fragment of a natural language to an artificial language is inferentially adequate just in case it exactly mirrors the inferential relations among the groups of sentences. We will also call an artificial language inferentially adequate if the "intended" or "standard" assignment is inferentially adequate. And it turns out that if we restrict ourself to certain simple sentences of English (ones without "finitely many" in it, or the like), then there is a inferentially adequate assignment of these sentences into $\mathcal{L}_{PA}$. How the assignment goes is not precisely specified, although I presume it could be. We do this on the fly, but in a systematic way. In general, we do it by rephrasing our English sentences so that they very closely mirror sentences of $\mathcal{L}_{PA}$, and assign them such sentences this way. It would be a mistake, though, to think that we simply assign English sentences ones in the formal language with the same truth conditions. Such sentence don't have truth conditions, by themselves. They are merely true in certain models, but this is secondary. Note that whether or not an assignment is inferentially adequate is not subject to mathematical proof since it relates to questions about the primitive notion of implication, which is not fully precise. But given presumed inferential adequacy, mathematical proofs can be carried out about what follows from what.[7]

An assignment of sentences of a formal language to those in a natural language can be inferentially adequate without getting the syntax of the sentences in the natural language right, and without getting the semantic features right. All that is required is that whatever features are gotten wrong are gotten wrong systematically, so that all errors cancel themselves out in the end. This is what makes $\mathcal{L}_{PA}$ so great. It is inferentially adequate for a very important fragment of ordinary talk about natural numbers. Of course, the fact that $\mathcal{L}_{PA}$ is inferentially adequate might be taken as evidence that it does get the syntax and semantics right as well. After all, wouldn't it be the easiest explanation why it gets the inferential relations right that it gets everything else right as well? There is something to be said for this line, and I think that it gets the inferential relations right is *some* evidence that it gets the rest right. But the other evidence, that it gets the syntax and semantics wrong, is much stronger. I take the case of the syntax to be obvious. The syntax of natural language even when restricted to simple talk about numbers is not the same as the syntax of $\mathcal{L}_{PA}$. The evidence for this is conclusive. Nonetheless, $\mathcal{L}_{PA}$ gets the inferential relations right. And we can see quite easily how that can be. It captures enough of the inferentially relevant parts of that language in a way that inferential relations can be mirrored, while the syntax and semantics is different.

Since $\mathcal{L}_{PA}$ is a precise formal language, and the relevant notion of consequence is a precise notion, this gives rise to the possibility of precise proofs about what follows from what. And given inferential adequacy, this will imply that the corresponding inferential relations hold between the corresponding English sentences. Here formal tools can do real work. And they can do all this work even if, as I claimed above, $\mathcal{L}_{PA}$ gets basically everything wrong about the syntax and semantics of the language of arithmetic. And to carry out proofs about inferential relationships, formal tools do not play a secondary role, they are essential. Without precision in the notion of consequence and the language to which it applies no such proof would be possible. Here, there is hope for a positive and substantial role for formal tools.

Our question about the role of formal tools in the large-scale philosophy of mathematics is the following: what role do proofs about inferential relations have in this endeavor? Why do proofs about inferential relations matter for the large-scale philosophical questions? To answer this question, we will have to look at two other questions: the question about the status of axioms, and the question about whether all of mathematics is philosophically equal. I will argue in the following section that if one holds a certain position on these two other questions, then formal tools can have such a substantial role. However, most philosophers don't have that view about the other two questions. (I hold this view, but I won't be able to defend it here.)

7 The status of axioms

Axioms are special mathematical statements. They are just like other mathematical statements in many ways: they are in the same language, involving the same mathematical expressions, and so on. But they are also, somehow, special, and distinct from other mathematical statements in the very same part of mathematics. In what sense they are distinct from other statements in the same part of mathematics is a major dividing line for how to understand that mathematical discipline. We can distinguish the *descriptive* from the *constitutive* conception of (particular) axioms. On the descriptive conception, axioms are special, but they are mostly special for us. Among all the true statements in that part of mathematics, axioms are ones that are especially compelling to us, and that are especially useful for deriving consequences. They might be a nice way to put together a group of statements that allow us to derive a lot of other ones, and that are each fairly simple, in particular for us. But there is otherwise no special connection between what is true in that part of mathematics and what is an axiom. Axioms are just some especially compelling and systematizing true statements. On the constitutive conception of axioms, on the other hand, there is such a connection. Here the axioms, somehow, determine what is the case in this domain. The axioms are the basis for all truth in the domain. All truth flows from the axioms. On the constitutive conception of axioms for a particular part of mathematics, axioms are not just especially compelling true statements, they are the source of all truth in that domain. Both conceptions of axioms hold that axioms are true statements, but they differ in their assessment of priority. For descriptive axioms the mathematical facts are prior to the axioms, whereas on the constitutive conception of axioms it is the axioms that are prior to the facts. This difference seems intuitive, but it is hard to spell it out more precisely. In fact, whether or not it can be spelled out properly is an open question, but I will count on it being clear enough for the following.

Whether axioms can ever be constitutive is problematic. How could it be that all truth flows from the axioms in the sense, for example, that what is true in that domain of mathematics is just what is implied by the axioms? After all, the axioms imply themselves, and so how could they guarantee their own truth? But we can get a sense how this could be by seeing an analogy with fiction (although constitutive axioms might not require an association with fictionalism in the philosophy of mathematics). What is true in a fiction flows from (in part at least) what is implied by the story. The story implies itself, and so it guarantees its own truth. Not literal truth, of course, but truth in the relevant sense. Even though the difference between these two conceptions of axioms is hard to make precise, and even though it isn't clear whether axioms can ever be constitutive, the difference is compelling enough to appear in many places in the philosophy of mathematics.[8] Let's elaborate on this with some examples.

It is plausible that the Peano Axioms are not constitutive for arithmetic. By the Peano Axioms, I simply mean the particular arithmetical statements that are these axioms, not the formal axiom system of PA. The Peano Axioms are thus ordinary mathematical statements, not ones in an artificial language. One reason for thinking that the Peano Axioms are not constitutive for arithmetic is the independence of some Π^0_1 statements $\forall x \Phi(x)$ where every instance $\Phi(n)$ can be proven from the Peano Axioms. Even though the universally quantified statement can't be proven from the axioms, it nonetheless is true, and so truth goes beyond what can be proven from the axioms. This argument of course assumes that constitutive axioms only determine facts that can be proven from them, which might not be so if we understand what facts axioms determine in a different way, for example, by bringing in second order logic. I would be hesitant to do so, but this is certainly debatable. All I am claiming is that it is plausible that the Peano Axioms are not constitutive for arithmetic.

A further reason against the Peano Axioms being constitutive for arithmetic requires a close reflection on what makes any axioms constitutive, which we won't be able to give here, but which we can nonetheless appreciate in outline. Being constitutive for a certain domain of truths is not the same as implying all these truths. Many subsets of all the truths will imply all the truths, but that does not make them constitutive axioms of the domain. More is required for the axioms to be constitutive. The constitutive axioms are those that play some role in that domain of mathematics that, explicitly or implicitly, are involved in its conception. And this seems to require that the axioms can't be an afterthought, something discovered long after the domain has been established as an active part of mathematics. Peano's Axioms seem more like a discovery after the fact, rather than something involved at the beginning. But again, all I am saying is that it is plausible that the Peano Axioms are not constitutive for arithmetic.

One tempting example of constitutive axioms are algebraic axioms, for example, the axioms of group theory. What it is to be a group is fully determined by these axioms since the notion of a group is defined by these axioms. However, the results of *group theory*, the mathematical discipline, do not have the group axioms as constitutive axioms. That there are any groups at all is not settled by the group axioms, nor do the group axioms alone imply almost any of the results of group theory. What is needed in addition is the rest of mathematics, or at least the relevant parts of it. So, when one wants to show that there is an infinite cyclic group one does this by pointing out that the integers under addition have all the required features. But that they do have these features is not settled by the group axioms. Whether algebraic axioms, axioms like the ones of group theory, are good examples of constitutive axioms for a mathematical discipline, or merely useful definitions, doesn't have to be settled here. Algebraic axioms are often brought up as examples of constitutive ones, and if they help to illustrate the difference, even if they in the end don't fall on the relevant side of the difference, then that should be enough for us now.

If there is a part of mathematics that has constitutive axioms, then formal tools will have a greater role in the large-scale philosophical project than if this is never true. If truth in a part of mathematics is determined by what follows from the axioms, since the axioms are constitutive for that part, then investigating what does and doesn't follow from the axioms is tied to investigating where the facts are and where they end. To find out, via a proof of an inferentially adequate formal language, that a certain statement is independent of the constitutive axioms, would shed quite a bit of light on the factual status of that claim. It would show that the statement is outside of the realm of the facts.

But if the axioms are descriptive, then nothing should be taken to follow from the fact that a certain statement is independent of the axioms. What is true in that domain is not in any special way tied to the axioms, and so no conclusion should be drawn from the statement having or not having a certain relation to the axioms. The status of the axioms is crucial for the consequences of certain results about what follows from what.[9]

8 An epistemic role?

Might formal tools not have an epistemic role even if axioms are descriptive? For example, couldn't they help us to acquire knowledge about what is or isn't true in a certain domain? No, not in the relevant way. A proof that something follows from the axioms when they are descriptive shows that this statement is true, just like the axioms, but this we could have established by just proving the statement from the axioms directly. Formal tools might aid in such proofs, but this is not the substantial role we were looking for. But what about a proof of the independence of a statement from the axioms? This also does not have as great an epistemic significance as sometimes claimed, if the axioms are descriptive. Nothing could be concluded about there being no fact of the matter with regard to that independent statement, nor can it be concluded that we can't know the answer. The axioms are not our sole source of evidence when they are descriptive. They are merely one way for us to write up what compels us in a systematic way.

No assumption can be made that the axioms are an exhaustive source of evidence. The epistemic role of formal tools so understood is nothing to write home about. If the axioms are constitutive, then things are different, in the way discussed above.

9 Pluralism

The overall role of formal tools is further tied to another question, which is itself a central question in the large-scale philosophy of mathematics. It is the question whether or not all of mathematics is philosophically equal. This question has an innocent reading, where the answer is clearly: no. Some parts of mathematics are simple, known with more certainty, and so on. There clearly are differences among different parts of mathematics that extend to their philosophical assessment. But the question is intended in a grander scheme: Among the large-scale philosophical pictures that we philosophers have of mathematics, does one apply to all of mathematics, or do different ones apply to different parts of mathematics? Let's call those who think that all of mathematics is equal with respect to the large-scale philosophical picture *monists*, and those who think that it is not *pluralists*. Most philosophers, I think it is fair to say, are monists. Usually, fictionalists are fictionalists about all of mathematics, and platonists or structuralists are platonists or structuralists about all of mathematics. A philosophy of mathematics is often conceived as a philosophical story of mathematics as a whole, intended to fit all the parts equally. But there are some prominent exceptions. Notably, Hilbert thought that some part of mathematics, the finitist part, was in large-scale philosophical ways different from the rest of mathematics. Similarly, contemporary predicativists, like Solomon Feferman,[10] hold that predicative mathematics is in important philosophical ways different from other non-predicative parts of mathematics. Those who are pluralists thus hold that at least two parts of mathematics are different in important philosophical ways, and they thus have to say (i) how they are different and (ii) how they nonetheless come together into one discipline: mathematics. The pluralists mentioned so far are *dualists*, they hold that from a large-scale philosophical point of view there is one crucial difference which divides mathematics into two parts: the finitist versus the rest, or the predicative versus the rest. But in principle, pluralists could hold that there are a variety of different large-scale philosophical accounts of different parts of mathematics.

One way to be a pluralist is to hold that for some parts of mathematics axioms are descriptive, while for others they are constitutive. This marks a large-scale difference between these two parts of mathematics, and it is one of the ways in which formal tools can have an even greater role. The issue is the following: If there is a part of mathematics that has constitutive axioms, then the question of the consistency of these axioms becomes of central importance. To see this, consider what happens when there are inconsistent axioms in a part of mathematics where the axioms are descriptive, and thus not constitutive. Here, the axioms are merely a simple way for us to systematize the truths in that part of mathematics. They might be epistemically compelling, especially powerful in allowing us to derive all kinds of important results from them, and so on. But if they are inconsistent, then this should not be taken to be the end of that part of mathematics. It just shows that we did badly in finding basic principles from which the statements that are true in that part follow. And it might show that what we take to be compelling in this field isn't even all true.[11] If the axioms are inconsistent, then we just have to find better ones. The domain of mathematical truths that the axioms where supposed to capture is unaffected by this. Our conception of that domain might be in trouble, but the domain is not. Inconsistent axioms reflect badly on us, not what they are about.

If the axioms are constitutive, then things are different. The axioms are constitutive for what is true in that domain. If the axioms themselves are inconsistent, then the domain itself is in deep trouble. Constitutive axioms can't be separated from the domain they are about in the way that descriptive axioms can be. If constitutive axioms are inconsistent, then the domain that they were supposed to be constitutive for goes down the drain. Without consistency no meaningful domain has been constituted by the axioms. The question of the consistency of descriptive axioms is of interest, but the question of the consistency of constitutive axioms is a matter of survival.

Given that we hold that formal tools like $\mathcal{L}_{PA}$ get the inferential relations of ordinary mathematical sentences right, it follows that the consistency of an ordinary mathematical theory can properly be captured by a statement about artificial languages, like $\mathcal{L}_{PA}$. A set of axioms, in ordinary mathematical language, is consistent just in case not everything follows from it. This is a claim about what follows from what, and it thus is properly captured with a claim about what follows from what in a formal language. And this claim, again, can be stated with mathematical precision since it is a claim about a precise language and a precise notion of consequence. And since it can be stated with mathematical precision, it in effect becomes a mathematical statement itself. It can be stated in the language of arithmetic, and it can be stated in the language of PA, as the formal consistency statement of PA, CON_{PA}.[12] This is central for our main issue here, for the following reason.

Suppose, we have a part of mathematics with constitutive axioms. Then, the consistency of these axioms is of large-scale philosophical importance since without it there is no meaningful domain at all in this discipline. No facts are determined by the axioms since no facts can be inconsistent. If the axioms are constitutive, then it is central that they are consistent. But the consistency claim of the axioms is, in effect, itself a mathematical claim. To show that the axioms are consistent is to show that a certain mathematical claim is true. But what about the part of mathematics in which this is supposed to be shown? Does it have constitutive axioms as well? If it does, then the proof of the consistency of the first axioms would depend on the consistency of the second ones in a way that is too close for comfort. If the constitutive axioms of the part of mathematics in which the claim of consistency is stated are inconsistent, then that very claim of consistency loses all content. It would be a mathematical claim in an ill-defined domain. Thus, if these axioms are inconsistent then this would not only put the proof of consistency in question (which it would whether or not the axioms are constitutive), but it would put the whole statement of consistency itself in question. It is a mathematical claim after all, and it would be a mathematical claim in a part of mathematics where the axioms are supposed to be constitutive, but are inconsistent. The statement of the consistency of the axioms itself would be a statement in a domain that has not been properly established, and thus the statement of consistency itself would be most dubious. It would be not just dubious with respect to its truth value, but with respect to whether a coherent statement has been made at all. The status of the axioms of arithmetic (and finite set theory, and similar parts of mathematics) thus has a special standing among all other mathematical disciplines, philosophically. Only if these axioms are not constitutive can we deal with constitutive axioms in other parts of mathematics without things turning out badly. If the axioms of arithmetic are constitutive, then the claim of their consistency is only fully meaningful if they are consistent. This does not give us any ground to stand on.

But what if this is true: suppose the axioms of the parts of mathematics where consistency claims live, arithmetic and the like, are not constitutive. Instead, there is an objective domain of mathematical facts, somehow, and it is not constitutively tied to any axioms. Suppose further that some other part of mathematics does have constitutive axioms. Then the claim of the consistency of these axioms is well-grounded since it is a claim in a domain of mathematics where we have established facts. And a proof in arithmetic of the consistency of the axioms of the other part of mathematics has large-scale philosophical implications. It would show that there is a meaningful domain of mathematical facts in this other part. And the proof that this is so would not itself be subject to the further question if the domain of mathematical facts relied on to prove it is itself legitimate. If such a scenario obtains, then formal tools will have a large-scale philosophical role, even though only a limited one. They would be required to show that a certain domain of mathematics is well-grounded and meaningful. Inconsistent constitutive axioms destroy their domain. A proof of their consistency establishes it as a legitimate part of mathematics.

It turns out that the view about arithmetic referred to above, and defended in Hofweber (2005a) very much fits into this picture. Arithmetic axioms are not constitutive, and there is an objective domain of arithmetical facts. However, the reason why this is so does not carry over to other, higher, parts of mathematics. Whatever the large-scale philosophical story is about them, it will have to be fundamentally different.

For formal tools to ever have this greater role it must be true that

1. Arithmetic, or a similar part of mathematics, is about a domain of mathematical facts, but does not have constitutive axioms.
2. Some other part of mathematics does have constitutive axioms.

This would give a role to the statement of consistency, which can be seen as an arithmetical statement and so part of the domain of facts. The role of formal tools is even greater if this is true:

3. You can prove in arithmetic that (at least a sufficiently large part of) the axioms of the other part are consistent.

I don't know how many philosophers believe all these to be the case, but those who do can maintain that formal tools have a greater role than those who deny it. I believe this is indeed true. The view of arithmetic defended in Hofweber (2005a) has the required features, in particular there is a domain of arithmetical facts that is in no way constitutively tied to any axioms. Which parts of mathematics have constitutive axioms is a difficult question, but I would put money on that some do. However, I can't hope to give a defensible candidate here, nor can I hope to give a defensible account of what constitutive axioms are more precisely.

The third requirement, a consistency proof in arithmetic, might seem to be the most dubious. Hasn't Gödel shown that this is hopeless? Yes, and no. What matters is an informal proof of consistency, and in it we are certainly entitled to rely on the consistency of, for example, PA. So, all that is necessary is to give a relative consistency proof in PA of a certain set of axioms, which we can then take to sufficiently establish the consistency of the axioms. A relative consistency proof can often be given for axioms that appear to be substantially stronger than arithmetic. A series of results along these lines have been established in "reductive" proof theory, using the technique of proof theoretic reduction.[13] But even if it is not possible to give a relative consistency proof of all of the axioms, it might be possible to give such a proof for a sufficiently large part of the axioms, a part large enough to be taken to establish a core of the domain of mathematics in question. Some of the constitutive axioms might be taken to form the core of the domain, while others might be constitutive for the particular way to develop the core, but also optional in that the core might be developed along other lines. If this is true, then a consistency proof of the core would still be of large-scale importance since it at least established the core. Thus, consistency proofs and proof theoretic reductions of subsystems of a certain set of axioms are also of large-scale importance on this line, although it remains to be seen what counts as the core and what is merely a development of the core.

10 Conclusion

Formal tools are easily a source of error, and in fact have been such a source. When it comes to the understanding of mathematical activity and mathematical language, formal tools merely have a secondary, representational role. They are at best used to represent results established by other means. This is particularly vivid and neglected when it comes to understanding the semantic function of number words, where in the philosophy of mathematics they are almost uniformly and without much discussion understood as referring expressions, in analogy with their treatment in formal theories of arithmetic. And all this even though there is much evidence that in ordinary discourse, including in mathematical discourse, they have a different semantic function. Whatever their function is in the end, this has to be established by empirical means, in the study of mind and language. Formal tools can then be used to represent the result, but they shouldn't inspire us to expect one result or another.

However, there are two questions that determine how big the role of formal tools is in the large-scale philosophy of mathematics. First, the question about the status of axioms, second the question about pluralism. In particular, if one holds that axioms are sometimes constitutive, but they are not for arithmetic, then formal tools will have a much larger, but still limited, role. The real questions will still have to be addressed with other means. Whether the axioms are constitutive is not a question that is to be settled with formal tools, nor is the question if all of mathematics is in relevant ways the same. These will be the more central and harder questions. And similarly, for many other of the central questions in the philosophy of mathematics, formal tools will play little or no role.

Even though our main example in the discussion of these claims was arithmetic, the same, I maintain, carries over to other parts of mathematics, like set theory. My claims that $\mathcal{L}_{PA}$ gets basically everything wrong about the language of arithmetic (except inferential relations) does, of course, not carry over to the language of ordinary set theory and $\mathcal{L}_{ZF}$.[14] Even here there are real issues about how much $\mathcal{L}_{ZF}$ gets right about the language of set theory, as it is ordinarily carried out, informally with the use of symbols. Whether or not expressions for sets are referential is an issue that is less obvious than one might think. Here, real work needs to be done in the study of mathematical language.

So, for the future, I hope that the field of the philosophy of mathematics will have a clearer sense of with what methods which questions should be addressed. In particular, I hope that we will have a clearer sense of which questions should be addressed with formal tools and why. Some of the questions we are dealing with are empirical questions, some are technical questions, and some are philosophical questions. Depending on where a particular question falls will have a great effect on how we can hope to answer it. Once it is clear what kind of question it is, we can see with what method it should be addressed. To see what kind of question we are dealing with is often hard. Sometimes real progress can be made by showing that a particular question can be understood as a technical question. The history of logic has given us a few such cases, but the success of these few cases has lead to treating many other questions as being too closely tied to the formal tools that were successfully used in those cases. If we keep these issues apart, then things will be better. We won't be anywhere near done, but things will be better.[15]

Notes

1. Although the language of Peano Arithmetic does not contain "4", it is commonly introduced as a new term with the definition "4 = S(S(S(S(0))))".
2. See Hilbert (1925). The philosophical significance of Hilbert's use of formal tools and their development in contemporary "reductive proof theory" is discussed in Hofweber (2000).
3. I will use "PA" for the axiom system, and $\mathcal{L}_{PA}$ for the language in which these axioms are given. The $\mathcal{L}_{PA}$ is a standard first-order language, that is, no free logic, no ϵ terms, and so on. These modifications soften some of the blows mentioned below, but don't avoid the general issue that the success of $\mathcal{L}_{PA}$ is independent of getting all these other things right or wrong. It will not matter for my points below just what precise the basic vocabulary of $\mathcal{L}_{PA}$ is taken to. We can include "$<$" or not, and either just have a constant for 0, or constants for both 0 and 1, or constants for all numbers. My main points merely rely on $\mathcal{L}_{PA}$ being a standard first-order language.
4. See Hofweber (2005a).
5. See Salmon (1986).
6. A different picture of number words is presented in Øystein Linnebo's contribution to this volume.
7. I am neglecting in this section that strictly speaking I don't believe that the formal notion of consequence used in $\mathcal{L}_{PA}$ is adequate for the informal notion of consequence. This has nothing to do with the shortcomings of first-order logic, but with another issue in the philosophy of logic which right now would just sidetrack us. See Hofweber (2007a, 2008). I am also putting aside those who hold that there is no clear enough set of facts about what implies what in terms of ordinary implication among English sentences. See, for example, Resnik (1997, 161 ff.) for such a stance toward facts about what follows from what.
8. The difference between descriptive and constitutive conceptions of axioms is related to Hellman's (2003) distinction between "assertory" and "algebraic" conceptions of mathematics, which is discussed in Mary Leng's article in this volume.
9. I am indebted to an anonymous referee for several helpful suggestions on this section.

10. See, for example, the essays in Feferman (1998).
11. I am assuming here that "axioms" in this sense, in which they can turn out to be inconsistent, don't have to be true. I am not taking it to be an option that mathematics itself might be inconsistent. For inconsistent mathematics, the same issue will arise for axioms to be non-trivial.
12. I am skipping here the issue of which one of various candidates should be seen to properly capture the consistency statement. Not because there isn't something important here, but just because that is not our main issue. I am simply assuming that it can be properly captured. See Feferman (1960).
13. See Feferman (1988) for a survey of such results.
14. $\mathcal{L}_{ZF}$ is the language of formal ZF set theory. Some parts of set theory are carried out in this language directly, and in those no mismatch can occur. But others are carried out in ordinary language, with the use of symbols, and here there is a real question about what features are correctly represented in $\mathcal{L}_{ZF}$.
15. Thanks to Øystein Linnebo, Mike Resnik, and an anonymous referee for many helpful comments on an earlier draft. A predecessor of this paper was presented at Stanford at a conference in honor of Solomon Feferman, on the occasion of his retirement. As one of his students, I would like to once more express my gratitude for his teaching, guidance, and support.

References

Feferman, S. (1960). Arithmetization of metamathematics in a general setting. *Fundamenta Mathematicae*, 49:35–92.

Feferman, S. (1988). Hilbert's program relativized. *Journal of Symbolic Logic*, 53: 364–384.

Feferman, S. (1998). *In the Light of Logic*. Oxford University Press, Oxford, UK.

Frege, G. (1884). *Die Grundlagen der Arithmetik: eine logisch mathematische Untersuchung über den Begriff der Zahl*. W. Koebner.

Hellman, G. (2003). Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 11(2):129–157.

Hilbert, D. (1925). Über das Unendliche. *Mathematische Annalen*, 95:161–190.

Hofweber, T. (2000). Proof-theoretic reduction as a philosopher's tool. *Erkenntnis*, 53:127–146.

Hofweber, T. (2005a). Number determiners, numbers, and arithmetic. *The Philosophical Review*, 114(2):179–225.

Hofweber, T. (2005b). A puzzle about ontology. *Noûs*, 39:256–283.

Hofweber, T. (2007a). The ideal of deductive logic. Unpublished manuscript.

Hofweber, T. (2007b). Innocent statements and their metaphysically loaded counterparts. *Philosophers' Imprint*, 7(1):1–32.

Hofweber, T. (2008). Validity, paradox, and the ideal of deductive logic. In Beall, J. (ed.), *Revenge of the Liar: New Essays on the Paradox*. Oxford University Press, Oxford, UK.

Resnik, M. D. (1997). *Mathematics as a Science of Patterns*. Oxford University Press.

Salmon, N. (1986). *Frege's Puzzle*. MIT Press, Cambridge, MA.