Commitment in Mathematics

Agustín Rayo
philosophy.ucsd.edu/arayo

May 27, 2005

The purpose of this paper is to develop a philosophical tool and argue that it can be used to address certain puzzles in the philosophy of logic and mathematics. Section 1 provides some preliminary motivation for the project. Section 2 supplies a statement of the proposal and section 4 a defence. Sections 3, 6 and 7 discuss the proposal's applications. Section 5 is a metaphysical aside.

1 The Irrelevance Problem

Here is a problem in the philosophy of mathematics I borrow from Yablo (2001):

> *The Irrelevance Problem*
>
> There is an apparent mismatch between the commitments carried by mathematical sentences and the beliefs whereby non-philosophers take assertions of such sentences to be true.
>
> It is natural to suppose, for instance, that 'the number of unicorns is 0' carries commitment to the number zero. But whether or not non-philosophers take assertions of this sentence to be true depends on whether they believe that the world contains unicorns, not on whether they believe that the world contains numbers.
>
> The problem is not restricted to applied mathematics. Whether or not non-philosophers take a sentence like 'there are infinitely many primes' to be true will usually depend on whether they have been exposed to the relevant proof (or heard it from an appropriate authority), but not on whether they believe that the world contains numbers.

One way of reacting to the Irrelevance Problem is by countenancing the idea that non-philosophers are systematically mistaken about the truth-conditions of mathematical assertions. An alternative is what I shall call the *common-sense* response. The common-sense response is based on the thought that in typical mathematical assertions mathematical vocabulary is used as a representational aid: it facilitates talk about the assertion's *real content*, which is what the assertion is really about. And real content is said to carry no problematic commitments. In the case of 'the number of unicorns is 0', for instance, real content might be said to concern unicorns, and whether there are any, but not numbers. The common-sense response then proceeds to argue that the commitments of a mathematical assertion are to be ascertained on the basis of the assertion's real content. Different versions of the common-sense response have different stories to tell about why it is that mathematical assertions should be associated with real contents, rather than the ontologically-offensive contents one might have expected.

The purpose of this section is to argue that the most straightforward implementations of the common-sense response to the Irrelevance Problem are burdened by an important limitation. This will help motivate the proposal I develop later in the paper.

1.1 Eliminativism

Eliminativism consists of three main claims. The first is that mathematical sentences are not to be taken at face value: their semantic structure[1] is not to be read off from their surface grammatical structure. In particular, 'the number of unicorns is 0' needn't have a semantic structure of the form '$a = b$', as its surface grammatical structure suggests. The second eliminativist claim is that the hidden semantic structure of mathematical sentences is such that mathematical objects needn't fall within the range of the variables in order for mathematical truths to be true and mathematical falsehoods to be false. Eliminativists are not, however, revisionists about truth-value. Their third claim is that mathematical sentences have standard truth-values, even though they have non-standard semantic structures: they have the very truth-values that would be assigned to them on a standard (Platonist) interpretation.

Eliminativism might be thought to deliver a version of the common-sense response to the Irrelevance Problem. For Eliminativists think of the numerical-terms displayed by the surface grammatical structure of a mathematical sentence as representational aids, with no direct correlate at the level of real content. In order for real content to be unveiled, one must go beyond surface grammatical structure and look at the sentence's true semantic structure.

I would like to focus on the question of what an eliminativist assignment of semantic structures to mathematical sentences might consist in. In the case of sentences such as 'the number of unicorns = 0' it is relatively easy to find a candidate semantic-structure that meets the eliminativist's needs. For instance:

$$\neg\exists x(\textsc{Unicorn}(x))$$

But it is not so obvious how to cash out the proposal when it comes to more complex sentences, like 'there is an n such that the number of unicorns is n' or 'there are infinitely

[1]The *semantic structure* of a sentence is the structure on which a compositional assignment of truth-conditions to that sentence is based. (More on semantic structure below.)

many primes'. What one would like is for eliminativists to set forth a function f, mapping sentences to sentences, that satisfies the following conditions:[2]

(a) f is recursive;

(b) for ϕ a mathematical sentence, $f(\phi)$ is true if and only if ϕ is true on a standard (Platonist) interpretation;

(c) mathematical objects needn't fall within the range of $f(\phi)$'s variables in order for condition (b) to obtain.

The eliminativist's claim could then be put by saying that the semantic structure of an arbitrary mathematical sentence ϕ corresponds to the surface grammatical structure of the sentence $f(\phi)$.

For the special case of arithmetic, there are several functions that satisfy (a)–(c), given suitable assumptions. Here is a particularly simple example.[3] For any sentence ϕ in the language of arithmetic, let $f(\phi)$ be the *universal Ramseyfication* of ϕ. In other words, let $f(\phi)$ be the universal closure of:

$$(\mathcal{A} \rightarrow \phi)^*$$

(where $\mathcal{A}$ is a the result of conjoining a suitable list of arithmetical axioms, and ψ^* is the result of uniformly replacing the arithmetical vocabulary in ψ by variables of appropriate type).

Assuming the legitimacy of higher-order quantification, it is easy to verify that f will satisfy conditions (a)–(c) provided there are infinitely many objects within the range of the variables. But the proviso is essential. If there are only finitely many objects within the range of the variables, the antecedent of '$(\mathcal{A} \rightarrow \phi)^*$' is unsatisfiable. So $f(\phi)$ is true for any arithmetical sentence ϕ. In particular, $f(\phi)$ is true even if ϕ is false on a standard (Platonist) interpretation, which is in violation on condition (b).

This suggests that the universally Ramseyfying eliminativism does not constitute any real progress when it comes to the Irrelevance Problem. For universally Ramseyfying eliminativists can certainly deny that the truth of 'the number of the unicorns is 0' requires any numbers; but only by conceding that the falsity of a sentence like 'the number of unicorns is 17' requires infinitely many objects. And in the context of the Irrelevance

[2] I assume that clause (a) is required by the eliminativist's claim that the semantic structure of ϕ is given by $f(\phi)$ (since given appropriate background information, there ought to be an effective procedure for getting from a sentence's surface grammatical structure to its semantic structure). If f is not assumed to be recursive, then constraints b and c are trivial, since f could be stipulated to assign '$\forall x(x = x)$' to any sentence that is true on a Platonist interpretation and '$\exists x(x \neq x)$' to any sentence that is not. The need for clauses (b) and (c) follows immediately from the second and third claims in our characterization of eliminativism.

[3] See Hellman (1989), Hellman (1990) and Parsons (1990). For other examples, see Hodes (1984), Rayo (2002a) and Yablo (2002). On all these examples, the 'suitable assumption' is that there be infinitely many things within the range of the quantifiers.

Problem such a concession is fatal. For whether or not non-philosophers take an assertion of 'the number of unicorns is 17' to be false does not depend on whether they believe that the world contains infinitely many things, any more than it depends on whether or not they believe that the world contains numbers. The Irrelevance Problem has not been solved; it has simply been relocated.

Couldn't the eliminativist refine her proposal by going modal? The simplest way to proceed would be by adding a box in front of the universal closure of '$(\mathcal{A} \rightarrow \phi)^*$'. That would allow f to satisfy condition (b), provided only that there are infinitely many objects within the range of the variables *relative to some possible world or other.*

This simple modification is only effective for sentences of *pure* arithmetic ('the number of unicorns is 0' would turn out to be false, for example, since there are infinite worlds with unicorns). In order to accommodate applied arithmetic, something more elaborate is required. One option is to claim that $f(\phi)$ is a counterfactual whose antecedent is a sentence to the effect that there are infinitely many objects and whose consequent is the universal closure of '$(\mathcal{A} \rightarrow \phi)^*$'. Counterfactuals bring in complications. But if all goes well, this proposal will deliver the result that f satisfies condition (b), provided only that there are infinitely many objects within the range of the variables relative to some possible world or other.

Unfortunately, neither version of the modal version of universally Ramseyfying eliminativism constitutes any real progress when it comes to the Irrelevance Problem. For whether or not non-philosophers take an assertion of 'the number of unicorns is 17' to be false does not depend on whether they believe that the universe *could* contain infinitely many things, any more than it depends on whether or not they believe that the universe *does* contain infinitely many things. The point is easily obscured by the fact that almost everyone believes that there could be infinitely many objects. But if you are in any doubt about it, consider how the proposal would play out in the case of set theory: f would only satisfy condition (b) if, relative to some possible world or other, there are *inaccessibly* many objects within the range of the variables.[4] But whether or not physicists take assertions of 'the state-function of an undisturbed system varies in accordance with Schrödinger's Equation' to be true or false depends on their beliefs about sub-atomic particles, not on whether they have any beliefs about whether the world could contain inaccessibly many things—a matter best left to their colleagues in the philosophy department.[5] Once more, the Irrelevance Problem has not been solved; universally Ramseyfying eliminativists have simply relocated the problem.

So far, attention has been focused on one particular version of eliminativism. Might there not be a different version of eliminativism that fares better with respect to the Irrelevance Problem? There is a formal result that places significant limits on the eliminativist's

[4]In the case of set-theory, universally Ramseyfying eliminativists face the additional difficulty that it is not clear whether one can have a categorical axiomatizations of set-theory. (If a global version of the second-order Reflection Principle holds, one can show that a categorical axiomatization is impossible. See Rayo (2005).) For present purposes, however, one might take $\mathcal{A}$ to be the result of conjoining the axioms of second-order ZFC.

[5]A state-function is, of course, a certain type of set-theoretic construction.

options. On the assumption that the domain of f is to contain every sentence of elementary arithmetic and that sentences in the range of f are to be couched in some first- or higher-order language with finite non-logical vocabulary, one can prove that it is *impossible* to find an f that satisfies conditions (a) and (b) when there are only finitely many objects within the range of the variables.[6]

The result continues to hold when the sentences in the range of f are allowed to contain any of the standard pieces of modal vocabulary, but this time what it shows is that it is impossible to find an f that satisfies conditions (a) and (b) when there is a finite upper bound to the number of objects a possible world can contain.[7] (And, of course, when it comes to addressing the Irrelevance Problem, it wouldn't be much progress if it turned out that status of arithmetical sentences depends on whether there could be at least, e.g. 10^{474} objects; what matters for the Irrelevance Problem is whether the auxiliary hypothesis is *relevant*, not whether it is plausible.)

It is worth emphasizing that the limits imposed by these results cannot be overcome by adopting a piecemeal approach, according to which different paraphrase functions are to used for different sentences. As long as there is a recursive procedure for specifying which paraphrase function is used for which sentences, the result will apply to the combined approach just as it applies to the simple approach.

It cannot be excluded that eliminativists will one day come up with an essentially novel strategy, which puts them in a position to address the Irrelevance Problem. But I think it would be a mistake to be optimistic, given the current state of play.

1.2 Fictionalism

Unlike eliminativists, fictionalists might be prepared to concede that 'the number of unicorns is 0' contains singular terms referring to numbers, and that serious assertions of 'the number of unicorns is 0' commit the speaker to the existence of numbers. What fictionalists deny is that typical assertions of mathematical sentences are 'serious', in the relevant respect. They claim that typical mathematical assertions are advanced 'in a spirit of make-believe'. Here is an example of how the story might go.[8]

[6] *Proof Sketch:* Let L be the language in which sentences in the range of f are couched, and assume, for *reductio*, that (a) and (b) are both satisfied in some model M with a finite domain. Since L is an nth-order language for some finite n, since L has finite non-logical vocabulary, and since M has a finite domain, the set of (Gödel numbers of) sentences of L that are true in M is recursive. By (a), this means that the set of (Gödel numbers of) arithmetical sentences that are mapped by f to true sentences in M is recursive. But, in light of (b), it follows that the set of (Gödel numbers of) true arithmetical sentences is recursive, contradicting Gödel's Theorem.

[7] *Proof Sketch:* A sentence containing any of the standard pieces of modal vocabulary (including counterfactuals) can be paraphrased as a sentence with no modal vocabulary but quantifiers ranging over possible worlds. When there is a finite upper bound on the number of objects a possible world can contain, any model for the resulting language is equivalent to a finite model. So the previous result applies. (This assumes that the accessibility relation for '$\Box$' takes indistinguishable worlds to be equally accessible (as is trivially the case in S5), and that the 'closeness' relation for counterfactuals takes indistinguishable worlds to be 'equally close'.)

[8] I follow Yablo (2001).

> There are circumstances having to do with John's nervousness in which it would be correct to say of John that he has butterflies in his stomach, regardless of whether, speaking literally, there are any insects inside John. Call these circumstances $\mathcal{C}$. A fictionalist about stomach-butterfly-talk would claim that a typical assertion of 'John has butterflies in his stomach' is advanced in a spirit of make believe, and that one must distinguish between the assertion's *fictional content* and its *literal content*. Fictional content is given by $\mathcal{C}$; literal content is determined by those circumstances in which 'John has butterflies in his stomach' is literally true. The fictionalist would insist that the proposition communicated by a typical assertion of 'John has butterflies in his stomach' is given by the assertion's fictional content, rather than its literal content. In particular, what is communicated by the assertion is true just in case the conditions imposed by fictional content are satisfied—just in case John experiences the right sort of nervousness—regardless of whether the conditions imposed by literal content are satisfied.
>
> According to mathematical fictionalism, number-talk is to be thought of on the model of stomach-butterfly-talk. Typical assertions of 'the number of unicorns is 0' are advanced in a spirit of make-believe, and one must distinguish between their fictional content and their literal content. Their *fictional* content might be determined, for instance, by those circumstances in which the following is true:
>
> $$\neg\exists x(\textsc{Unicorn}(x)),$$
>
> whether or not there are any numbers. Their *literal* content, on the other hand, is determined by those circumstances in which 'the number of unicorns is 0' is literally true, and is therefore sensitive to the existence of numbers. The fictionalist would insist that the proposition communicated by a typical assertion of 'the number of unicorns is 0' is given by the assertion's fictional content, rather than its literal content. In particular, what is communicated by the assertion is true just in case the conditions imposed by fictional content are satisfied—just in case $\neg\exists x(\textsc{Unicorn}(x))$—regardless of whether the conditions imposed by literal content are satisfied.

A more general version of mathematical fictionalism might be characterized in terms of the following three claims. The first is that typical assertions of mathematical sentences are advanced in a spirit of make-believe, and that the proposition communicated by such assertions is given by the assertions' fictional content, rather than their literal content. The second fictionalist claim is that the fictional contents associated with typical mathematical assertions are free from problematic involvement with mathematical ontology. (For instance, the existence of mathematical objects is irrelevant to the question of whether or not the fictional content of typical assertions of 'the number of unicorns is 0' is satisfied.) Fictionalists are not, however, revisionists about truth-value.[9] The third fictionalist claim is that even though the propositions communicated by typical assertions of a mathematical sentences do not, in general, correspond to the assertions' literal content, their truth-values are standard: they have the very truth-values that would be assigned to the relevant sentences on a standard (Platonist) interpretation.

Fictionalism might be thought to deliver a version of the common-sense response to the Irrelevance Problem. Fictionalists take the real content of typical mathematical assertions to be their fictional content, rather than their literal content. Numerical vocabulary is simply a representational-aid used in the specification of real content, which is free from any offending commitments.

As illustrated above, it is relatively straightforward for the fictionalist to specify what the fictional content of a typical assertion of 'the number of unicorns is 0' might consist in. But, as in the case of eliminativism, it is not so obvious how to cash out the proposal when it comes to more complex sentences such as 'there is an n such that the number of unicorns is n' or 'there are infinitely many primes'.

Some fictionalists have been careful to specify a function f satisfying conditions (a)–(c) from section 1.1.[10] The fictionalist claim can then be put by saying that the fictional content—and hence the *real* content—of typical assertions of a mathematical sentence ϕ corresponds to the literal content of $f(\phi)$. Yablo (2002), for example, provides the following characterization of f for the special case of arithmetic.[11] One begins by specifying f for certain quantifierless sentences:

- $f(\ulcorner Num_x[Px] = n\urcorner) = \ulcorner\exists_n!x(Px)\urcorner$
- $f(\ulcorner Num_x[Px] \neq n\urcorner) = \ulcorner\neg\exists_n!x(Px)\urcorner$
- $\ulcorner f(n = m\urcorner) = \ulcorner\exists_n!x(Fx) \to \exists_m!x(Fx)\urcorner$
- $\ulcorner f(n \neq m\urcorner) = \ulcorner\exists_n!x(Fx) \to \neg\exists_m!x(Fx)\urcorner$
- $f(\ulcorner n + m = s\urcorner) = \ulcorner(\exists_n!x(Fx) \wedge \exists_m!x(Gx) \wedge \neg\exists x(Fx \wedge Gx)) \to \exists_s!x(Fx \vee Gx)\urcorner$
- $f(\ulcorner n + m \neq s\urcorner) = \ulcorner(\exists_n!x(Fx) \wedge \exists_m!x(Gx) \wedge \neg\exists x(Fx \wedge Gx)) \to \neg\exists_s!x(Fx \vee Gx)\urcorner$
- $f(\ulcorner n \times m = s\urcorner) =$
  $\ulcorner(\exists_n!x(F_1x) \wedge \ldots \exists_n!x(F_mx) \wedge \neg\exists x(F_1x \wedge \ldots F_mx)) \to \exists_s!x(F_1x \vee \ldots F_mx)\urcorner$
- $f(\ulcorner n \times m \neq s\urcorner) =$
  $\ulcorner(\exists_n!x(F_1x) \wedge \ldots \exists_n!x(F_mx) \wedge \neg\exists x(F_1x \wedge \ldots F_mx)) \to \neg\exists_s!x(F_1x \vee \ldots F_mx)\urcorner$

[9] There is, of course, a variant of what I here call fictionalism that does propose revisions in truth-value. Such a view faces a different set of challenges, which I shall not consider here.

[10] I assume that clause (a) would be presupposed by any plausible story about how it is that that finite beings like ourselves are able to grasp fictional content. The need for clauses (b) and (c) follows immediately from the second and third claims in our characterization of fictionalism.

[11] Yablo also provides a characterization of f for the case of set theory. But it is not immune from the sort of difficulty I discuss below.

The next step is to stipulate that $f(\ulcorner \exists x(\phi(x))\urcorner)$ is to be the infinite disjunction of each of the $f(\ulcorner \phi(n)\urcorner)$ and that $f(\ulcorner \forall x(\phi(x))\urcorner)$ is to be the infinite conjunction of each of the $f(\ulcorner \phi(n)\urcorner)$. This doesn't quite deliver the result that f is defined for every arithmetical sentence. But, as Yablo notes, every sentence in the language in the language of first-order pure and applied arithmetic is logically equivalent to some sentence for which f is defined. So we needn't worry about the omissions.

Yablo's characterization of f yields the following result: provided there are infinitely many objects within the range of the variables, conditions (a)–(c) from section 1.1 are all satisfied. Unfortunately, the proviso is essential. If there are only n objects within the range of the variables, then, e.g. $f(\ulcorner n+1=n\urcorner)$ is trivially true, even though $\ulcorner n+1=n\urcorner$ is false on the standard (Platonist) interpretation. So condition (b) is violated.

This means that a fictionalist proposal based on Yablo's f constitutes no real progress when it comes to the Irrelevance Problem. It is certainly true that Yablonian fictionalists can deny that typical mathematical assertions involve any commitment to numbers; but only by conceding that the falsity of typical assertions of '$\exists n(n+1=n)$' requires infinitely many objects within the range of the variables. And in the context of the Irrelevance Problem such a concession is fatal. For whether or not non-philosophers take an assertion of '$\exists n(n+1=n)$' to be false does not depend on whether they believe that the world contains infinitely many individuals, any more than it depends on whether or not they believe that the world contains numbers. So the Irrelevance Problem has not been solved; it has simply been relocated.

Of course, fictionalists could attempt to specify f by using some method other than Yablo's. But it is important to be clear that the formal result of section 1.1 applies in the case of fictionalism just as it applied in the case of eliminativism.

1.3 Other Options

There are other ways in which the common-sense response to the Irrelevance Problem might be implemented. Here are some examples:

1. Pragmatic Instrumentalism[12]

   Suppose we are trying to determine whether guests at the party are having a good time. You say 'the guy in the red shirt looks bored'. I reply 'actually, that's a burgundy shirt'. You respond impatiently: 'all I meant is that that guy is bored'. A pragmatic instrumentalist would claim that the reason for your impatience is that the content a speaker intends to convey with an assertion is the assertion's *instrumental content*: the contribution the assertion makes towards forwarding the goals of the conversation. Since the goal of our conversation is to determine whether guests are having a good time, part of the instrumental content of your original assertion is that someone is bored. But it is no part of the assertion's instrumental content

[12] See Eklund (forthcoming). Steve Yablo discussed a similar view in a presentation of 'Non-catastrophic Presupposition Failuire' (University of California, San Diego, November 2004).

that the person in question is wearing a red shirt, since that doesn't forward our conversational goals.

A similar story can be told about typical mathematical assertions. John wants to know whether there are unicorns. You say 'the number of unicorns is 0'. A nominalist interjects: 'actually, that's not right because there aren't really any numbers'. You respond impatiently: 'all I meant is that there are no unicorns'. Pragmatic instrumentalists would concede that the proposition expressed by your original assertion implies that there are numbers. But they would insist that instrumental content was all your assertion intended to convey, and that the existence of numbers is no part of your assertion's instrumental content because it doesn't forward the relevant conversational goals.

More generally, pragmatic instrumentalists claim that the content set forth by a typical mathematical assertion is the assertion's instrumental content, and that instrumental content is free from offending commitments. But pragmatic instrumentalists are not revisionists about truth-value. They also claim that the instrumental contents of typical mathematical assertions have the very truth-values that would be assigned to the relevant sentences on a standard (Platonist) interpretation.

Pragmatic Instrumentalism might be thought to deliver a version of the common-sense response to the Irrelevance Problem. Pragmatic Instrumentalists take the real content of typical mathematical assertions to be their instrumental content. Numerical vocabulary is simply a representational-aid used in the specification of real content, which is free from any offending commitments.

2. Empiricism

   According to the brand of empiricism I would like to consider,[13] a typical mathematical assertion does not commit the speaker to the *truth* of the proposition expressed, it only commits speakers to the proposition's *empirical adequacy*. Let the *empirical content* of an assertion consist of whatever conditions need to be fulfilled in order for the proposition expressed by the assertion to be empirically adequate. One can then put the empiricist's claim by saying that the content set forth by a typical mathematical assertion is the assertion's empirical content, and that empirical content is free from offending commitments. But pragmatic instrumentalists are not revisionists about truth-value. They also claim that the empirical contents of typical mathematical assertions have the very truth-values that would be assigned to the relevant sentences on a standard (Platonist) interpretation.

   This brand of empiricism might be thought to deliver a version of the common-sense response to the Irrelevance Problem. Empiricists take the real content of typical mathematical assertions to be their empirical content. Numerical vocabulary is simply a representational-aid used in the specification of real content, which is free from any offending commitments.

[13] It is loosely based on Van Fraassen (1980).

3. Reconceptualization

   Neo-Fregeans have defended the idea that semantic content may be 'conceptualized' or 'carved up' in different ways.[14] They claim, in particular, that 'the number of the Fs is the number of the Gs' has the same content as 'the Fs are in one-one correspondence with the Gs', only 'conceptualized' in a different way.

   By going beyond the standard Neo-Fregean agenda, one could try to argue that the reconceptualization thesis yields a version of the common-sense response whereby the real content of typical mathematical assertions is given by the result of reconceptualizing the asserted sentences.

4. Azzouni-content

   Jody Azzouni has developed a theory of ontological commitment and existence.[15] It is based on two main ideas. The first is the thesis that "ontological commitment is to be carried not by the quantifiers but by an existence predicate". Thus, although quantifiers are taken to play inference-licensing and anaphora-supporting roles, the truth of sentences of the form '$\exists x\, Fx$' does not generally license the conclusion that the world contains Fs: one needs to know, in addition, whether '$\forall x(Fx \rightarrow Ex)$' is true, where '$E$' is the existence predicate. The second idea consists of a criterion for deciding when sentences of the form 'Ex' are true. Specifically, he suggests that '$\forall x(Ex \leftrightarrow Ix)$' should be taken to be true, where 'I' means something like 'not ontologically dependent on any linguistic or psychological process'.[16] By arguing that numbers are dependent on our linguistic and psychological processes, Azzouni uses these two ideas to defend the view that 'there are numbers' is true even though numbers don't exist.

   Let the *Azzouni-content* of an assertion consist of whatever conditions need to be fulfilled by existent objects order for the assertion to be counted as true by Azzouni's proposal. Although Azzouni himself doesn't set out to do so, one might wish to set forth a version of the common sense response by arguing that the real content of typical mathematical assertions is their Azzouni-content. Numerical vocabulary is simply a representational-aid used in the specification of real content, which is free from any offending commitments.

Much as before, a problem emerges when one tries to explain in detail what real content consists in according to such views. For although it is easy to identify real content in certain special cases, the proposals under consideration do not make clear how real content is supposed to be specified in general. One could attempt to set forth a function f satisfying

[14] See Hale and Wright (2000) pp. 148-150, Wright (1997) pp. 207-209, Hale (1997), Hale (1999) and Wright (1983). The original idea is due to Frege; see Frege (1884) §64.

[15] See, especially, Azzouni (1997), Azzouni (1998) and Azzouni (2004).

[16] It is argued, however, that there are competing criteria "which can't be rationally adjudicated among" (Azzouni (2004), p. 82). The quotation above is taken from Azzouni (2004), p. 49. For a statement of the second idea, see pp. 112-3.

conditions (a) and (b) from section 1.1, and go on to claim that the real content of typical assertions of a mathematical sentence ϕ corresponds to the content that is straightforwardly expressed by $f(\phi)$.[17] But the formal result of section 1.1 will apply here just as it applied before.[18] (And, just as before, there are limits to what modal maneuvering can achieve. One could certainly claim that the real content of, e.g. 'the number of unicorns is 0' is something like 'as far as our conversational goals are concerned, the world is just as it would be if there were numbers and the number of the unicorns were 0', or 'as far as empirical testability is concerned, the world is just as it would be if there were numbers and the number of unicorns were 0'. But in the context of the Irrelevance Problem, a commitment to the existence of numbers is just as problematic as a commitment to the existence of numbers relative to some possible world or other.)

I have argued that the most obvious implementations of the common sense response to the Irrelevance Problem are burdened by an important limitation. I hope to have made clear, moreover, that the problem is largely independent of the philosophical positions motivating any particular implementation: it stems instead from an inherent difficulty in the project of identifying the real content of an arithmetical sentence ϕ (or typical assertions thereof) with the content expressed by some related sentence $f(\phi)$.

2 The proposal

It seems to me that philosophers have often made an unwarranted assumption in thinking about the relationship between commitment and semantics.[19] In this section I shall try to identify the assumption and develop an alternate way of thinking about commitment and

[17] As before, I assume that clause (a) would be presupposed by any plausible story about how it is that that finite beings like ourselves are able to grasp real content. The need for clause (b) follows immediately from our characterizations of pragmatic instrumentalism and empiricism.

[18] The machinery developed in Hofweber (forthcoming) and Hofweber (typescript) can be used to construct a proposal that might allow one to escape the limitation imposed by the formal result. The proposal would consist of three parts. The first is the thesis that a (suitable reading of a) sentence like 'two unicorns and two (more) unicorns make four unicorns' carries no commitment to unicorns (or possible unicorns). The second is the thesis that what is said by (a suitable reading of) a statement like 'two and two are four' is something like: two X and two (more) X are four X, for whatever X. The third is the thesis that (a suitable reading of) 'there is a number n such that . . . n . . .' has the same truth-conditions as the infinite disjunction 'either . . . zero . . . , or . . . one . . . , or . . . two, . . . , or . . . ', and (a suitable reading of) 'every number n is such that . . . n . . .' has the same truth-conditions as the infinite conjunction '. . . zero . . . , and . . . one . . . , and . . . two, . . . , and . . . '. By putting the first two theses together, one can characterize a function f satisfying conditions (a) and (b) from section 1.1 for every quantifier-free sentence in the language of arithmetic. And one can use the third thesis to extend the domain of f in such a way that every sentence in the language of arithmetic is equivalent to some sentence for which f is defined. Although the proposal strikes me as promising, most of Hofweber's efforts so far have focussed on arguing for ideas leading to theses two and three. A full assessment of the proposal must await further discussion of the (potentially controversial) first thesis.

[19] Notable exceptions are Azzouni (1997), Azzouni (2004), Hofweber (forthcoming) and Hofweber (typescript).

semantics.

2.1 Containment Assumptions

On one way of thinking about semantics, the semantic value of an expression is—at a first approximation[20]—a function from possible worlds to referents: the semantic value of a singular term is a function from possible worlds to objects; the semantic value of a predicate is a function from possible worlds to extensions; the semantic value of a sentence is a function from possible worlds to truth-values; the semantic value of a quantifier-expression is a function from possible worlds to sets of sets of objects; the semantic value of an n-place sentential connective is a function from possible worlds to functions from n-tuples of truth-values to truth-values; and so forth. The point I wish to draw attention to is that semantic values are sometimes taken to play a dual role in our semantic theories. On the one hand, they play a *compositional role*: the role of enabling a compositional assignment of truth-conditions on the basis of semantic structure; on the other, they play an *ontological role*: the role of determining which objects must be contained in the world in order for a given sentence to be true. Some examples should make this clear.

Let me begin by considering a case in which a semantic value plays a compositional role without playing an ontological role. Consider the sentence 'Susan runs'. On the standard way of proceeding, the semantic value of 'Susan runs' is the function f_{SR}, which assigns the value TRUE to a world w if Susan runs at w and the value FALSE otherwise. It is clear that f_{SR} plays a compositional role: it enables one to say, for example, that 'it is not the case that Susan runs' is true at world w if and only if it is not the case that $f_{SR}(w) = \text{TRUE}$. But f_{SR} does *not* play an ontological role, in the following sense: it is not the case that w must *contain* the truth-value $f_{SR}(w)$ in order for 'Susan runs' to be true at w. Accordingly, $f_{SR}(w)$ is best thought of as a piece of *semantic machinery*, rather than a carrier of ontological commitment.

Compare this with the case of names and predicates. On the standard way of proceeding, the semantic value of the name 'Susan' is the function f_{Susan}, which maps each possible world that contains Susan to Susan (and is otherwise undefined), and the semantic value of the predicate 'runs' is the function f_{runs}, which maps each possible world w to the set objects that run at w. As before, it is clear that f_{Susan} and f_{runs} play compositional roles: one can say, for example, that 'Susan runs' is true at world w just in case $f_{Susan}(w) \in f_{runs}(w)$. But f_{Susan} also plays an ontological role: w must *contain* $f_{Susan}(w)$ in order for 'Susan runs' to be true at w. The ontological role of $f_{runs}(w)$ is somewhat different. Although w needn't contain $f_{runs}(w)$ in order for 'Susan runs' to be true at w, it must contain every individual in $f_{runs}(w)$. One can think of $f_{runs}(w)$ as part semantic machinery and part carrier of ontological commitment.

The case of quantifiers is more like the case of predicates than like the case of names or the case of sentences. On the standard way of proceeding, the semantic value of the

[20]This is overly simplistic in various respects. It neglects contextual dependence, for example, and does not allow for the possibility of intensional modifiers such as 'alleged' or 'necessarily'. I ignore the additional complexity for ease of presentation. For a more developed picture, see Lewis (1970).

existential quantifier is the function $f_\exists$, which maps each world w to the set of non-empty sets of individuals x such that x is contained in w. So, although w needn't contain $f_\exists(w)$ in order for, e.g. '$\exists x$ (x runs)' to be true at w, it must contain every individual in every member of $f_\exists(w)$. Again, $f_\exists(w)$ might be thought of as part semantic machinery and part carrier of ontological commitment.

The case of sentential connectives is like the case of sentences: there is a compositional role with no ontological role. On the standard way of proceeding, the semantic value of 'it is not the case that' is the function $f_\neg$, which maps each world to the function that takes TRUE to FALSE and FALSE to TRUE. $f_\neg$ plays no ontological role because it is not the case that w must *contain* the function $f_\neg(w)$ in order for sentences involving negations to be true at w.

Let me introduce some notation. Say that a *containment assumption* is a specification of the ontological role that is to be played by a given piece of semantic machinery. Thus:

- the standard containment assumption for the semantic values of sentences is that the referent assigned to a possible world w (i.e. a truth-value) needn't be contained in w;
- the standard containment assumption for the semantic values of names is that the referent assigned to a possible world w (i.e. an individual) must be contained in w;
- the standard containment assumption for the semantic values of predicates is that, although the referent assigned to a possible world w (i.e. a set of individuals) needn't be an object contained in w, every member of the referent must be contained in w;
- the standard containment assumption for the semantic values of quantifiers is that, although the referent assigned to a possible world w (i.e. a set of non-empty sets of individuals) needn't be an object contained in w, every member of a member of the referent must be contained in w;
- the standard containment assumption for the semantic values of the sentential connectives is that the referent assigned to a possible world w (i.e. a function from n-tuples of truth-values to truth-values) needn't be contained in w.

When it comes to sentences like 'Susan runs', the standard containment assumptions are eminently sensible. But I would like to suggest that there are certain kinds of sentences—including mathematical sentences—for which the standard containment assumptions should be given up.[21]

[21] A word about possible-world-talk. My proposal does not presuppose any particular story about how it is that possible worlds represent. (For example, it does not presuppose, as in Lewis (1986), that a possible world represents the world as containing water just in case the possible world itself contains water.) It is for ease of exposition that I say 'w contains x' instead of 'w represents the world as containing x' and 'x is P at w' instead of 'w represents the world as being such that x is P'. Talk of containment is itself a notational abbreviation: 'w represents the world as containing x' is true just in case there is some F such that 'w represents the world as being such that x is F' is true.

2.2 A Criterion of Ontological Commitment

In his classic paper on ontological commitment, Quine proposed that a first-order sentence carries commitment to Fs if and only if Fs must be admitted among the values of the variables in order for the sentence to be true.[22] Quine's criterion is inappropriate for present purposes because it is tailored to conform to the standard containment assumptions. I would like to suggest that we work with the following instead:[23]

> *The Semantic Criterion*
> A sentence ϕ carries commitment to Fs just in case: (*i*) one's semantic theory entails that ϕ is only true at worlds containing Gs, and (*ii*) necessarily, all Gs are Fs.[24]

In Quine's criterion the values of the variables are assumed to play a dual role: on the one hand they play a compositional role by enabling a compositional assignment of truth-conditions; on the other, they play an ontological role by determining which objects must be contained in the world in order for a sentence to be true. The Semantic Criterion, by contrast, makes no presuppositions about the ontological role of the various pieces of semantic machinery. It lets the matter be decided by the containment assumptions embedded in one's semantic theory.[25]

2.3 Arithmetic

In this section I will illustrate how non-standard commitment assumptions can be put to use in the case of mathematical discourse. We will confine our attention to three expressions

[22] See Quine (1948).

[23] For a related criterion see Cartwright (1954).

[24] Here 'entails' is taken to mean 'logically entails given auxiliary assumptions'. An auxiliary assumption may be any true sentence in the language of one's semantic theory, except for sentences including locutions of the form 'p is true at world w' (or notational variants).

[25] Even in the special case in which one's semantic theory is assumed to be based on the standard containment assumptions, there are certain cases in which Quine's criterion and the Semantic Criterion come apart. Just how this happens depends on how Quine's criterion is disambiguated.

Suppose, first, that Quine's 'must' is cashed out in terms of metaphysical necessity (i.e. 'Fs must be admitted among the values of the variables in order for the sentence to be true' is read 'necessarily, every model of the sentence in which the non-logical vocabulary receives its intended interpretation has a domain containing Fs'). Then, on the assumption that Prince Charles couldn't have had a different mother, it is a consequence of Quine's criterion that, e.g. 'Charles is taller than Elizabeth' is committed to mothers, even thought Elizabeth might not have been a mother. But this is not a consequence of the Semantic Criterion. For, in the absence of assumptions about what is true in the various possible worlds, one's semantic theory will not entail (in the relevant sense) that there are no worlds in which Charles is taller than Elizabeth but there are no Gs (where G is such that, necessarily, all Gs are mothers). So it will not entail (in the relevant sense) that 'Charles is taller than Elizabeth' can only be true in worlds containing Gs.

Now suppose that Quine's 'must' is cashed out in terms of logical consequence (i.e. 'Fs must be admitted among the values of the variables in order for the sentence to be true' is read 'every model of the sentence has a domain with objects that are Fs according to the model'). Then a sentence like 'there are numbers' is committed to numbers but not, e.g. to abstract objects. This result is averted on the Semantic Criterion because, necessarily, every number is an abstract object.

in the language of applied arithmetic: the numeral '0', the non-arithmetical predicate 'is a unicorn' and the number operator '$Num(F, n)$' (read 'the number of Fs is n').

If one were to supply a semantic theory for arithmetic on the basis of the standard containment assumptions, one might proceed by assigning the semantic values to '0', 'is an elephant' and '*Num*' in the following manner:

(a) the semantic value of '0' is the function f_0 that assigns the number zero to each possible world containing the number zero (and is otherwise undefined).

(b) the semantic value of 'is a unicorn' is a function f_u that assigns to each possible world w the set of individuals x such that x is a unicorn at w.

(c) the semantic value of 'Num' is the function f_{Num} that assigns to each possible world w the set of ordered pairs $\langle n, x\rangle$ such that: (1) x is a set of individuals contained in w, (2) there are precisely n objects in x, and (3) w contains the number n.

Given a semantic theory based on (a)–(c), it follows from the Semantic Criterion that 'the number of unicorns is 0' is committed to numbers. Such a semantic theory will also yield the result that 'the number of unicorns is 0' is true in w just in case w contains no unicorns, but only given a metaphysical assumption: the assumption that every possible world that contains no unicorns also contains the number zero.

The situation changes dramatically when (a) and (c) are reformulated in accordance with *non-standard* containment assumptions:

(a') the semantic value of '0' is the function f_0 that assigns the number zero to each possible world w, *regardless of whether* w *contains any numbers.*

(c') the semantic value of 'Num' is the function f_{Num} that assigns to each possible world w the set of ordered pairs $\langle n, x\rangle$ such that: (1) x is a set of individuals contained in w, and (2) there are precisely n objects in x; *all of this, regardless of whether* w *contains any numbers.*

Given a semantic theory based on (a'), (b) and (c'), it follows from the Semantic Criterion that 'the number of unicorns is 0' carries no commitment to numbers. (It also follows that 'the number of unicorns is 0' carries no commitment to infinitely many individuals.[26]) And such a semantic theory yields the result that 'the number of unicorns is 0' is true in w just

[26]On its current formulation, the Semantic Criterion cannot be used to assess the question of whether a sentence carries commitment to infinitely many individuals. For 'F' is assumed to be a first-order predicate. But it can easily be extended:

> A sentence ϕ carries commitment to objects collectively satisfying $\mathcal{F}$ just in case: (i) one's semantic theory entails that ϕ is only true at worlds containing individuals collectively satisfying $\mathcal{G}$, and (ii) necessarily, if some objects are $\mathcal{G}$ then they are $\mathcal{F}$.

where $\mathcal{F}$ and $\mathcal{G}$ are second-level predicates, such as 'are infinitely many in number'.

in case w contains no unicorns, but now the result holds independently of any metaphysical assumptions.

The result can be generalized. (Appendix A contains an unabridged version of the relevant semantic theory.) With the non-standard containment assumptions in place, the Semantic Criterion implies that arithmetical sentences generally carry no problematic ontological commitments. And the semantic theory assigns the desired truth-conditions to every sentence in the language of arithmetic whether or not one assumes that possible worlds contain any abstract objects and whether or not one assumes they contain infinitely many individuals.

With the non-standard containment assumptions in place, the semantic values of arithmetical expressions might be compared to the semantic values of the logical connectives: what we have a piece of semantic machinery that plays a compositional role, but no ontological role. Think of the matter as follows. Even though the semantic value of 'it is not the case that' (relative to a world) is a function from truth-values to truth-values, a sentence like 'it is not the case that Susan runs' *is not about functions*; it is about Susan, and whether she runs. Something similar is true of the arithmetical vocabulary when the non-standard containment assumptions are in place. Even though the semantic value of '0' (relative to a world) is a number, a sentence like 'the number of unicorns is 0' *is not about numbers*; it is about whether there are any unicorns.

2.4 Truth-conditions

Say that the *truth-set* of a sentence is the set of possible worlds at which the sentence is true. I would like to suggest that the *truth-conditions* assigned to a sentence by a semantic theory ought to be identified with the condition a world must meet according to the semantic theory in order to be a member of the sentence's truth-set. The purpose of this section is to explain how such a condition might be ascertained.

To see that the matter is not straightforward, note is that both of the following are consequences of a semantic theory based on clauses (a'), (b) and (c') from section 2.3:

1. A possible world w is in the truth-set of 'the number of unicorns is 0' just in case there is a set x such that: (1) the members of x are all and only the individuals that are contained in w and are unicorns according to w, and (2) the number of x's members is 0.

2. A possible world w is in the truth-set of 'the number of unicorns is 0' just in case w contains no unicorns.

On the face of it, 1 and 2 are not equivalent, since only the former involves reference to sets and numbers. So how could either of them capture *the* condition a possible world must meet, according to the relevant semantic theory, in order to be a member of the truth-set of 'the number of unicorns is 0'?

I shall say that two formulations *express the same condition* just in case they are interderivable on the basis of auxiliary assumptions, where an auxiliary assumption may be any true sentence in the language of one's semantic theory, except for sentences including locutions of the form 'p is true at world w' (or notational variants). 1 and 2 are therefore counted as distinct formulations of the same condition. But the following, which is a consequence of a semantic theory based on (a)–(c), is counted as a formulation of a different condition:

3. A possible world w is in the truth-set of 'the number of unicorns is 0' just in case there is a set x such that: (1) the members of x are all and only the individuals that are contained in w and are unicorns according to w, (2) the number of x's members is 0, and (3) w contains the number 0.

For 3 cannot be derived from 1 or 2 without making substantial assumptions about what is true according to the various possible worlds. (In particular, one would need the assumption that every world containing no unicorns is also a world containing the number zero.)

This sets the stage for the following observation. Although it is certainly true that a semantic theory based on clauses (a'), (b) and (c') from section 2.3 contains number-talk, talk of any objects with the structure of the numbers would have done just as well, in the sense that the resulting semantic theory would have imposed the very same conditions on truth-set membership. Suppose, for example, that number-talk had been replaced by von Neumann-ordinal-talk. Suppose, in particular, that clauses (a') and (c') had been replaced by the following:

(a^*) the semantic value of '0' is the function f_0 that assigns to each possible world w the empty set, regardless of whether w contains any sets.

(c^*) the semantic value of 'Num' is the function f_{Num} that assigns to each possible world w the set of ordered pairs $\langle \alpha, x \rangle$ such that: (0) α is a finite Von Neumann ordinal, (1) x is a set of individuals contained in w, and (2) there is a one-one correspondence between α and x; all of this, regardless of whether w contains any sets.

Then the following would have then been a consequence of one's semantic theory:

4. A possible world w is in the truth-set of 'the number of unicorns is 0' just in case there is a set x such that: (1) the members of x are all and only the individuals that are contained in w and are unicorns according to w, and (2) x can be put in one-one correspondence with the empty set.

And 4 can be shown to express the same condition as 1 and 2, since it is interderivable with them on the basis of auxiliary assumptions but no assumptions about what is true according to the various possible worlds. Notice, however, that the analogous result is *not* available when it comes to a semantic theory based on clauses (a)–(c). Were number-talk to be replaced by talk of von Neumann ordinals, one would get the following instead of 3:

5. A possible world w is in the truth-set of 'the number of unicorns is 0' just in case there is a set x such that: (1) the members of x are all and only the individuals that are contained in w and are unicorns according to w, (2) x is in one-one correspondence with the empty set, and (3) w contains the empty set.

And 5 is not interderivable with 3 without substantial assumptions about what is true according to the various possible worlds. (In particular, one would need the assumption that every world containing no unicorns contains the number zero just in case it contains the empty set.) So 5 and 3 do not express the same condition on truth-set membership.

When the truth-conditions assigned to a sentence by a semantic theory are identified with the condition a world must meet according to the semantic theory in order to be a member of the sentence's truth-set, we get the following result. According to a semantic theory based on the standard containment assumptions, the numbers are part of the truth-conditions of arithmetical expressions. But according to a semantic theory based on the non-standard containment assumptions they are not. They are better thought of as a (replaceable) part of the semantic machinery whereby truth-conditions are assigned to arithmetical sentences by a particular semantic theory.

Moreover, on a semantic theory based on the non-standard containment assumptions, 'the number of unicorns is zero' and 'there are no unicorns' have the same truth-conditions: in both cases the condition imposed on a possible world w to be a member of the relevant truth-set is that there be no x such that w contains x amongst its unicorns. The difference between them is a difference in semantic structure. Similarly, arithmetical truths like '5+7 = 12' or 'there are infinitely many primes' have the same truth-conditions as '$\forall x\, x = x$': in each case the condition imposed on a possible world w to be a member of the relevant truth-set is trivial. Again, the difference between them is a difference in semantic structure. In this respect, the difference between '5+7 = 12' and 'there are infinitely many primes', or between 'the number of unicorns is zero' and 'there are no unicorns', is not unlike the difference between '$\phi \wedge \psi$' and '$\neg(\neg\phi \vee \neg\psi)$'.

2.5 Anything Goes?

Might one use a semantic theory based on non-standard containment assumptions to argue that a sentence like 'Susan runs' has less ontological commitments than one might have thought? Not without paying a hefty price. For a semantic theory that alters the commitments of 'Susan runs' can be expected to assign the sentence unreasonable truth-conditions.

Consider an example. The following assignments of semantic value are based on the standard containment assumptions:

(A) the semantic value of 'Susan' is the function f_{Susan}, which assigns Susan to each possible world containing Susan (and is otherwise undefined).

(B) the semantic value of 'runs' is the function f_{runs}, which assigns to each possible world w the set of individuals x such that x runs at w.

It is straightforward enough to produce a version of (A) that is based on non-standard containment assumptions:

(A') the semantic value of 'Susan' is the function f_{Susan}, which assigns Susan to each possible world w, *regardless of whether* w *contains Susan.*

But what should one do with (B)? If it is left unmodified, then 'Susan runs' receives the right truth-conditions, but carries commitment to Susan, and to runners.[27] If, on the other hand, f_{runs} is taken to assign worlds a certain set x regardless of which objects the worlds contain, then 'Susan runs' will turn out to be necessarily true or necessarily false (depending on whether Susan is a member of x). So 'Susan runs' is assigned the wrong truth-conditions.

Here is a more devious example. Suppose one thinks that the existence of Susan supervenes on the existence of particles 'arranged Susanishly'. Then one might modify (A) and (B) as follows:

(A^*) the semantic value of 'Susan' is the function f_{Susan}, which assigns Susan to each possible world containing particles arranged Susanishly (and is otherwise undefined).

(B^*) the semantic value of 'runs' is the function f_{runs}, which assigns to each possible world w the set of sets of particles that are arranged runningly at w.

With (A^*) and (B^*) in place, the Semantic Criterion will yield the result that 'Susan runs' is committed to particles arranged Susanishly and to particles arranged runningly, but not to Susan or to runners. The problem is that 'Susan runs' gets assigned the wrong truth-conditions. For whereas a semantic theory based on (A) and (B) yields:

1. A possible world w is in the truth-set of 'Susan runs' just in case: (1) w contains Susan, and (2) Susan runs at w;

a semantic theory based on (A^*) and (B^*) yields:

2. A possible world w is in the truth-set of 'Susan runs' just in case there is a set of particles such that: (1) each of the particles is contained in w, (2) the particles are arranged Susanishly at w, and (3) the particles are arranged runningly at w.

And 1 and 2 do not express the same condition, since they are not interderivable without substantial assumptions about what is true according to the various possible worlds. (In particular, one would need the assumption that a world contains Susan just in case it contains particles arranged Susanishly.)

Whereas 1 expresses the desired truth-condition that 'Susan runs' be true only in worlds where Susan runs, 2 expresses the undesired truth-condition that 'Susan runs' be true only in worlds where particles arranged Susanishly are arranged runningly.

[27] This relies on the fact that 'x is contained in w' follows from 'x runs at w'. See footnote 21.

3 Back to the Irrelevance Problem

A friend of the common-sense response to the Irrelevance Problem can claim that the 'real content' of an arithmetical sentence is to be identified with the truth-conditions that the sentence gets assigned according to a semantic theory based on the non-standard containment assumptions.

Although certain *specifications* of the condition for membership in the truth-set of a sentence like 'the number of unicorns is 0' involve number-talk, the condition itself imposes no problematic requirements on the world. When non-standard containment assumptions are in place, all that is required for a possible world w to be in the truth-set of 'the number of unicorns is 0' is that w contain no unicorns (or, equivalently, that the number of individuals x such that w contains x amongst its apostles be 0); it is not required, in particular, that w contain numbers or infinities, or have any other offensive characteristics.

The numerical vocabulary in typical arithmetical sentences is a representational aid: it enables the sentences to have the right sorts of truth-conditions while enjoying an interesting semantic structure. It allows a sentence like 'the number of unicorns is 0' to have the same truth-conditions as 'there are no unicorns', and sentences like '5+7 = 12' or 'there are infinitely many primes' to have the same truth-conditions as '$\forall x\, x = x$'. But, at the same time, it endows arithmetical sentences with semantic structures that are succinct, perspicuous and rich in inferential possibilities.

There are a couple of issues that deserve comment:

1. It is a consequence of Gödel's Incompleteness Theorem that there is no recursive function assigning each truth of pure arithmetic to the trivial condition on truth-set membership and each falsity of pure arithmetic to the impossible condition on truth-set membership. Does it follow that a general specification of truth-conditions for the language of arithmetic must be beyond the reach of finite beings like ourselves? Fortunately not. All it takes to produce the relevant specification is a recursive function that assigns to each sentence in the language of arithmetic a *specification* of truth-conditions (rather than the truth-conditions themselves, whatever that would amount to). And there are certainly such a functions. To produce a compositional semantic theory for the language is, in effect, to characterize such a function.

   Of course, one may be in a position to specify the truth-conditions of a sentence without being able to put the information in a useful form. It is easy enough, for instance, to specify truth-conditions for a sentence GC expressing Goldbach's Conjecture. To wit: if every even number greater than 2 is the sum of two primes then every possible world is a member of GC's truth-set; otherwise no possible world is a member of GC's truth-set. But, of course, the information that is delivered by such a specification of truth-conditions is not in a very useful form. Conspicuously, it does not immediately enable us to determine whether GC is true.

2. I emphasized throughout section 1 that familiar implementations of the common-sense response are subject to an important limitation. I observed, in particular, that

there is an inherent difficulty in the project of defining a function f that maps each arithmetical sentence to some first- or higher-order sentence and going on to claim that the real content of typical assertions of ϕ is the content that is straightforwardly expressed by $f(\phi)$. How does the present proposal escape this difficulty?

The difficulty is averted because the present proposal specifies real content in an essentially different way. Rather than mapping arithmetical sentences to sentences that are meant to *express* real content, arithmetical sentence are mapped to sentences that are meant to *state* what real content consists in by stating a condition for truth-set membership.

4 Why accept non-standard containment assumptions?

A semantic theory based on non-standard containment assumptions is a philosophical tool. It can be used for different purposes, and by proponents of different views. A fictionalist, for instance, could claim that the fictional content of a mathematical assertion is to be identified with the truth-conditions that are assigned to the asserted sentence by a semantic theory based on non-standard containment assumptions. (And, of course, Pragmatic Instrumentalists could make a similar claim about instrumental content, Empiricists about empirical content, Neo-Fregeans about the content of reconceptualizations and friends of Azzouni's proposal about Azzouni-content.) As far as the purposes of this paper are concerned, any such implementation of the proposal would do.

In this section I would like to illustrate one particular way in which the use of a semantic theory based on non-standard containment assumptions might be defended. I do not claim that it is the only one, or the best.

4.1 The Minimalist Conception of Semantics

According to the implementation I would like to consider, the *literal* truth-conditions of arithmetical sentences should be characterized by a semantic theory based on non-standard containment assumptions. To illustrate how such a claim might be motivated, I will make use of a particular conception of semantic theorizing. The purpose of this subsection is to say what the conception consists in.

Our semantic theorizing is subject to constraints coming from two different directions. On the one hand, it is constrained by the grammatical structure of the sentences in our language; on the other, it is constrained by the way in which sentences are put to use.

Let me explain. The grammatical structure of our sentences imposes constraints on our semantic theorizing because one of the things we expect from a semantic theory is a compositional assignment of truth-conditions that applies to every sentence in the language and that is such that the following obtains:

*c*1 There is a surveyable connection between a sentence's (surface) grammatical structure and the semantic structure on the basis of which the compositional assignment of truth-conditions takes place.

At the same time, we want a semantic theory to account for the way in which sentences are put to use. For instance, we might like it to yield the following results:

*c*2 There is a significant correlation between the content of an assertion, the context in which the assertion takes place and the truth-conditions of the sentence asserted.

*c*3 There is an significant correlation between making an assertion with content $\mathcal{C}$ and having beliefs that are incompatible with $\mathcal{C}$'s negation.[28]

*c*4 There is a significant correlation between witnessing an assertion with content $\mathcal{C}$ and updating one's beliefs so that they are incompatible with $\mathcal{C}$'s negation.

According to the conception of semantics I would like to consider, the grammatical structure of sentences and the way in which sentences are put to use are the *only* two external constraints a semantic theory is answerable to. When internal considerations such as simplicity are set aside, all it takes for a semantic theory to be fully adequate is for these two constraints to be satisfied. (Whether or not *c*2–*c*4 constitute the best way of spelling-out the second constraint is a separate issue.) The crucial feature of this way of thinking about semantics is that it regards the assignment of semantic properties to sentential constituents as constrained only by its role in the assignment of semantic properties to *sentences*. There are no independently given principles about the role that the semantic values of sentential constituents should play in a semantic theory. Although the view is closely related to Frege's Context Principle,[29] here I shall refer to it as the *Minimalist Conception* of semantics.

It is important to be clear that the Minimalist Conception is compatible with semantic externalism. Provided there are suitable differences in the contents of the beliefs of Earthlings and Twin-Earthlings, one can expect there to be a difference in what 'water is wet' says about the world as used by Earthlings and what it says about the world as used by Twin-Earthlings. For analogous reasons, the Minimalist Conception is compatible with a causal theory of reference. Provided that beliefs that are suitably related to sentences involving 'Gödel' are about the individual at the end of the relevant causal chain, one can expect that 'Gödel' will refer to the individual at the end of the relevant causal chain.

4.2 How to Choose Between Rival Containment Assumptions

On a conception of semantics according to which the semantic values of various kinds of expressions must play pre-established roles, a semantic theory based on non-standard commitment assumptions might be ruled out from the start. But not on the Minimalist Conception. On the Minimalist Conception a semantic theory governed by non-standard containment assumptions might well be acceptable. What matters is whether it does

[28] With some important qualifications, one might think that the 'significant correlation' in *t*2 and *t*3 is a convention. See Lewis (1969) and Lewis (1973).

[29] "Never ask for the meaning of a word in isolation, but only in the context of a proposition". See the introduction to Frege (1884).

justice to the grammatical structure of our sentences and to the way in which sentences are put to use. In this subsection we will use the Minimalist Conception to assess two semantic theories for the langauge of arithmetic. The first is a semantic theory based on (a)–(c) from section 2.3, which embody the standard containment assumptions; the second is a semantic theory based on (a'), (b) and (c') from section 2.3, which embody the non-standard containment assumptions.

Both of our semantic theories treat the language of arithmetic as a formal language. They make no distinction between the grammatical structure of a sentence and the semantic structure that is used to produce a compositional assignment of truth-conditions for that sentence. So both of our semantic theories do an equally good job of answering to the grammatical structure of the original sentences.

But the semantic theories differ significantly when it comes to the assignment of truth-conditions. Consider 'the number of unicorns is zero'. According to the semantic theory based on non-standard containment assumptions, the condition a possible world must meet in order to be a member of the sentence's truth-set is that of containing no unicorns; according to the semantic theory based on standard containment assumptions the condition a possible world must meet in order to be a member of the sentence's truth-set is that of containing no unicorns *and* containing the number zero.

Determining which of these assignments of truth-conditions is to be preferred is no easy task. It would presumably involve a combination of serious theorizing and careful empirical research. But here is one respect in which the semantic theory based on non-standard containment assumptions seems to deliver better results. Whether or not ordinary speakers take an assertion of 'the number of unicorns is 0' to be true usually depends on whether they believe that there are unicorns, but not on whether they believe that there are numbers. So insofar as there ought to be a significant correlation between the truth-conditions of a sentence and what speakers must believe in order to take assertions of the sentence to be true (as $c2$–$c4$ suggest), there is a *prima facie* reason for accepting the view that the truth-conditions of 'the number of unicorns is 0' concern the question of whether there are any unicorns, but not the question of whether there are any numbers.

It would be a mistake to give this argument too much weight. As the study of pragmatics makes clear—and as fictionalists and pragmatic instrumentalists are keen to emphasize—the connection between the content of an assertion and the truth-conditions of the sentence asserted can be highly complex. And there are internal considerations such as simplicity to be reckoned with. At the end of the day, it might turn out that a semantic theory based on the standard containment assumptions delivers the best results, all things considered. But it is a consequence of the Minimalist Conception that it is on the basis of considerations such as the ones have been discussed in this section that the matter must be resolved. A semantic theory based on the non-standard containment assumptions cannot be ruled-out from the outset.

5 Metaphysics

The proposal I have been defending makes no metaphysical claims. It is a thesis about the commitments of sentences, not a thesis about what the world is like. In particular, no claims have been made about whether or not the world contains any numbers. A couple of metaphysical remarks are nonetheless in order.

5.1 Nominalism

That the proposal I have been defending is compatible with the view that the world contains numbers should be straightforward enough. (Just because a sentence like 'the number of unicorns is 0' carries no commitment to numbers it doesn't follow that the world doesn't contain any.) But it is not so obvious that the proposal is also compatible with Arithmetical Nominalism: the view that the world doesn't contain any numbers. For the semantic theory I set forth to buttress the claim that 'the number of unicorns is 0' carries no commitment to numbers made heavy use of number-talk. So one might be tempted to think that the semantic theory can only be taken seriously on the assumption that the world contains numbers.

The temptation should be resisted. Just as one can apply a semantic theory based on non-standard containment assumptions to the language of arithmetic and get the conclusion that a sentence like 'the number of unicorns is 0' carries no offensive commitments, one can apply a semantic theory based on non-standard containment assumptions to the language in which one's semantic theory for the language of arithmetic is couched and get the conclusion that one's semantic theory for the language of arithmetic carries no offensive commitments.[30] And there is as much reason for applying a semantics based on non-standard containment assumptions in the case of the metalanguage as there was in the case of the object language. The metatheory, after all, is just a piece of applied mathematics.

I have argued that Arithmetical Nominalism is consistent with my proposal, not that Arithmetical Nominalism is true. The latter would require more than just a thesis about the commitments of sentences.

5.2 Meinongianism

The semantic theory I set forth for the language of arithmetic makes claims such as the following:

($*$) The semantic value of '0' (relative to a possible world) is the number 0.

I noted above that by supplying a semantic theory based on the non-standard containment assumptions for the metalanguage in which ($*$) is stated, one gets the result that ($*$) carries no commitment to numbers. But now suppose Arithmetical Nominalism is true. Is one

[30] For a closely related observation, see Azzouni (1997) pp. 206-7 n. 19, and Azzouni (2004) p. 55.

thereby committed to the Meinongian conclusion that the truth of (∗) requires '0' to bear the reference relation to a non-existent object?

Certainly not. Drawing the Meinongian conclusion would be like drawing the conclusion that, in the absence of numbers, the truth of 'the number of planets is 9' would require the planets to bear the relation of being-numbered-by to a non-existent object. On the proposal I have been defending, the truth-conditions of 'the number of planets is 9' *do not concern numbers*: they concern planets, and how many of them there are. Whether or not the planets bear the relevant relation to an (existent or non-existent) number is neither here nor there when it comes to the question of whether 'the number of planets is 9' is true. Similarly, when one applies a semantic theory based on the non-standard containment to the language of (∗), the truth-conditions of (∗) *do not concern numbers*: they concern the manner in which truth-conditions are to be associated with sentences involving the expression '0'. Whether or not '0' bears a particular relation to an (existent or non-existent) number is neither here nor there when it comes to the question of whether (∗) is true.

In general, the difference between Meinongianism and the proposal I have been defending can be characterized as follows. Meinongianism is a metaphysical claim; the proposal I have been defending is a semantic thesis. Whereas a Meinongian about numbers would claim that numbers are non-existent objects, I claim that—whatever their status vis-a-vis-existence—numbers are not what arithmetical sentences are generally about any more than truth-functions are what sentences containing logical connectives are generally about. For the Meinongian, truth-conditions are standard: the truth of 'numbers are so-and-so' requires that numbers be so-and-so; what is non-standard is the metaphysics: numbers are so-and-so even though they don't exist. According to my proposal, it is the truth-conditions that are non-standard: the truth of 'numbers are so-and-so' requires nothing of numbers; the metaphysics of numbers is therefore neither here nor there.[31]

6 Other Puzzles in the Philosophy of Mathematics

A semantic theory based on the non-standard containment assumptions can be used to address other puzzles in the philosophy of mathematics. Here are a couple of examples:

1. *Indeterminacy*

   Some philosophers have been troubled by the observation that the practice of mathematicians would appear to be compatible with many different assignments of reference to the mathematical terms.[32]

   If the present proposal is along the right lines, then there is no direct route from the indeterminacy of reference to the indeterminacy of truth-conditions. We have seen,

[31] For a contemporary defence of Meinongianism, see Parsons (1980), Routley (1982) and Priest (forthcoming). For a Meinongian possible-worlds semantics, see Priest (forthcoming), § 3.3.

[32] The canonical formulation of this worry is Benacerraf (1965).

for example, that when non-standard containment assumptions are in place 'the number of unicorns is 0' receives the same truth-conditions regardless of whether the referents of numerical terms are taken to be numbers or von Neumann ordinals.

Since mathematical terms play a compositional role without playing an ontological role, the observation that the truth conditions of 'the number of unicorns is 0' are compatible with different ways of characterizing the semantic values of mathematical terms should be no more problematic than the observation that the truth-conditions of 'it is not the case that Susan runs' are compatible with different ways of characterizing the semantic value of 'it is not the case that'. Like the semantic value of a logical connective, the semantic value of a mathematical term is a piece of semantic machinery. The only constraint on its adequacy is that it enable a compositional assignment of truth conditions with the right sorts of properties.

2. *Knowledge*

Some philosophers have argued that it is difficult to understand how mathematical knowledge is possible, on the grounds that the truth-conditions of mathematical sentences concern entities in a Platonic realm, from which we are causally disconnected.[33]

From the perspective of the present proposal, the difficulty never gets off the ground. For we get the result that, just like the truth conditions of a sentence like 'it is not the case that Susan runs' do not concern truth-functions, the truth conditions of a sentence like 'the number of unicorns is 0' do not concern numbers. All we need to know about the world in order to determine whether, e.g. 'the number of unicorns is 0' is true is whether there are any unicorns. Knowing how things stand in the causally disconnected realm of abstract objects is neither here nor there.

We get the result, moreover, that knowing whether a sentence in the language of *pure* arithmetic is true requires no knowledge whatsoever about what the world is like. For, on a semantic theory based on the non-standard containment assumptions, the truth-condition associated with a sentence of pure arithmetic is either the trivial condition or the impossible condition. So all that is required for determining whether a sentence in the language of pure arithmetic is true is putting its truth-condition in a form that makes triviality or impossibility apparent. Knowing what the world is like is unnecessary for determining whether the condition is met.

Of course, the matter of putting the truth-condition of a mathematical statement in a useful form may be highly non-trivial. It may require sophisticated mathematical computations. The point is that it doesn't call for any knowledge of how things stand in the causally disconnected realm of abstract objects.

3. *Applicability*

[33]The canonical statement of this sort of argument is Benacerraf (1973).

Some philosophers have been troubled by the following question:[34] given that numbers are causally disconnected from the natural world, how could knowing the truth of sentences about *numbers* (e.g. 'the number of planets is 9') possibly be relevant to knowing the truth of sentences about the natural world (e.g. 'there are nine planets')?

From the present perspective, the answer is straightforward: a sentence like 'the number of planets is 9' is not about numbers any more than a sentence like 'it is not the case that Susan runs' is about truth-functions; it is about planets, and how many of them there are. In fact, it has precisely the same truth-conditions as 'there are nine planets'. So it is no surprise that knowing the truth of the one could be relevant to knowing the truth of the other.

7 Further Applications

It would be a mistake to suppose that the adoption of non-standard containment assumptions holds the key to every situation in which philosophers feel uncomfortable with the face-value commitments of an assertion. But I think there are some additional cases in which the proposal I have been advocating might be applied. The purpose of this section is to supply some examples.

7.1 Set theory

What goes for arithmetic goes for set theory. By giving up the standard containment assumptions it is possible to produce a semantic theory that assigns the desired truth-conditions to every sentence in the language of pure and applied set theory, and yields the conclusion that set theoretic sentences carry no problematic ontological commitments.

7.2 Second-order quantification

Advocates of the picture of second-order quantification set forth by Quine have held that a second-order sentence such as '$\exists X(Xa)$' is best understood as carrying commitment to *sets* of individuals in the relevant domain of discourse (or classes of such individuals, or properties of such individuals, or predicates applying to such individuals, of 'plural entities' of such individuals).

The Quinean picture has its detractors.[35] Some philosophers have followed Boolos (1984) in thinking that *monadic* second-order quantification might be understood in terms of plural quantification, and that plural quantification over a certain domain carries no commitment to individuals outside that domain. But many friends of plural quantification have

[34] For a discussion of applicability problems, see Steiner (1998).

[35] For the Quinean side of the debate see Quine (1986) ch. 5, Resnik (1988), Parsons (1990) and Linnebo (2003) (among others). For the rival side of the debate see Boolos (1984), Boolos (1985a), Boolos (1985b), McGee (1997), Hossack (2000), McGee (2000), Rayo and Yablo (2001), Oliver and Smiley (2001), Rayo (2002b) and Williamson (2003) (among others).

also followed Boolos in thinking that *dyadic* second-order quantification is best understood in terms of plural quantification over ordered pairs. Accordingly, they believe that a second-order sentence such as '$\exists R(Raa)$' is best understood as carrying commitment to *ordered pairs* of individuals in the relevant domain of discourse.

I would like to suggest that both the Quinean picture and the Boolosian account of dyadic second-order quantifiers rely on containment assumptions that are far from obvious. To see this, consider the following assignments of semantic value to '$\exists$':

(α) The semantic value of '$\exists$'—when binding first-order variables—is the function $f_{\exists_1}$, which assigns to each possible world w the set of non-empty sets of individuals contained in w.

(β_1) The semantic value of '$\exists$'—when binding monadic second-order variables—is the function $f_{\exists_2^1}$, which assigns to each possible world w the set of non-empty sets of sets of individuals contained in w.

(β_2) The semantic value of '$\exists$'—when binding dyadic second-order variables—is the function $f_{\exists_2^2}$, which assigns to each possible world w the set of non-empty sets of sets of pairs $\langle x_1, x_2\rangle$ such that x_1 and x_2 are contained in w.

and similarly for quantifiers binding n-place second-order variables

A semantic theory based on (α) and the (β_i) assigns the desired truth-conditions to every second-order sentence, whether or not one assumes that possible worlds contain n-tuples, or sets of any other kind. (See Appendix B for details.) So one gets the result that neither '$\exists X(Xa)$' nor '$\exists R(Raa)$' carry commitment to sets (provided, of course, that 'a' does not refer to a set).

In order for the Semantic Criterion to yield the result that '$\exists X(Xa)$' is committed to sets regardless of whether 'a' refers to a set, one would have to modify (β_1) as follows:

(β_1^*) The semantic value of '$\exists$'—when binding monadic second-order variables—is the function $f_{\exists_2^1}$, which assigns to each possible world w the set of non-empty sets of sets x of individuals such that w *contains the set* x.

And there is no obvious reason for accepting the containment assumption that has been incorporated into (β_1^*). In order for the Semantic Criterion to yield the result that '$\exists R(Raa)$' is committed to ordered pairs regardless of whether 'a' refers to an ordered pair, one would have to modify (β_2) as follows:

(β_2^*) The semantic value of '$\exists$'—when binding dyadic second-order variables—is the function $f_{\exists_2^2}$, which assigns to each possible world w the set of non-empty sets of sets of pairs $\langle x_1, x_2\rangle$ such that w *contains the ordered pair* $\langle x_1, x_2\rangle$.

And, again, there is no obvious reason for accepting the containment assumption that has been incorporated into (β_2^*).

It is worth noting that second-order quantifies are not an isolated case. On the most natural way of supplying a semantic theory for non-standard quantifiers of other kinds—e.g. quantifiers binding variables that take *sentence* positions—problematic commitments only ensue against the background of non-obvious containment assumptions.

7.3 Adverbial modification

Davidson (1967) offers a first-order paraphrase for certain sentences involving adverbial modifiers. It is suggested, for example, that a sentence such as

> Susan runs quickly

might be paraphrased as (a first-order version of) the following:

> There is an event e such that: (1) e is a running, (2) Susan is the agent of e and (3) e is quick.

This makes it tempting to supply a semantic theory for sentences like 'Susan runs quickly' that is based on the following assignments of semantic value:

- the semantic value of the name 'Susan' is the function f_{Susan}, which assigns Susan to each possible event that has Susan as an agent, and remains undefined for each possible event that does not have Susan as an agent.
- the semantic value of the one-place predicate 'runs' is a function f_{runs}, which assigns to a possible event e the value TRUE if e is a running, and the value FALSE otherwise.
- the semantic value of the adverbial modifier 'quickly' is a function f_{quickly}, which assigns to a possible event e the value TRUE if e is quick, and the value FALSE otherwise.

A possible world w can then be said to be a member of the truth-set of a sentence of the form $\ulcorner n\ F\text{s}\ A\text{ly}\urcorner$ (where $\ulcorner n \urcorner$ is a name and $\ulcorner F \urcorner$ is a one-place predicate and $\ulcorner A \urcorner$ is an adverbial modifier for $\ulcorner Fs \urcorner$) just in case the following condition is met:

($*$) w contains a possible event e such that: (1) f_F assigns to e the value TRUE, (2) f_n is an agent of e, and (3) f_A assigns to e the value TRUE.

It follows from a semantic theory based on ($*$) that a possible world can only be a member of the truth-set of 'Susan runs quickly' if it contains an event of the relevant kind. So—according to the Semantic Criterion—'Susan runs quickly' carries commitment to events of the relevant kind.

The source of the problem is that ($*$) involves a containment assumption of sorts: the assumption that it can only be true at a possible world that a given event *takes place* if the event is *contained* in the world. But the two are not the same. All it takes for Susan's running to take place is for Susan to run—whether or not the world contains the event itself as an additional piece of ontology. In order to avoid the assumption, ($*$) may be reformulated as follows:

($**$) an event e with the following characteristics takes place at w: (1) f_F assigns to e the value TRUE, (2) f_n is an agent of e, and (3) f_A assigns to e the value TRUE; *all of this whether or not* w *contains* e.

With the revised principle in place, 'Susan runs quickly' can be shown to carry no commitment to events.

I don't mean to endorse a Davidsonian account of adverbial modifiers. The point is simply that Davidsoninans can make use of non-standard containment assumptions to avoid the unpleasant conclusion that a sentence like 'Susan runs quickly' carries commitment to events.

Appendices

A A Semantic Theory for Arithmetic

For ease of reference I surround each non-standard feature of the semantic theory with with a box. Our object-language, $\mathfrak{L}$, consists of the following symbols:

1. arithmetical variables: 'n_1', 'n_2', ...;
2. non-arithmetical variables: 'x_1', 'x_2', ...;
3. arithmetical vocabulary: the individual constant '0', the one-place function-letter 's' and the two place function-letters '+' and '×';
4. non-arithmetical vocabulary: the one-place predicate letters 'F_1', 'F_2', ...;
5. the identity symbol '=';
6. the mixed two-place predicate-letter 'Num';
7. the quantifier-symbols '∃' and '∀';
8. the monadic operator '¬'
9. the dyadic operators '∧', '∨' and '⊃';
10. the auxiliaries '(' and ')'.

Terms and formulas are defined as follows:

1. any arithmetical variable is an arithmetical term;
2. '0' is an arithmetical term;
3. if $\ulcorner\tau_1\urcorner$ and $\ulcorner\tau_2\urcorner$ are arithmetical terms, then $\ulcorner s(\tau_1)\urcorner$, $\ulcorner\tau_1 + \tau_2\urcorner$ and $\ulcorner\tau_1 \times \tau_2\urcorner$ are arithmetical terms;
4. nothing else is an arithmetical term;
5. non-arithmetical variables are the only non-arithmetical terms;
6. if $\ulcorner\tau_1\urcorner$ and $\ulcorner\tau_2\urcorner$ are both arithmetical terms, then '$\ulcorner\tau_1 = \tau_2\urcorner$' is a formula;
7. if $\ulcorner\tau_1\urcorner$ and $\ulcorner\tau_2\urcorner$ are both non-arithmetical terms, then '$\ulcorner\tau_1 = \tau_2\urcorner$' is a formula;
8. if $\ulcorner\tau\urcorner$ is a non-arithmetical term, then $\ulcorner F_i(\tau)\urcorner$ is a formula;

9. if $\ulcorner\tau\urcorner$ is an arithmetical term, then $\ulcorner Num(F_i, \tau)\urcorner$ is a formula;[36]

10. if $\ulcorner v\urcorner$ is a variable and $\ulcorner\phi\urcorner$ is a formula, then $\ulcorner\exists v(\phi)\urcorner$ and $\ulcorner\forall v(\phi)\urcorner$ are formulas;

11. if $\ulcorner\phi\urcorner$ and '$\ulcorner\psi\urcorner$' are formulas, then $\ulcorner\neg\phi\urcorner$, $\ulcorner(\phi \wedge \psi)\urcorner$, $\ulcorner(\phi \vee \psi)\urcorner$ and $\ulcorner(\phi \supset \psi)\urcorner$ are formulas;

12. nothing else is a formula.

As usual, a *sentence* is a formula with no free variables.

The semantic theory I would like to propose for $\mathfrak{L}$ proceeds as follows. We begin by assigning semantic values to the basic vocabulary:

1. the semantic value of '0' is the function f_0, which assigns to each possible world w the number zero, *regardless of whether w represents the world as containing any numbers.*

2. the semantic value of 's' is the function f_s, which assigns to each possible world w the function that takes each number to its successor, *regardless of whether w represents the world as containing any numbers.*

3. the semantic value of '+' is the function f_+, which assigns to each possible world w the function that takes each pair of numbers to their sum, *regardless of whether w represents the world as containing any numbers.*

4. the semantic value of '×' is the function $f_\times$, which assigns to each possible world w the function that takes each pair of numbers to their product, *regardless of whether w represents the world as containing any numbers.*

5. the semantic value of each $\ulcorner F_i\urcorner$ is a function f_{F_i} that assigns to each possible world w a set of individuals that w represents the world as containing. (For instance, the semantic value of 'F_7' might be the function that assigns to each possible world w the set of individuals that w represents the world as containing amongst its elephants.)

6. the semantic value of 'Num' is the function f_{Num}, which assigns to each possible world w the set of ordered pairs $\langle n, x\rangle$ such that: (1) x is a set of individuals y such that w represents the world as containing y, and (2) there are precisely n objects in x; *all of this, regardless of whether w represents the world as containing any numbers.*

7. the semantic value of '=' is the function that assigns to each possible world w the set containing all and only: (a) pairs $\langle n, n\rangle$ where n is a number, and (b) pairs $\langle x, x\rangle$ where w represents the world as containing x amongst its self-identical things.

[36] I limit the second argument-place of 'Num' to *atomic* non-arithmetical predicates to keep things simple. Allowing non-atomic non-arithmetical predicates is straightforward enough. Allowing arithmetical predicates requires a trick. See Rayo (2002a).

8. $\boxed{\text{the semantic value of '$\exists$' is the function } f_\exists\text{, which assigns to each possible world } w \text{ the set containing all and only: } (a) \text{ non-empty sets of numbers, and } (b) \text{ non-empty sets of individuals } x \text{ such that } w \text{ represents the world as containing } x.}$

9. $\boxed{\text{the semantic value of '$\forall$' is the function } f_\forall\text{, which assigns to each possible world } w \text{ the set containing all and only: } (a) \text{ the set of numbers, and } (b) \text{ the set of individuals } x \text{ such that } w \text{ represents the world as containing } x.}$

10. the semantic value of '$\neg$' is the function that assigns to each possible world the function that assigns $True$ to $False$ and $False$ to $True$.

11. the semantic value of '$\wedge$' is the function that assigns to each possible world the function that assigns $True$ to the pair $\langle True, True\rangle$ and $False$ to any other pair of truth-values.

12. the semantic value of '$\vee$' is the function that assigns to each possible world the function that assigns $False$ to the pair $\langle False, False\rangle$ and $True$ to any other pair of truth-values.

13. the semantic value of '$\supset$' is the function that assigns to each possible world the function that assigns $False$ to the pair $\langle True, False\rangle$ and the $True$ to any other pair of truth-values.

Next, we introduce the standard notion of a variable assignment:

> A *variable assignment* is a function that assigns to each ordered pair $\langle \ulcorner n_i \urcorner, w\rangle$ a number and to each ordered pair $\langle \ulcorner x_i \urcorner, w\rangle$ an object that is contained in the world as represented by w.

We may then characterize a denotation function $\delta_{\langle a,w\rangle}$ (for a a variable assignment and w a possible world), in the usual way:

1. if v is a variable, then $\delta_{\langle a,w\rangle}(v) = a(\langle v, w\rangle)$;

2. $\delta_{\langle a,w\rangle}(\text{'0'}) = f_0(w) = 0$;

3. if $\ulcorner s(\tau_1)\urcorner$ and $\ulcorner s(\tau_1)\urcorner$ are arithmetical terms, then

$$\delta_{\langle a,w\rangle}(\ulcorner s(\tau_1)\urcorner) = f_s(w)(\delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner)) = (\delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner))';$$
$$\delta_{\langle a,w\rangle}(\ulcorner \tau_1 + \tau_2 \urcorner) = f_+(w)(\big\langle \delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner), \delta_{\langle a,w\rangle}(\ulcorner \tau_2 \urcorner)\big\rangle) = \delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner) + \delta_{\langle a,w\rangle}(\ulcorner \tau_2 \urcorner);$$
$$\delta_{\langle a,w\rangle}(\ulcorner \tau_1 \times \tau_2 \urcorner) = f_\times(w)(\big\langle \delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner), \delta_{\langle a,w\rangle}(\ulcorner \tau_2 \urcorner)\big\rangle) = \delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner) \times \delta_{\langle a,w\rangle}(\ulcorner \tau_2 \urcorner);$$

Finally, we characterize the notion of satisfaction relative to $\langle a, w\rangle$ (for a a variable assignment and w a possible world), in the usual way:

1. $\ulcorner F_i(x_j) \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if $\delta_{\langle a,w\rangle}(\ulcorner x_j \urcorner) \in f_{F_i}(w)$;

2. if $\ulcorner \tau_1 = \tau_2 \urcorner$ is a formula, then $\ulcorner \tau_1 = \tau_2 \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if $\left\langle \delta_{\langle a,w\rangle}(\ulcorner \tau_1 \urcorner), \delta_{\langle a,w\rangle}(\ulcorner \tau_2 \urcorner) \right\rangle \in f_{=}(w)$;

3. if $\ulcorner Num(F_i, \tau) \urcorner$ is a formula, then $\ulcorner Num(F_i, \tau) \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if $\left\langle \delta_{\langle a,w\rangle}(\ulcorner \tau \urcorner), f_{F_i}(w) \right\rangle \in f_{Num}(w)$;

4. if $\ulcorner Q \urcorner$ is a quantifier-symbol, $\ulcorner v \urcorner$ is a variable and ϕ is a formula, then $\ulcorner Qv(\phi) \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if $f_Q(w)$ contains the set of individuals x such that a variable assignment differing from a at most in assigning x to $\langle \ulcorner v \urcorner, w\rangle$ satisfies $\ulcorner \phi \urcorner$;

5. if $*$ is a monadic operator and $\ulcorner \phi \urcorner$ is a formula, then $\ulcorner * \phi \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if either (a) $f(w)(True) = True$ and $\ulcorner \phi \urcorner$ is satisfied by $\langle a, w\rangle$, or (b) $f(w)(False) = True$ and $\ulcorner \phi \urcorner$ is not satisfied by $\langle a, w\rangle$;

6. if $*$ is a dyadic operator and $\ulcorner \phi \urcorner$ and $\ulcorner \psi \urcorner$ are formulas, then $\ulcorner \phi * \psi \urcorner$ is satisfied by $\langle a, w\rangle$ if and only if either (a) $f(w)(\langle True, True\rangle) = True$ and $\ulcorner \phi \urcorner$ and $\ulcorner \psi \urcorner$ are both satisfied by $\langle a, w\rangle$, or (b) $f(w)(\langle True, False\rangle) = True$ and $\ulcorner \phi \urcorner$ is satisfied by $\langle a, w\rangle$ but $\ulcorner \psi \urcorner$ is not, or (c) $f(w)(\langle False, True\rangle) = True$ and $\ulcorner \psi \urcorner$ is satisfied by $\langle a, w\rangle$ but $\ulcorner \phi \urcorner$ is not, or (d) $f(w)(\langle False, False\rangle) = True$ and neither $\ulcorner \phi \urcorner$ nor $\ulcorner \psi \urcorner$ is satisfied by $\langle a, w\rangle$.

The truth-set of a sentence ϕ can then be defined as the set of possible worlds w such that ϕ is satisfied by $\langle a, w\rangle$ for every assignment a. And a sentence can be said to be *true* if the actual world is a member of its truth-set.

B A Semantic Theory for Second-order Languages

For ease of reference I surround each non-standard feature of the semantic theory with with a box. Our object-language, $\mathfrak{L}$, consists of the following symbols:

1. first-order variables: 'x_1', 'x_2', . . .;

2. for each $n \geq 1$ n-place second-order variables: 'X^n_1', 'X^n_2', . . .;

3. non-logical vocabulary: the one-place predicate letters 'F_1', 'F_2', . . .;

4. the identity symbol '$=$';

5. the quantifier-symbols '$\exists$' and '$\forall$';

6. the monadic operator '$\neg$'

7. the dyadic operators '$\wedge$', '$\vee$' and '$\supset$';

8. the auxiliaries '(' and ')'.

Formulas are defined as follows:

1. $\ulcorner x_i = x_j \urcorner$, $\ulcorner F_i(x_i) \urcorner$ and $\ulcorner X_i(x_j) \urcorner$ are formulas;

2. if $\ulcorner v \urcorner$ is a variable and $\ulcorner \phi \urcorner$ is a formula, then $\ulcorner \exists v(\phi) \urcorner$ and $\ulcorner \forall v(\phi) \urcorner$ are formulas;

3. if $\ulcorner \phi \urcorner$ and '$\ulcorner \psi \urcorner$' are formulas, then $\ulcorner \neg \phi \urcorner$, $\ulcorner (\phi \wedge \psi) \urcorner$, $\ulcorner (\phi \vee \psi) \urcorner$ and $\ulcorner (\phi \supset \psi) \urcorner$ are formulas;

4. nothing else is a formula.

As usual, a *sentence* is a formula with no free variables.

The semantic theory I would like to propose for $\mathfrak{L}$ proceeds as follows. We begin by assigning semantic values to the basic vocabulary:

1. the semantic value of each $\ulcorner F_i \urcorner$ is a function f_{F_i} that assigns to each possible world w a set of individuals that w represents the world as containing. (For instance, the semantic value of 'F_7' might be the function that assigns to each possible world w the set of individuals that w represents the world as containing amongst its elephants.)

2. the semantic value of '=' is the function that assigns to each possible world w the set of pairs $\langle x, x \rangle$ where w represents the world as containing x amongst its self-identical things.

3. The semantic value of '$\exists$'—when binding first-order variables—is the function $f_{\exists_1}$, which assigns to each possible world w the set of non-empty sets of individuals x such that w represents the world as containing x.

4. The semantic value of '$\forall$'—when binding first-order variables—is the function $f_{\forall_1}$, which assigns to each possible world w the set consisting of the set of all individuals x such that w represents the world as containing x.

5. The semantic value of '$\exists$'—when binding monadic second-order variables—is the function $f_{\exists^1_2}$, which assigns to each possible world w the set of non-empty sets of sets of individuals x such that w represents the world as containing x; *all of this regardless of whether w contains any sets.*

6. The semantic value of '$\forall$'—when binding monadic second-order variables—is the function $f_{\forall^1_2}$, which assigns to each possible world w the set consisting of the sets of all sets of individuals x such that w represents the world as containing x; *all of this regardless of whether w contains any sets.*

7. The semantic value of '$\exists$'—when binding n-place second-order variables ($n \geq 2$)—is the function $f_{\exists^n_2}$, which assigns to each possible world w the set of non-empty sets of sets of n-tuples $\langle x_1, \ldots, x_n \rangle$ such that w represents the world as containing $x_1, \ldots, x_n$; *all of this regardless of whether w contains any n-tuples.*

8. The semantic value of '$\forall$'—when binding n-place second-order variables ($n \geq 2$)—is the function $f_{\forall_2^n}$, which assigns to each possible world w the set consisting of the set of all sets of n-tuples $\langle x_1, \ldots, x_n \rangle$ such that w represents the world as containing $x_1, \ldots, x_n$; *all of this regardless of whether w contains any n-tuples.*

9. the semantic value of '$\neg$' is the function that assigns to each possible world the function that assigns $True$ to $False$ and $False$ to $True$.

10. the semantic value of '$\wedge$' is the function that assigns to each possible world the function that assigns $True$ to the pair $\langle True, True \rangle$ and $False$ to any other pair of truth-values.

11. the semantic value of '$\vee$' is the function that assigns to each possible world the function that assigns $False$ to the pair $\langle False, False \rangle$ and $True$ to any other pair of truth-values.

12. the semantic value of '$\supset$' is the function that assigns to each possible world the function that assigns $False$ to the pair $\langle True, False \rangle$ and the $True$ to any other pair of truth-values.

Next, we introduce the notion of a variable assignment:

A *variable assignment* is a function that assigns to each ordered pair $\langle \ulcorner x_i \urcorner, w \rangle$ an object z such that w represents the world as containing z, to each ordered pair $\langle \ulcorner X_i^1 \urcorner, w \rangle$ a set of objects z such that w represents the world as containing z, and to each ordered pair $\langle \ulcorner X_i^n \urcorner, w \rangle$ (for $n \geq 2$) a set of n-tuples $\langle z_1, \ldots, z_n \rangle$ such that w represents the world as containing $z_1, \ldots, z_n$; *all of this whether or not w represents the world as containing any sets.*

Finally, we characterize the notion of satisfaction relative to $\langle a, w \rangle$ (for a a variable assignment and w a possible world), in the usual way:

1. $\ulcorner F_i(x_j) \urcorner$ is satisfied by $\langle a, w \rangle$ if and only if $a \langle \ulcorner x_j \urcorner, w \rangle \in f_{F_i}(w)$;

2. if $\ulcorner x_1 = x_2 \urcorner$ is a formula, then $\ulcorner x_1 = x_2 \urcorner$ is satisfied by $\langle a, w \rangle$ if and only if $\langle a \langle \ulcorner x_i \urcorner, w \rangle, a \langle \ulcorner x_j \urcorner, w \rangle \rangle \in f_{=}(w)$;

3. $X_i(x_j)$ is satisfied by $\langle a, w \rangle$ if and only if $a \langle \ulcorner x_j \urcorner, w \rangle \in a \langle \ulcorner X_i \urcorner, w \rangle$;

4. if $\ulcorner Q \urcorner$ is a quantifier-symbol, $\ulcorner v \urcorner$ is a variable and ϕ is a formula, then $\ulcorner Qv(\phi) \urcorner$ is satisfied by $\langle a, w \rangle$ if and only if $f_Q(w)$ contains the set of x such that a variable assignment differing from a at most in assigning x to $\langle \ulcorner v \urcorner, w \rangle$ satisfies $\ulcorner \phi \urcorner$;

5. if $*$ is a monadic operator and $\ulcorner \phi \urcorner$ is a formula, then $\ulcorner * \phi \urcorner$ is satisfied by $\langle a, w \rangle$ if and only if either (a) $f(w)(True) = True$ and $\ulcorner \phi \urcorner$ is satisfied by $\langle a, w \rangle$, or (b) $f(w)(False) = True$ and $\ulcorner \phi \urcorner$ is not satisfied by $\langle a, w \rangle$;

6. if $*$ is a dyadic operator and $\ulcorner\phi\urcorner$ and $\ulcorner\psi\urcorner$ are formulas, then $\ulcorner\phi * \psi\urcorner$ is satisfied by $\langle a, w\rangle$ if and only if either (a) $f(w)(\langle True, True\rangle) = True$ and $\ulcorner\phi\urcorner$ and $\ulcorner\psi\urcorner$ are both satisfied by $\langle a, w\rangle$, or (b) $f(w)(\langle True, False\rangle) = True$ and $\ulcorner\phi\urcorner$ is satisfied by $\langle a, w\rangle$ but $\ulcorner\psi\urcorner$ is not, or (c) $f(w)(\langle False, True\rangle) = True$ and $\ulcorner\psi\urcorner$ is satisfied by $\langle a, w\rangle$ but $\ulcorner\phi\urcorner$ is not, or (d) $f(w)(\langle False, False\rangle) = True$ and neither $\ulcorner\phi\urcorner$ nor $\ulcorner\psi\urcorner$ is satisfied by $\langle a, w\rangle$.

The truth-set of a sentence ϕ can then be defined as the set of possible worlds w such that ϕ is satisfied by $\langle a, w\rangle$ for every assignment a. And a sentence can be said to be *true* if the actual world is a member of its truth-set.

References

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

Azzouni, J. (1997) "Applied Mathematics, Existential Commitment and the Quine-Putnam Indispensability Thesis," *Philosophia Mathematica* 5, 193–209.

Azzouni, J. (1998) "On "'On What There Is'," *Philosophical Quarterly* 79, 1–18.

Azzouni, J. (2004) *Deflating Existential Consequence: A Case for Nominalism*, OUP, New York.

Benacerraf, P. (1965) "What Numbers Could not Be," *The Philosophical Review* 74, 47–73. Reprinted in Benacerraf and Putnam (1983).

Benacerraf, P. (1973) "Mathematical Truth," *Journal of Philosophy* 70, 661–79. Reprinted in Benacerraf and Putnam (1983).

Benacerraf, P., and H. Putnam, eds. (1983) *Philosophy of Mathematics*, Cambridge University Press, Cambridge, second edition.

Boghossian, P., and C. Peacocke, eds. (2000) *New Essays on the A Priori*, Clarendon Press.

Boolos, G. (1984) "To Be is to Be a Value of a Variable (or to be Some Values of Some Variables)," *The Journal of Philosophy* 81, 430–49. Reprinted in Boolos (1998).

Boolos, G. (1985a) "Nominalist Platonism," *Philosophical Review* 94, 327–44. Reprinted in Boolos (1998).

Boolos, G. (1985b) "Reading the *Begriffsschrift*," *Mind* 94, 331–334. Reprinted in Boolos (1998).

Boolos, G. (1998) *Logic, Logic and Logic*, Harvard, Cambridge, Massachusetts.

Cartwright, R. (1954) "Ontology and the theory of meaning," *Philosophy of Science* 21, 316–325. Reprinted in Cartwright (1987).

Cartwright, R. (1987) *Philosophical Essays*, MIT Press, Cambridge, MA.

Davidson, D. (1967) "The Logical Form of Action Sentences." Reprinted in Donald Davidson, *Essays on Actions and Events*.

Eklund, M. (forthcoming) "Fiction, Indifference and Ontology," *Philosophy and Phenomenological Research*.

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

Hale, B. (1997) "Grundlagen §64," *Proceedings of the Aristotelian Society* 97, 243–61. Reprinted in Hale and Wright (2001), 91-116.

Hale, B. (1999) "Arithmetic Reflection without Intuition," *Aristotelian Society, Supplementary Volume* 73, 75–98.

Hale, B., and C. Wright (2000) "Implicit Definition and the A Priori." In Boghossian and Peacocke (2000), 286-319. Reprinted in Hale and Wright (2001), 117-150.

Hale, B., and C. Wright (2001) *The Reason's Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics*, Clarendon Press.

Heck, R., ed. (1997) *Language, Thought and Logic*, Clarendon Press, Oxford.

Hellman, G. (1989) *Mathematics Without Numbers*, Clarendon Press, Oxford.

Hellman, G. (1990) "Modal-Structural Mathematics" 307–30. In Irvine (1990).

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

Hofweber, T. (forthcoming) "A puzzle about ontology," *Noûs*.

Hofweber, T. (typescript) "Number determiners, numbers and arithmetic."

Hossack, K. (2000) "Plurals and Complexes," *The British Journal for the Philosophy of Science* 51:3, 411–443.

Irvine, A., ed. (1990) *Physicalism in Mathematics*, Dordrecht.

Lewis, D. (1969) *Convention: A Philosophical Study*, Harvard University Press, Cambridge, MA.

Lewis, D. (1970) "General Semantics," *Synthese* 22, 18–67. Reprinted in Lewis (1983).

Lewis, D. (1973) "Languages and Language," *Versus* 4, 2–21. Appears in Italian translation as 'Lingue e lingua'. Reprinted in English in Lewis (1983).

Lewis, D. (1983) *Philosophical Papers, Volume I*, Oxford.

Lewis, D. (1986) *On the Plurality of Worlds*, Blackwell, Oxford and New York.

Linnebo, Ø. (2003) "Plural Quantification Exposed," *Noûs* 37, 71–92.

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

McGee, V. (2000) "'Everything'." In Sher and Tieszen (2000).

Oliver, A., and T. Smiley (2001) "Strategies for a Logic of Plurals," *Philosophical Quarterly* 51, 289–306.

Parsons, C. (1990) "The Structuralist View of Mathematical Objects," *Synthese* 84, 303–346.

Parsons, T. (1980) *Nonexistent Objects*, Yale University Press, New Haven and London.

Priest, G. (forthcoming) *Towards Non-Being: the Logic and Metaphysics of Intentionality.*

Quine, W. V. (1948) "On what there is." Reprinted in Willard V. Quine, *From a Logical Point of View.*

Quine, W. V. (1986) *Philosophy of Logic, Second Edition*, Harvard, Cambridge, Massachusetts.

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

Rayo, A. (2002b) "Word and Objects," *Noûs* 36, 436–464.

Rayo, A. (2005) "Logicism Reconsidered." In Shapiro (2005).

Rayo, A., and S. Yablo (2001) "Nominalism Through De-Nominalization," *Noûs* 35:1, 74–92.

Resnik, M. (1988) "Second-Order Logic Still Wild," *Journal of Philosophy* 85:2, 75–87.

Routley, R. (1982) "On What There Is Not," *Philosophy and Phenomenological Research* 43, 151–77.

Shapiro, S., ed. (2005) *The Oxford Handbook for Logic and the Philosophy of Mathematics*, Clarendon Press, Oxford.

Sher, G., and R. Tieszen, eds. (2000) *Between Logic and Intuition*, Cambridge University Press, New York and Cambridge.

Steiner, M. (1998) *The Applicability of Mathematics as a Philosophical Problem*, Harvard University Press, Cambridge, MA.

Van Fraassen, B. C. (1980) *The Scientific Image*, Clarendon Press, Oxford.

Williamson, T. (2003) "Everything" 415–465. In Hawthorne, J. and D. Zimmerman, eds. *Philosophical Perspectives 17: Language and Philosophical Linguistics*, Blackwell, Oxford.

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

Wright, C. (1997) "The Philosophical Significance of Frege's Theorem." In Heck (1997), 201-244.

Yablo, S. (2001) "Go Figure: A Path Through Fictionalism," *Midwest Studies in Philosophy* XXV, 72–102.

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