Towards a Trivialist Account of Mathematics

Agustín Rayo
MIT

March 12, 2009

The aim of this paper is to defend mathematical trivialism—the view that the truths of pure mathematics have trivial truth-conditions and the falsities of pure mathematics have trivial falsity-conditions.

I doubt there can be an easy argument for trivialism, for two reasons. The first is that the debate got off to a bad start. Discussions in the philosophy of mathematics tend to presuppose a certain conception of the conceptual landscape that makes little room for trivialism. I think this conception is mistaken, and that once it is set aside trivialism can be seen to be a plausible position. But old habits die hard. The second reason is that trivialists face an important challenge. They need to explain what the point of mathematical knowledge could be if mathematics deals with trivialities. I think there is a good answer to the challenge, and that the resulting picture of mathematical knowledge is independently attractive. But it is a picture that is not easy to set up, and is unlikely to seem compelling at first.

When Otávio and Øystein invited me to participate in this volume, I was faced with a choice. I could write a paper outlining the entire case for trivialism in very broad strokes, or I could focus on some particular portion of the argument, discussing it in detail at the expense of the rest of the material. In the end I decided to do the former. The result is a

For their many helpful comments I would like to thank Ross Cameron, Roy Cook, Matti Eklund, Caspar Hare, John Heil, Øystein Linnebo, Alejandro Pérez Carballo, Brad Skow and an anonymous reviewer for this volume.

paper that is somewhat impressionistic. Although many of the omissions are addressed in other work, you will find that crucial moves are made with insufficient discussion and that arguments are often sketchy. Please bear with me. In return, I will try to give you a sense of the 'big picture' underlying the combination of views I defend.

1 Against Conventional Wisdom

Platonists and nominalists disagree about *ontology*: Platonism is the view that there are mathematical objects; nominalism is the view that there aren't any. Committalists and noncommittalists, in contrast, disagree about the *truth-conditions* of mathematical statements: committalism is the view that everyday mathematical assertions carry commitment to mathematical objects; noncommittalism is the view that they don't. These two distinctions yield a four-fold partition of logical space, but the two most popular positions are Platonism+committalism and nominalism+noncommittalism. *Error-theories* (i.e. nominalism+committalism) and *irrelevance-theories* (i.e. Platonism+noncommittalism) are both consistent, but they are not as well represented in the literature. (A notable exception is the error-theory espoused in Field (1980).)

Conventional wisdom has it that each of the two most popular positions has an advantage over the other: nominalism+noncommittalism does a better job of accounting for mathematical knowledge, since it doesn't need to explain how we could have knowledge of the abstract realm; but Platonism+committalism does a better job accounting for mathematical discourse, since it doesn't need to postulate a non-standard semantics (or a less-than-straightforward connection between what is communicated by an assertion and the truth-conditions of the sentence asserted).

It seems to me that this is an unhelpful picture of the terrain. It focuses on the wrong distinction when it comes to epistemology, and it relies on a questionable conception of the way language works when it comes to mathematical discourse. I shall discuss each of these points in turn.

Say that a sentence has *trivial* truth-conditions if any scenario in which the truth-conditions fail to be satisfied would be unintelligible. (More on the relevant notion of intelligibility below.) We have no trouble making sense of a scenario in which there are no elephants, so we should take 'there are elephants' to have *non*-trivial truth-conditions. But (most of us) are unable to make sense of a scenario in which something fails to be self-identical. So we should take a logical truth like '$\forall x(x = x)$' to have *trivial* truth-conditions. (Some dialetheists would disagree.)

Trivialism is the view that true sentences of pure mathematics have trivial truth-conditions (and that false sentences of pure mathematics have trivial falsity-conditions). According to the trivialist, *nothing* is required of the world in order for the truth-conditions of a mathematical truth to be satisfied: there is no intelligible possibility that the world would need to steer clear of in order to cooperate with the demands of mathematical truth. This means, in particular, that there is no need to go to the world to check whether any requirements have been met in order to determine whether a given mathematical truth is true. So once one gets clear about the sentence's truth conditions—clear enough to know that they are trivial—one has done all that needs to be done to establish the sentence's truth. (Keep in mind that getting clear about the truth-conditions of a given mathematical sentence can be highly non-trivial. So determining whether the sentence is true is not, in general, a trivial affair.)

For the trivialist, our knowledge of pure mathematics can be understood on the model of our knowledge of pure logic. (More on this below.) The non-trivialist, on the other hand, owes us an account of what is required of the world in order for the truth-conditions of a given mathematical truth to be satisfied, and an explanation of how we might be in a position to check whether the relevant requirement has been met. So when it comes to the task of accounting for mathematical knowledge, the trivialist has an advantage over the non-trivialist. But it is important to note—and this is where conventional wisdom proves unhelpful—that the distinction between trivialism and non-trivialism *cuts across* the distinction between nominalism+noncommittalism and Platonism+committalism. In particular, one can be a nominalist and a noncommittalist without being a trivialist, and one can be a Platonist and a committalist while being a trivialist.

Suppose, for example, that one is a nominalist and embraces a version of noncommittalism whereby the truth-conditions of a sentence of pure mathematics are given by its *universal Ramseyfication*. (If ϕ is an arithmetical sentence, its universal Ramseyfication is the universal closure of $\ulcorner(\mathscr{A} \rightarrow \phi)^{*}\urcorner$, where $\mathscr{A}$ is the conjunction of a suitable list of axioms and ψ^{*} is the result of uniformly substituting variables for mathematical vocabulary in ψ.) Then as long as one is able to make sense of a finite world, one will take oneself to be a non-trivialist. To see this, consider an arbitrary arithmetical falsehood, $\mathscr{F}$. Since $\mathscr{A}$ can only be true if there are infinitely many objects, the universal closure of $\ulcorner(\mathscr{A} \rightarrow \mathscr{F})^{*}\urcorner$ can only be false if there are infinitely many objects. So the falsity-conditions of $\mathscr{F}$ will fail to be satisfied if the world is finite, and are therefore non-trivial.[1] Of course, a non-trivialist non-committalist might have a story to tell about how it is that the truth-conditions of mathematical truths (and the falsity-conditions of mathematical falsehoods) can be known to be satisfied, even though they are non-trivial. The Universal Ramseyfier, for example, might have a story to tell about why it is that we're entitled to the assumption that there are infinitely many things. The point is that, unlike the trivialist, she needs a story to tell—non-committalism does not, by itself, deliver epistemological innocence.

We have seen that one can be a nominalist and a noncommittalist without being a

[1]Other versions of non-trivialist non-committalism include Hodes (1984), Hodes (1990), Fine (2002) II.5, Rayo (2002) and Yablo (2002), as well as Bueno and Leng's contributions to this volume. The proposal in McGee (1993) may or may not be interpreted as noncommittalist, but it is certainly non-trivialist. An example of trivialist noncommittalism is Hofweber (2005). Assessing modal versions of non-committalism can be tricky—see my 'On Specifying Truth-conditions' for discussion.

trivialist. But one can also be Platonist and a committalist while being a trivialist. Traditional Platonists think that even though numbers exist, it is intelligible that they not exist. *Subtle* Platonists maintain, in contrast, that for the number of the Fs to be n *just is* for there to be n Fs. Accordingly, the view that there are no numbers is not just false, but unintelligible. ("Suppose there are no numbers. For the number of Fs to be 0 *just is* for there to be no Fs. So the number 0 must exist after all!"). This means that even if one is a committalist one should think that *nothing* is required of the world in order for the truth-conditions of 'there are numbers' to be satisfied. For there is no intelligible possibility that the world needs to steer clear of. So there is room for a committalist who is also a subtle Platonist to be a trivialist. (Each of the following can be interpreted as defending a version of subtle Platonism: Frege (1884), Parsons (1983), Wright (1983), Stalnaker (1996) and Linnebo's contribution to this volume.)

Moral: if one is concerned with mathematical knowledge, the most interesting place to look is not the contrast between nominalism+noncommittalism and Platonism+committalism. It is the contrast between trivialism and non-trivialism.

I have argued that conventional wisdom delivers a potentially misleading picture of the epistemological terrain. I would now like to explain why I think it delivers a misguided picture of the linguistic terrain. According to conventional wisdom, Platonism+committalism does a better job than nominalism+noncommittalism when it comes to accounting for mathematical discourse. For consider a mathematical sentence such as '2 is prime'. Proponents of Platonism+committalism can take '2' to refer to a particular object—the number 2—and claim that the sentence is true just in case that object has the property expressed by 'is prime'. But proponents of nominalism+noncommittalism think there are no numbers. So they lack a straightforward way of specifying truth-conditions for '2 is prime'. They must either claim that the logical structure of mathematical statements shouldn't be taken at face value, or claim that the information conveyed by mathematical assertions is very different from what the sentences asserted literally say.

The problem with this way of approaching the issue is that it is based on a questionable picture of the workings of language: the idea that there is a certain kind of correspondence between the structure of language and the structure of the world. More specifically, what is presupposed is this: (1) there is a particular carving of the world into objects which is more apt, in some metaphysical sense, than any potential rival—a carving that is in accord with the world's true 'metaphysical structure'; (2) to each legitimate singular term there must correspond an object carved out by the world's metaphysical structure; and (3) satisfaction of the truth-conditions of an atomic sentence requires that the objects corresponding to singular terms in the sentence bear the property expressed by the sentence's predicate. (Should one be a deflationist about properties, and claim that for an object to have the property of Fness *just is* for the object to be F? Or should one admit properties as separate items in one's ontology? Different versions of the view will address the issue in different ways.)

This conception of language is a close cousin of the 'picture theory' that Wittgenstein defended in the *Tractatus*.[2] And it seems to me that it ought to be rejected for just the reason Wittgenstein rejected the picture theory in his later writings. Namely: if one looks at the way language is actually used, one sees that usage is not beholden to the constraint that an atomic sentence can only be true if its logical structure is in suitable correspondence with the structure of the world.

Assertions are tools for communication. Suppose you are organizing a dinner party and are thinking about seating arrangements. I say 'There will be an odd number of people at the table'. In doing so, my objective is not to represent the structure of reality as somehow corresponding to the logical structure of the sentence I uttered. In particular, I do not mean to commit myself to a non-trivial ontological thesis about numbers, and go on to

[2]Here I have in mind a traditionalist interpretation of the *Tractatus*, as in Hacker (1986) and Pears (1987). See, however, Goldfarb (1997).

represent numbers as bearing a certain relation to people and the table. All I want to do is help you discriminate amongst the possibilities at hand. Suppose, for example, that you are trying to decide between using the round table and using the rectangular table. Then the point of my assertion will be fully satisfied if I get you to opt for the former, and get you to understand why this is the right decision—a rectangular table will make for awkward seating. If I also happen to succeed in limning the structure of reality with my assertion, by choosing a sentence with the right logical structure, that is no part of what I set out to do.

If *assertions* of sentences involving mathematical vocabulary are not intended to limn the structure of the world, what could be the motivation for thinking that the truth-conditions of the sentences themselves play this role? As far as I can tell, it is nothing over and above the idea that the logical structure of atomic sentences should correspond to the structure of the world. Remove this idea and there is no motivation left.

It is true that our language is compositional. But the role of compositionality is to allow for the production of large numbers of sentences from a restricted lexicon. To claim that compositionality plays the *additional* role of allowing for the representation of the structure of reality is to set forth a doctrine that is not supported by our linguistic usage. It is to start out with a preconception of the way language ought to work, and impose it on our linguistic theorizing from the outside—from beyond what is motivated by the project of making sense of our linguistic practice. (For further discussion of Tractarian conceptions of language, see Heil (2003).)

We can now see where conventional wisdom goes wrong. The idea was supposed to be that Platonism+committalism is able to take the logical structure of a sentence like '2 is prime' at face value because it has a matching ontology to offer: the reference of '2' is taken to exist and is taken to have the property expressed by 'is prime'. Nominalism+noncommittalism, on the other hand, lacks the needed ontology, so it must choose between treating ‘2 is prime’ as saying something literally false and meddling with its logical structure. But once one abandons the doctrine that the logical structure of an atomic sentence must correspond to the structure of reality, there is room for a distinction between the *semantic values* of expressions occurring in a sentence—a piece of theoretical machinery used to explain how the meanings of complex expressions depend on the meanings of its parts—and the objects that must exist in order for the sentence’s truth-conditions to be satisfied. In particular, a friend of nominalism+noncommittalism can take the logical structure of ‘2 is prime’ at face value, and assign semantic values to ‘2’ and ‘is prime’ in the course of developing a compositional semantic theory, while resisting the conclusion that satisfaction of the sentence’s *truth-conditions* requires that the semantic value of ‘2’ exist and instantiate the semantic value of ‘is prime’. This is a tricky point, so I shall dwell on it further.

Consider a sentence like ‘roses are red *and* violets are blue’. Standard semantic theories assign a semantic value to ‘and’—a certain kind of *function*—but it would be a mistake to go from this to the conclusion that the sentence carries commitment to functions: that part of what is required of the world in order for the sentence’s truth-conditions to be satisfied is that there be functions. To take the additional step would be to misjudge the role of semantic values in our semantic theorizing. The point of assigning a semantic value to ‘and’ is that we want our semantics to be *compositional*—we want a systematic way of determining the semantic properties of sentences of the form $\ulcorner\phi$ and $\psi\urcorner$ on the basis of the semantic properties of ϕ and ψ—not to get the result that sentences involving ‘and’ count the semantic value amongst their ontological commitments.

Most philosophers take a similar attitude towards the semantic values of predicates. Consider a sentence like ‘Susan runs’. Standard semantic theories assign a semantic value to ‘runs’—in the simplest case, an extension—but most of us would want to resist the conclusion that ‘Susan runs’ carries commitment to extensions: that part of what is required of the world in order for the sentence's truth-conditions to be satisfied is that there be extensions. Again, the point of assigning a semantic value to 'runs' is that we want our semantics to be *compositional*—we want a systematic way of determining the semantic properties of expressions of the form ⌜t runs⌝ on the basis of the semantic properties of t—not to get the result that sentences involving 'runs' count the semantic value amongst their ontological commitments.

When it comes to singular terms, however, it is common for philosophers to think of semantic values as playing a broader role. Philosophers often presuppose that the semantic value of a singular term t does more than just deliver a systematic way of determining the semantic properties of sentences of the form ⌜t Fs⌝ on the basis of the semantic properties of F. There is the additional requirement that the semantic value of t be counted amongst the ontological commitments of ⌜t Fs⌝. How is this expanded role for the semantic values of singular terms supposed to be motivated? As far as I can tell, it is nothing over and above the idea that the logical structure of atomic sentences should correspond to the structure of the world. Remove this idea and there is no motivation left.

This is not to say, of course, that a sentence of the form ⌜t Fs⌝ should *always* remain uncommitted to the semantic value of t. In many cases—when t is 'Susan' and F is 'runs', for example—commitment to the semantic value of t is the right result. The point I wish to make is that commitment to the semantic value of t shouldn't be regarded as *automatic*. When appropriate, it can be secured by assigning the right semantic values to t and F, and specifying the right rule for extracting truth-conditions from the semantic values of sentences. But one shouldn't assume that a sentence of the form ⌜t Fs⌝ must carry commitment to a certain object merely on the grounds that that object has been assigned as t's semantic value. As in the case of lexical items falling under different syntactic categories, the reason for assigning semantic values to singular terms is to allow for compositionality, not to secure a correspondence between the logical structure of our sentences and the structure of the world.

In a semantic theory where the role of semantic values is exhausted by considerations of compositionality—together with the principle that the semantic value of a sentence must determine truth-conditions for the sentence, relative to suitable contextual parameters—it is possible to develop a semantics for mathematical discourse that runs contrary to conventional wisdom. One can take the logical structure of mathematical sentence at face value and still get the conclusion that all that is required of the world in order for the truth-conditions of 'the number of the planets is eight' to be satisfied is that there be eight planets, and the conclusion that any true sentence in the language of pure mathematics gets assigned trivial truth-conditions and any false sentence in the language of pure mathematics gets assigned trivial falsity-conditions. (See my 'On Specifying Truth-Conditions' for details. For a related discussion, see Hofweber's contribution to this volume.)

I hear some complaints:

1. *Objection:* You claim that one can give a semantic theory according to which mathematical sentences carry no objectionable commitments. But the semantic theory itself will carry all sorts of commitments (functions, extensions, etc.). So we're stuck with the problematic commitments anyway! How is this any help to a friend of nominalism?

   Response: Suppose we've established the view that sentences of the object-language carry no commitment to abstract objects. Then we've established that there can be a type of quantifier that has the same syntax and inferential patterns as the committalist's quantifier but carries no commitment to abstract objects. A friend of nominalism+noncommittalism will claim that she uses such quantifiers when she uses the object-language—and, therefore, that she incurs no problematic commitments in ordinary mathematical discourse. But she will wish to make an additional claim: she will claim that she also uses the special quantifier when she does *semantics*. So she can consistently claim that problematic commitments are incurred neither when she uses the object-language to do mathematics nor when she uses the metalanguage to do semantics. (Whether she can use this move to convince the unconvinced is a different matter—again, see *OSTC* for details.)

2. *Objection:* Suppose I buy your semantics, and you convince me that all that is required of the world in order for the truth-conditions of 'the number of the planets is eight' to be satisfied is that there be eight planets. Why should I conclude from this that 'the number of the planets is 8' carries no problematic commitments? After all, everyone should believe that 'the number of the planets is eight' is true just in case the number of the planets is eight. Isn't this enough for the conclusion that 'the number of the planets is eight' carries commitment to numbers?"

   Response: In asking about the ontological commitments of a sentence there is a potential ambiguity. On one reading of the question, one wants to know which of the objects that are carved out by the world's metaphysical structure need to exist in order for the sentence's truth-conditions to be satisfied. If this is how you think about the matter, you should say that the truth-conditions of 'the number of the planets is eight' can be stated in 'more fundamental' and 'less fundamental' terms. The more fundamental statement—the one that tells us which of the objects carved out by the world's metaphysical structure must exist in order for the sentence to be true—is the one delivered by the proposed semantics: that there be eight planets. A less fundamental statement—one that makes no effort to limn metaphysical structure—is that the number of the planets be eight. (I am not myself able to understand what people mean by 'metaphysical structure' or 'metaphysically fundamental', but Cameron (forthcoming) and Williams (typescript) have developed the proposal in this direction.) On a different reading of the question—the reading that I prefer—one is unconcerned with metaphysical structure. In asking about a sentence's ontological commitments all one wants is an informative statement of how the world must be in order for the sentence's truth-conditions to be satisfied. On this second reading, the proposed semantics can be used to argue that even though one can accurately specify the truth-conditions of 'the number of the planets is eight' by saying 'that the number of the planets be eight', it would be just as accurate to say 'that there be eight planets'. Neither of these statements counts as 'more fundamental' than the other: for the number of the planets to be eight *just is* for there to be eight planets. And it is in this sense that 'there are eight planets' can be said to carry no problematic commitments.

Moral: the claim that only committalists are in a position to take the logical structure of mathematical sentences at face value is based on a questionable conception of language. When the problematic assumptions are dropped, there is no obstacle to taking logical structure at face value while being a non-committalist—or, indeed, a trivialist.

⋆ ⋆ ⋆

Where does all of this leave us? First and foremost, I would like to urge you to consider becoming a trivialist. In doing so, you would put yourself in a position to give a satisfying answer to a question that has been a source of endless woe for epistemologists of mathematics: how do we know that the world satisfies the requirements that would need to be satisfied in order for the truths of pure mathematics to be true and its falsehoods to be false? While the non-trivialist is searching for a way to justify non-trivial claims to the effect that the realm of abstract objects has a certain property, or that the world is infinite, you will confidently proclaim: "The relevant requirements can be known to be satisfied because *nothing* is required of the world in order for a truth of pure mathematics to be true, and *nothing* is required of the world in order for a falsity of pure mathematics to be false."

If you were beholden to conventional wisdom, you would worry that the price to be paid for trivialism is semantic awkwardness—you would fear that by becoming a trivialist you would have had to choose between meddling with the logical form of mathematical statements and claiming that they are used to convey information which is very different from what they literally say. But now you know that such fears would be doubly mistaken. You know, first of all, that as long as one is a *subtle* Platonist one can be a trivialist even if one is also a committalist (and there was never any worry about *committalists* being led into semantic awkwardness). You know, moreover, that the claim that *non*committalists are faced with semantic awkwardness is based on a questionable conception of language. So you know that even if one was unhappy about subtle Platonism, one could be a trivialist by becoming a noncommittalist and abandoning the idea that truth can only be achieved through a correspondence between the structure of language and the structure of reality.

The previous paragraph identifies two ways of being a trivialist without plunging into semantic awkwardness: subtle Platonism and post-Tractarian noncommittalism. It is worth emphasizing that these views more closely related than one might think. Suppose you are a subtle Platonist. You believe that for the number of the Fs to be n *just is* for there to be n Fs. So when the committalist claims that satisfaction of the truth-conditions of 'the number of the planets is 8' requires of the world that the planets be numbered by the number 8, and the noncommittalist counters that all that is required is that there be eight planets, you will see them as stating the very same requirement—for the committalist's requirement to be met *just is* for the noncommittalist requirement to be met.

2 Knowing trivial truths

Suppose trivialism is right and the truths of pure mathematics have trivial truth-conditions. What could the point of mathematical knowledge be? The purpose of this section is to answer that question. But it will take some time to set things up—the crucial discussion won't take place until section 2.3.

2.1 Intelligibility

Let a *story* be a set of sentences in some language we understand. I shall assume that stories are read *de re*: that every name used by the story is used to say of the name's actual bearer how it is according to the story, and that every predicate used by the story is used to attribute the property actually expressed by the predicate to characters in the story. Accordingly, in order for a story that says 'Hesperus is covered with water' to be true it must be the case that Venus itself is covered with H_2O. (I shall ignore names that are actually empty, such as 'Sherlock Holmes', and predicates that are actually empty, such as '. . . is composed of phlogiston' or '. . . is a unicorn'.)

Sometimes we describe a story as unintelligible on the grounds that it is too complicated for us to understand. That is not the notion of unintelligibility I have in mind here. As I understand the term, a story is *unintelligible* for a subject if her best effort to make sense of a scenario in which the story is true would yield something she regards as *incoherent*. (Intelligibility can then be defined as non-unintelligibility.) Let me give you some examples of what I have in mind.

Consider a story that says 'A fortnight elapsed in only 13 days'. My best effort to make sense of a scenario in which this story is true ends in incoherence. 'Fortnight' *means* 'period of 14 days'. So a scenario verifying 'A fortnight elapsed in only 13 days' would have to be a scenario in which a period of 14 days lasts only 13 days, which is something

I regard as incoherent. (Of course, it would be easy enough to make sense of a scenario in which *language* is used in such a way that the expression 'a fortnight elapsed in only 13 days' is true. But that won't help with the question of whether the original story is intelligible, in the relevant sense.)

The preceding example might tempt you to think that only 'conceptually inconsistent' stories count as unintelligible. But consider a story that says 'Hesperus is not Phosphorus'—presumably an example of a 'conceptually consistent' statement. My best effort to make sense of a scenario in which this story is true ends in incoherence. For a scenario in which the story is true would have to be a scenario in which Hesperus itself (i.e. Venus) fails to be identical with Phosphorus itself (i.e. Venus), and the nonselfidentity of Venus is something I regard as incoherent. Another example: consider a story that says 'there is a lake with water but no H_2O'. For something to contain water *just is* for it to contain H_2O. So a scenario verifying 'there is a lake with water but no H_2O' would have to be a scenario in which something that is filled with H_2O fails to be filled with H_2O—which I regard as incoherent. (As before, it would be easy enough to make sense of a scenario in which *language* is used in such a way that the expression 'Hesperus is not Phosphorus' or 'there is a lake with water but no H_2O' is true. But that is irrelevant to the issue at hand.)

> *Objection:* Given that you insist on a *de re* reading of stories, I can see why you want to treat 'Hesperus is not Phosphorus' and 'there is a lake with water but no H_2O' as unintelligible. But consider a scenario in which the first celestial body to be visible in the evenings is not the last celestial body to disappear in the morning. I grant you that 'Hesperus is not Phosphorus' is not literally true in this scenario. But surely there is some derived sense of 'verify' such that 'Hesperus is not Phosphorus' is verified by this scenario. So there is a certain sense in which 'Hesperus is not Phosphorus' is intelligible after all.
>
> *Reply:* In order to claim that this alternate notion of intelligibility is well-

defined when we go beyond toy examples like ‘Hesperus is Phosphorus’, one needs a substantial assumption: the assumption that every name and every predicate has a ‘narrow content’ or ‘primary intension’ which can be used to determine which scenarios will count as verifying a given sentence in the derived sense of ‘verify’. I myself am skeptical of this assumption, since I find no evidence for it in linguistic practice. (As far as I can tell, all that is required by our actual linguistic usage is the ability to determine which of a highly restricted set of contextually salient possibilities would count the sentence asserted as expressing a true proposition—see my ‘Vague Representation’ for details.) But nothing in this paper hinges on rejecting the assumption. If you think the alternate characterization of intelligibility is legitimate, that’s fine. Just keep in mind that it’s not the one that will be relevant in this paper.

Moral: As it is understood here, the intelligibility of a story (for a subject) is a highly non-*a priori* matter. For what one finds unintelligible depends on whether one believes that Hesperus is Phosphorus, or that water is H_2O. And knowledge of such truths is far from *a priori*.

2.2 Intelligibility and Identity

I would like to suggest that there is a close connection between intelligibility and *identity*.

Statements of the form ‘$a = b$’ are identity statements. But they are only a special case. Consider the following sentences:

Sibling

To be a sibling *just is* to share a parent.

[In symbols: ‘$\text{Sibling}(x) \equiv_x \exists y \exists z(\text{Parent}(z, x) \wedge \text{Parent}(z, y) \wedge x \neq y)$’]

HEAT

To be hot *just is* to have high mean kinetic energy.

[In symbols: 'Hot(x) $\equiv_x$ High-Mean-Kinetic-Energy(x)']

WATER

To be composed of water *just is* to be composed of H_2O.

[In symbols: 'Composed-of-water(x) $\equiv_x$ Composed-of-H_2O(x)']

In these three sentences the expression 'just is' (or its formalization '$\equiv_x$') is functioning as an identity-predicate of sorts. To accept '$F(x) \equiv_x G(x)$' is not simply to accept that all and only the Fs are Gs. If you accept SIBLING, for example, you believe that there is *no difference* between being a sibling and sharing a parent with someone; you believe that if someone is a sibling it is *thereby* the case that she shares a parent with someone. (Compare: if you accept 'Hesperus is Phosphorus', you believe that someone who travels to Hesperus has *thereby* traveled to Phosphorus.)

One might be tempted to describe SIBLING, HEAT and WATER as expressing identities amongst *properties* (e.g. 'the property of being a sibling = the property of sharing a parent'.) I have no qualms with this description, as long as property-talk is understood in a suitably deflationary way. But I will avoid property-talk here because it is potentially misleading. It might be taken to suggest that one should only assert SIBLING if one is prepared to countenance a traditional Platonism about properties—the view that even though it is intelligible that there be no properties, we are lucky enough to have them. The truth of SIBLING, as I understand it, is totally independent of such a view. If one wishes to characterize the difference between '$\equiv_x$' and the standard first-order identity predicate '$=$', the safe thing to say is that whereas '$=$' takes a singular-term in each of its argument-places, '$\equiv_x$' takes a first-order predicate in each of its argument places. I shall therefore refer to sentences of the form $\ulcorner a = b \urcorner$ (where a and b are singular terms) as *first-order* identity statements, and sentences of the form $\ulcorner \phi(x) \equiv_x \psi(x) \urcorner$ (where $\phi(x)$ and $\psi(x)$ are first-order predicates) as *second-order* identity statements.

Sometimes one is in a position to endorse something in the vicinity of a second-order identity statement even though one has only partial information. Suppose you know that the chemical composition of water includes oxygen but don't know what else is involved. You can still say:

> Part of what it is to be composed of water is to contain oxygen.
>
> [In symbols: 'Composed-of-water(x) $\ll_x$ Contains-Oxygen(x).']

I shall call this as a *semi-identity statement*. Think of it as a more idiomatic a way of saying:

> To be composed of water *just is* (to contain oxygen and to be composed of water).

(Please note that it is no part of the view that '$F(x) \ll_x G(x)$' entails that something is an F 'in virtue' of being a G, or that being a G is 'more fundamental' than being an F, or that being G is part of the 'essence' of an F.)

As in the case of second-order identity statements, it is tempting to think of semi-identity statements in terms of properties (e.g. 'the property of being water has the property of containing oxygen as a part'.) Again, I have no objection to this sort of description, as long as property-talk is taken in a suitably deflationary spirit. But I will avoid it here because of its potential to mislead.

Second-order identity statements can be dispensed with in the presence of semi-identity statements. Water, for example, is equivalent to the conjunction of 'part of what it is to be water is H_2O' and 'part of what it is to be H_2O is to be water'. And, in general, '$F(x) \equiv_x G(x)$' is equivalent to the conjunction of '$F(x) \ll_x G(x)$' and '$G(x) \ll_x F(x)$'.

Note, moreover, that the content of a first-order identity-statement '$a = b$' can be expressed by way of the second-order identity-statement '$x = a \equiv_x x = b$'—to be *a just is* to be b. (More precisely: '$a = b$' is equivalent to the conjunction of '$x = a \equiv_x x = b$' and '$\exists x(a = x)$'.) This means that semi-identity statements can be used to do the work of both first- and second-order identity statements.

A few paragraphs back I hinted at a close connection between intelligibility and identity. I can now tell you what I take the connection to be. To wit: the sole source of unintelligibility for a subject is inconsistency with the semi-identities she accepts. (More precisely: a story will be unintelligible for a subject just in case she is in a position to derive, using inferences she takes to be logically valid, something she regards as incoherent from the result of adding the story to the set of semi-identities she accepts.)

Three observations: (1) In suggesting a connection between identity and intelligibility I do not mean to suggest that one of these notions is 'prior' to the other. The claim is that the two notions are connected, and that one can better understand them by understanding the connection. (2) So far we have focused our attention on intelligibility for a subject. What about intelligibility *simpliciter*? If you think there is an objective fact of the matter about which semi-identities are true you can go on to say that a story is intelligible *simpliciter* just in case it is logically consistent with the set of true semi-identities. (3) So far we have focused on the intelligibility of a story. What about the intelligibility of a *scenario*? The usual way of picking out a scenario is by setting forth a story (or some other kind of representation); in this special case, one can say that a scenario is intelligible just in case the story used to pick it out is intelligible. (For a more detailed discussion of intelligibility, see my 'An Account of Possibility'.)

The connection between identity and intelligibility can help us get a grip on the elusive notion of a sentence's truth-conditions. To see this, note that one can think of a sentence's truth-conditions as a requirement imposed on the world—the requirement that the world be a certain way. Knowing whether a scenario we take to be intelligible must fail to obtain in order for the requirement to be met is valuable because it gives us an understanding of how the world would need to be in order for the requirement to be satisfied. But knowing whether a scenario we take to be *un*intelligible must fail to obtain in order for the requirement to be to be met is not very helpful. For when one is unable to make sense of a scenario, the claim that it must fail to obtain gives one no understanding of how the world would need to be in order for the requirement to be satisfied. The lesson is that one can model a sentence's truth-conditions as a partition of the space of scenarios one takes to be intelligible. Accordingly, the connection between semi-identity and intelligibility yields a connection between semi-identity and a sentence's truth-conditions: since the scenarios one regards as intelligible will depend on the semi-identities one accepts, and since a sentence's truth-conditions can be modeled as a set of intelligible scenarios, one's beliefs about the range of possible truth-conditions will depend on the semi-identities one accepts.

A first consequence of this observation is that a subject is committed to seeing any sentence she takes to follow from semi-identity statements she accepts as having trivial truth-conditions. If, for example, you think that to be composed of water *just is* to be composed of H_2O, then you are committed to thinking that nothing is required of the world in order for the truth-conditions of 'if a lake contains water, it contains H_2O' to be satisfied. And if you think that to be Hesperus *just is* to be Phosphorus, you are committed to thinking that nothing is required of the world in order for the truth conditions of 'if Hesperus is a planet, Phosphorus is a planet' to be satisfied.

It is worth considering some additional consequences of the connection between semi-identity and intelligibility. The semi-identity operator '$\ll$' can bind more than one variable. For instance:

$$\text{Sisters}(x, y) \ll_{x,y} \exists z(\text{Parent}(z, x) \wedge \text{Parent}(z, y))$$

[*Read:* part of what it is for x and y to be sisters is for x and y to share a parent.]

But it can also bind no variables at all:

a wedding takes place $\ll$ someone gets married

[*Read:* part of what it is for a wedding to take place is for someone to get married.

The same is true of second-order identity statements. For instance:

a wedding takes place $\equiv$ someone gets married

[*Read:* for a wedding to take place *just is* for someone to get married.]

In light of the connection between semi-identity statements and truth-conditions, this yields the result that a subject who accepts 'for a wedding to take place *just is* for someone to get married' is committed to the view that 'a wedding takes place' and 'someone gets married' have the same truth-conditions.

Now let $\top$ be a sentence which is known to have trivial truth-conditions. One can use $\ulcorner p \equiv \top \urcorner$ to capture the thought that p has trivial truth-conditions. For instance:

$\forall x(x = x) \equiv \top$

[*Read:* that everything is identical is trivially the case.]

Second-level identity is an equivalence relation. So if one accepts $\ulcorner p \equiv \top \urcorner$ and accepts $\ulcorner q \equiv \top \urcorner$ one is committed to accepting $\ulcorner p \equiv q \urcorner$ and $\ulcorner q \equiv p \urcorner$. If one takes seriously the idea that a logical truth has trivial truth-conditions, then part of what one does when one recognizes p as a logical truth is to accept $\ulcorner p \equiv \top \urcorner$. So one is committed to accepting $\ulcorner p \equiv q \urcorner$ for any p and q which one recognizes as logical truths. (When I say 'logical truth' I mean 'truth of a free logic', since we want to avoid the result that $\ulcorner \exists x(x = c) \urcorner$ has trivial truth-conditions whenever c is a proper name.)

The same goes for mathematics, according to the trivialist. Suppose you think that nothing is required of the world in order for the truth-conditions of 'there are numbers' to be satisfied—equivalently:

there are numbers $\equiv \top$

[*Read:* that there be numbers is trivially the case.]

You have thereby committed yourself to:

there are numbers $\equiv \forall x(x = x)$

[*Read:* for there to be numbers *just is* for everything to be self-identical.]

For if it is trivially the case that there are numbers and it is trivially the case that everything is self identical, then there is no more to the world's satisfying the one than there is to the world's satisfying the other: if the world is such that there are numbers it is *thereby* such that everything is self-identical, and if it is such that everything is self-identical it is *thereby* such that there are numbers. So 'there are numbers' and '$\forall x(x = x)$' have the same truth-conditions. More generally, every true sentence of pure mathematics has the same truth-conditions as any other.

Moral: suppose you buy into the idea that a story is intelligible for a subject just in case the subject takes the story to be consistent with the set of semi-identities she accepts. Then you should also buy into the idea that a subject's views about the range of possible truth-conditions—and therefore her views about which sentences share their truth-conditions—will be shaped by the identity statements she accepts.

2.3 Knowledge

The discussion in the preceding sections suggests that our cognitive attitudes towards semi-identity statements and ordinary statements play different roles: by accepting semi-identity statements we fix the limits of what we take to be intelligible; by accepting ordinary statements we partition the space of intelligible scenarios into regions that are treated as candidates for truth and regions that are ruled out as false. In this section I will sketch an epistemological picture based on this idea.

In rough outline, the picture is this. In an effort to satisfy our goals, we develop strategies for interacting with the world. Fruitful strategies allow us to control what the world is like and predict how it will evolve under specified circumstances. They also allow us to direct our research in ways that lead to the development of further fruitful strategies. In order to articulate the strategies we adopt, we do three things at once: firstly, we develop a *language* within which to formulate theoretical questions; secondly we set forth *theoretical claims* addressing some of these questions; finally, we endorse a family of *semi-identity statements*. The third task is connected to the other two because the semi-identity statements we endorse help determine which theoretical questions are worth investigating and which are not. To a certain extent, the endeavor is a holistic one. It is sometimes possible to vary the semi-identity statements one endorses without significantly affecting the success of one's methods of inquiry, provided one makes compensating adjustments in the theoretical claims one accepts. (For a more detailed discussion, see my 'An Account of Possibility'.)

Consider, for example, our acceptance of 'part of what it is to be water is to contain hydrogen and oxygen'. It affects our overall theorizing by ruling out certain questions as pointless while allowing others as fruitful. (See Block and Stalnaker (1999).) For example, we would regard it as wrong-headed to try to understand *why* every portion of water contains oxygen and hydrogen. ("Water *just is* H_2O!") But we see it as a worthwhile endeavor to try to understand *why* liquid water is colorless. Now suppose we vary the water-related semi-identities we accept. Say we reject 'part of what it is to be water is to contain hydrogen and oxygen' and accept instead 'part of what it is to be water is to be a colorless liquid'. Then we will see a different range of questions as worth pursuing. It will now seem pointless to ask why liquid water is colorless ("that's just part of what it *is* to be water!") but interesting to ask why every portion of water contains oxygen and hydrogen. The change would also affect the way certain theoretical claims are formulated.

For instance, we would need to reformulate the principle that water at sea-level freezes at 0 degrees celsius. Nevertheless, someone sufficiently committed to the alternate semi-identity might be able make enough adjustments elsewhere in the system to secure a successful articulation of her methods of inquiry.

Is there an objective fact of the matter about which choice of semi-identity statements is correct? My own view is that there is not. All one can say is that different choices are more or less amenable to the development of a successful articulation of one's methods of inquiry, given the way the world is. The empirical facts make it the case that acceptance of 'part of what it is to be water is to contain hydrogen and oxygen' leads to particularly fruitful theorizing, and one has strong reasons to accept it on this basis. But there is no more to be said on its behalf. And when it comes to other semi-identity statements, the empirical pressures are even milder. For example, whether it is a good idea to accept 'part of what it is to be a reptile is to have a certain lineage' rather than 'part of what it is to be a reptile is to have a certain phenotype' might to a large extent depend on the purposes at hand. I would like to emphasize, however, that nothing in this paper will turn on taking a pragmatic attitude towards semi-identity statements. If you think there is an objective fact of the matter about which semi-identities are correct, that's fine for present purposes.

Let us now turn our attention to logical truth. Suppose you learn that it fails to be the case that $\neg p$. Then, if you are a friend of classical logic, you will be in a position to conclude that p. But not just that: you will think that your understanding of why it fails to be the case that p is *already* an understanding of why it is the case that p. There is no need to add an explanation of why the *transition* is valid, since there is nothing to be explained: for it to fail to be the case that $\neg p$ *just is* for it to be the case that p. There is no intelligible scenario in which the transition fails, whose absence would need to be accounted for. Asking: "I can see that if it is not the case that $\neg p$ it will be the case that p, but *why* is this so?" is as wrong-headed as asking: "I can see that water is H_2O, but *why* is this so?".

The point generalizes to more complex logical truths. When one treats a sentence ϕ as *logically* true, one does more than simply treat it as true. One is, in effect, accepting the higher-order identity statement $\ulcorner \phi \equiv \top \urcorner$ (which is equivalent to accepting $\ulcorner \psi \equiv \theta \urcorner$ when ϕ is of the form $\ulcorner \psi \leftrightarrow \theta \urcorner$, and equivalent to accepting $\ulcorner \psi \ll \theta \urcorner$ when ϕ is of the form $\ulcorner \psi \rightarrow \theta \urcorner$). This means that the result of treating $\ulcorner \psi \rightarrow \theta \urcorner$ as a logical truth is not just that one will take oneself to be justified in accepting θ whenever one feels justified in acepting ψ. One will think that one's understanding of why ψ's truth-conditions are satisfied is *already* an understanding of why θ's truth-conditions are satisfied. As before, there is no need to add an explanation of why the *transition* from ψ to θ is valid: that θ's truth-conditions be satisfied is *part of what it is* for ψ's truth-conditions to be satisfied. Of course, when $\ulcorner \psi \rightarrow \theta \urcorner$ is sufficiently complex, coming to recognize it as a logical truth may be a highly non-trivial process. And throughout that process one might be justified in asking oneself *why* $\ulcorner \psi \rightarrow \theta \urcorner$ is a logical truth. But one's query can be addressed by setting forth a sufficiently illuminating proof. And once one has understood such a proof, one will see that there is no intelligible scenario in which the transition from ϕ to θ fails, and therefore that one's understanding of why ψ's truth-conditions are satisfied is *already* an understanding of why θ's truth-conditions are satisfied.

We are now in a position to answer an important question about logical knowledge: if the truths of pure logic have trivial truth-conditions, what could the point of knowing a logical truth be? The answer is that in learning a logical truth one increases one's ability to distinguish between intelligible and unintelligible scenarios, and therefore one's ability to use old information in new ways. (See Stalnaker (1984) ch. 5, and Stalnaker (1999) chs. 13 and 14.) An example will help illustrate the point.

Suppose that there are seventeen apples, and that you have counted them. This gives you a certain range of abilities. You are able to determine whether you got short-changed at the market, or whether there are enough apples for your recipe. You are also able to answer questions of the form 'How many apples?' One might represent that you have such a range of abilities by claiming that you know that there are seventeen apples [in symbols: $\exists_{17}!x(\text{Apple}(x))$]. Now suppose that there are twenty-nine pears, and you have also counted them. This, again, gives you a distinctive range of abilities, a fact that might be represented by claiming that you know that there are twenty-nine pears [in symbols: $\exists_{29}!x(\text{Pear}(x))$]. Perhaps you are able to combine these two cognitive accomplishments in the service of a single task. You might, for instance, be in a position to determine whether there are more apples than pears. But other tasks might elude you. Say you know that every relevant piece of fruit is an apple or a pear, and that no piece of fruit is both an apple and a pear. Then you have all the information you need to answer questions of the form 'How many pieces of fruit?'. But you may still not be in a position to use the information at you disposal for that particular task, at least not immediately. What is missing is knowledge of a logical truth:

$$\exists_{17}!x(\text{Apple}(x)) \wedge \exists_{29}!x(\text{Pear}(x)) \wedge \neg\exists x(\text{Apple}(x) \wedge \text{Pear}(x)) \rightarrow \exists_{46}!x(\text{Apple}(x) \vee \text{Pear}(x))$$

In performing the relevant computation, do you acquire novel information about the world? It is tempting to say that you do, since you will learn there are forty six pieces of fruit. But I think the right thing to say is that you don't. For you already knew that every piece of fruit is an apple or a pear (but not both) and that there are seventeen apples and twenty nine pears, and *part of what it is* for that to be the case is that there be forty six pieces of fruit. In carrying out the computation, your cognitive accomplishment consists not in the acquisition of new information, but in the ability to deploy old information in new ways. Before you carry out the computation you are unsure about whether a scenario in which there are, say, *thirty*-six pieces of fruit could be genuinely intelligible while respecting the information you already had about apples and pears. What the computation reveals is that it is not. You have increased your ability to distinguish between intelligible and unintelligible scenarios, and this gives you the ability to see how to answer questions of the form ‘How many pieces of fruit?’ in light of the information you had at your disposal all along. (I have greatly benefitted from discussion with Adam Elga on these topics.)

When one embraces a logical system one adopts a framework for settling questions of intelligibility. In deciding which logic to accept one must therefore strike a delicate balance. If one’s logic is too strong, it will commit one to treating as unintelligible scenarios that might have been useful in making sense of the world. By weakening one’s logic one opens the door to a larger range of intelligible scenarios, all of them candidates for truth. In discriminating amongst them one will have to explain why one favors the ones one favors. The relevant explanations may sometimes lead to fruitful theorizing about the world. But they may also prove burdensome. Consider a friend of intuitionistic logic, who denies that for it to fail to be the case that $\neg p$ *just is* for it to be the case that p. She thinks it might be worthwhile to ask why it is the case that p even if you fully understand why it is not the case that $\neg p$. In the best case scenario, making room for an answer will lead to fruitful theorizing. But things may not go that well. One might come to see the newfound conceptual space between a sentence and its double negation as a pointless distraction, demanding explanations in places where there is nothing fruitful to be said. (For a particularly insightful discussion of intuitionistic logic, see Wright (2001).)

There can sometimes be empirical pressure of a more or less direct kind in favor of a particular semi-identity statement—think of ‘part of what it is to be water is to contain hydrogen and oxygen’. But when it comes to choosing a logical system, the decision to accept the relevant semi-identities is likely to be driven by considerations of a more general nature. We want a framework for settling questions of intelligibility that is flexible enough to allow for interesting questions to be posed but constrained enough to make our insights transferable to a large range of contexts.

Could one be an *objectivist* about logic, and claim that there is an objective fact of the matter about which logical system is correct? Perhaps the thought is that the world has one true 'logical structure', and that a logical system is correct to the extent that it does justice to the logical structure of the world. (How do we know which logical system is objectively correct? Maybe we get evidence of correctness when a logical system delivers a useful framework for settling questions of intelligibility.) Nothing I have said is in tension with making such additional claims. But, as far as I can tell, there they would be unmotivated without a prior commitment to the objectivist standpoint.

Let us finally turn to the case of mathematics. If you are a non-trivialist, you think there is a world of difference between logic and mathematics. Whereas the truths of pure logic have trivial truth-conditions (and its falsities have trivial-falsity conditions), there are intelligible scenarios that fail to satisfy the truth-conditions of some mathematical truth (or the falsity conditions of some mathematical falsehood). Accordingly, one needs some sort of entitlement to the view that the rogue scenarios fail to obtain before one can claim to know that the relevant mathematical truths are true (or that the relevant mathematical falsehoods are false). For the trivialist, on the other hand, there is no deep difference between logic and mathematics. As in the case of logic, mathematical truths have trivial truth-conditions (and mathematical falsities have trivial falsity-conditions). The difference is simply that the language of mathematics enjoys expressive resources that the language of logic lacks.

These enhanced expressive resources are important in two ways. First, they allow us to articulate requirements on the world that cannot be articulated in the language of pure logic, or that can only be articulated with significant awkwardness. By using the sentence '$\#_x\text{Apple}(x) = \#_x\text{Pear}(x)$', for example, one can express the thought that there be just as many apples than pears—something that cannot be done within the language of first-order logic in any straightforward sense. (Can the trivialist generalize this point, and give a recipe that specifies ontologically innocent truth-conditions for arbitrary mathematical sentences? This is not as easy as one might think, but see my *OSTC*.)

Second, the enhanced expressive resources of mathematics improve our ability to sort out the intelligible from the unintelligible. Consider the rather unlovely logical truth that I mentioned a few paragraphs back:

$$\exists_{17}!x(\text{Apple}(x)) \wedge \exists_{29}!x(\text{Pear}(x)) \wedge \neg\exists x(\text{Apple}(x) \wedge \text{Pear}(x)) \rightarrow \exists_{46}!x(\text{Apple}(x) \vee \text{Pear}(x))$$

By accepting this sentence one acquires the ability to rule out as unintelligible a scenario in which there are seventeen apples, twenty-nine pears and anything other than forty-six apple-or-pears. But in accepting the (far simpler) mathematical sentence '$17 + 29 = 46$', one acquires a more general ability—the ability to rule out as unintelligible a scenario in which there are seventeen Fs, twenty-nine Gs and anything other than forty-six F-or-Gs (provided no Fs are Gs). And, of course, an improved ability to sort out the intelligible from the unintelligible is important because it gives us an improved ability to transfer insights from one context to another. To pick a simple example, knowledge of the basic facts of multiplication puts you in a position to use the insight gained from counting the rows and the insight gained from counting the columns for the purposes of answering questions of the form 'How many tiles?'. And, of course, this is only the beginning.

Moral: even if the trivialist believes that the truths of pure mathematics have trivial-truth conditions, she is able to explain why mathematical knowledge is worthwhile.

References

Block, N., and R. Stalnaker (1999) "Conceptual Analysis, Dualism, and the Explanatory Gap," *Philosophical Review* 108, 1–46.

Boolos, G., ed. (1990) *Meaning and Method: Essays in Honor of Hilary Putnam*, Camridge University Press, Cambridge.

Bottani, A., M. Carrara, and P. Giaretta, eds. (2002) *Individuals, essence, and identity: themes of analytic metaphysics*, Kluwer Academic, Dordrecht and Boston.

Cameron, R. (forthcoming) "Quantification, Naturalness and Ontology." In Hazlett (forthcoming).

Field, H. (1980) *Science Without Numbers*, Basil Blackwell and Princeton University Press, Oxford and Princeton.

Fine, K. (2002) *The Limits of Abstraction*, Oxford University Press, Oxford.

Frege, G. (1884) *Die Grundlagen der Arithmetik*. English Translation by J.L. Austin, *The Foundations of Arithmetic*, Northwestern University Press, Evanston, IL, 1980.

Goldfarb, W. (1997) "Metaphysics and Nonsense: On Cora Diamond's *The Realistic Spirit*," *Journal of Philosophical Research* 22, 57–73.

Hacker, P. (1986) *Insight and Illusion: Themes in the Philosophy of Wittgenstein*, Oxford University Press, Oxford.

Hazlett, A., ed. (forthcoming) *New Waves in Metaphysics*, Palgrave-Macmillan.

Heil, J. (2003) *From an Ontological Point of View*, Clarendon Press, Oxford.

Hodes, H. T. (1984) "Logicism and the Ontological Commitments of Arithmetic," *Journal of Philosophy* 81:3, 123–149.

Hodes, H. T. (1990) "Ontological Commitments: Thick and Thin." In Boolos (1990), 347-407.

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

McGee, V. (1993) "A Semantic Conception of Truth?" *Philosophical Topics* 21, 83–111.

Morton, A., and S. Stich, eds. (1996) *Benacerraf and his Critics*, Basil Blackwell, Oxford.

Parsons, C. (1983) *Mathematics in Philosophy*, Cornell University Press, Ithaca, NY.

Pears, D. (1987) *The False Prison: A Study of the Development of Wittgenstein's Philosophy*, volume 1, Oxford University Press, Oxford.

Rayo, A. (2002) "Frege's Unofficial Arithmetic," *The Journal of Symbolic Logic* 67, 1623–38.

Rayo, A. (2008a) "On Specifying Truth-Conditions," *The Philosophical Review* 117.

Rayo, A. (2008b) "Vague Representation," *Mind* 117, 329–373.

Stalnaker, R. C. (1984) *Inquiry*, MIT Press, Cambridge, MA.

Stalnaker, R. C. (1996) "On What Possible Worlds Could Not Be." In Morton and Stich (1996), 103–119. Reprinted in Stalnaker (2003), 40–54.

Stalnaker, R. C. (1999) *Context and Content*, Oxford University Press, Oxford.

Stalnaker, R. C. (2003) *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays*, Clarendon Press, Oxford.

Williams, J. R. G. (typescript) "Fundamental and derivative truths."

Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*, Routledge and Kegan Paul, London. Tanslation by C.K. Ogden. Published as "Logisch-Philosophische Abhandlung", in *Annalen der Naturphilosophische* Vol. XIV, 3/4, 1921, pp. 184–262.

Wright, C. (1983) *Frege's Conception of Numbers as Objects*, Aberdeen University Press, Aberdeen.

Wright, C. (2001) "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism," *Mind* 110, 45–98.

Yablo, S. (2002) "Abstract Objects: A Case Study," *Noûs* 36, supp. 1, 255–286. Originally appeared in Bottani et al. (2002).