

MIT Open Access Articles

Reply to Critics

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Rayo, Agustín. "Reply to Critics." *Inquiry*, vol. 57, no. 4, July 2014, pp. 498–534.

As Published: <http://dx.doi.org/10.1080/0020174X.2014.905036>

Publisher: Informa UK Limited

Persistent URL: <http://hdl.handle.net/1721.1/116283>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Reply to Critics

Agustín Rayo

April 9, 2014

This project has been a joy to work on. Cameron, Eklund, Hofweber, Linnebo, Russell and Sider wrote six splendid essays, filled with interesting ideas and thoughtful criticism. I cannot thank them enough.

The essays overlap in interesting ways, so I decided to organize my replies thematically, rather than writing a separate reply for each critic. There are three main headings: ‘World and Language’, ‘World’ and ‘Language’.

1 World and Language

1.1 Compositionism (Eklund)

Much of the book is premised on *compositionism*: a doctrine about the relationship between our language and the world it represents.

The basic idea behind compositionism is straightforward. Let L be an uninterpreted first-order language, and consider an assignment of truth-conditions to sentences in L . Our assignment must respect *logical entailments*,¹ but we are otherwise allowed to pick any assignment we like. According to compositionism, each of the following claims is true:

¹We require, in particular, that if ψ is a logical consequence of ϕ , then the truth conditions assigned to ϕ must demand at least as much of the world as the truth-conditions assigned to ψ . (See *The Construction of Logical Space*, Section 1.3 for a more detailed constraint.) I assume that we are operating with a negative free logic, so as to avoid trivializing the condition on reference that is mentioned below.

- C1 Should one stipulate that the sentences of L are to be interpreted on the basis of the chosen assignment of truth conditions, there is nothing to stop the stipulation from rendering the vocabulary in L meaningful, and doing so in such a way that each sentence of L comes to have its assigned truth-conditions.
- C2 When the sentences of L are so interpreted, all it takes for a singular term t of L to refer to an object in the world—all it takes for t to be non-empty—is for the world to satisfy the truth-conditions that were assigned to $\ulcorner \exists x(x = t) \urcorner$ (or some inferential analogue).

What makes compositionism a substantial view is that the constraints it imposes on a possible assignment of truth-conditions are so very weak. Suppose, for example, that L contains singular terms of the form ‘direction $^*(t)$ ’. It follows from C1 that—as long as we are careful to respect logical entailments—there is nothing to stop us from stipulating that ‘direction $^*(a) = \text{direction}^*(b)$ ’ is to be true just in case line a is parallel to line b . Moreover, on the most natural way of extending this assignment to other sentences of the language,² the sentence ‘ $\exists x(x = \text{direction}^*(a))$ ’ gets assigned the truth-condition that there be a *line* identical to a . And since this condition will be satisfied by the world as long as line a exists, it follows from C2 that the term ‘direction $^*(a)$ ’ will refer to an

²One can stipulate that a sentence ϕ of L is to have the same truth-conditions as its ‘nominalization’ $[\phi]^N$, where nominalizations are defined as follows:

- $[\ulcorner \text{direction}^*(a) = \text{direction}^*(b) \urcorner]^N = \ulcorner a \text{ is parallel to } b \urcorner$.
- $[\ulcorner x_i = \text{direction}^*(a) \urcorner]^N = \ulcorner z_i \text{ is parallel to } a \urcorner$.
- $[\ulcorner x_i = x_j \urcorner]^N = \ulcorner z_i \text{ is parallel to } z_j \urcorner$.
- $[\ulcorner \exists x_i(\phi) \urcorner]^N = \ulcorner \text{there is a line } z_i \text{ such that } ([\phi]^N) \urcorner$.
- $[\ulcorner \phi \wedge \psi \urcorner]^N = \text{the conjunction of } [\phi]^N \text{ and } [\psi]^N$.
- $[\ulcorner \neg\phi \urcorner]^N = \text{the negation of } [\phi]^N$.

This assignment of truth-conditions is not defined for every sentence of the language. And, in particular, it is not defined for ‘mixed’ identity statements, such as ‘ $a = \text{direction}^*(a)$ ’. Although this could easily be remedied by treating mixed identities as false, an important feature of compositionism is that it does not presuppose that the vocabulary in L can only be meaningful if every sentence in L has well-defined truth-conditions. (See *The Construction of Logical Space*, Sections 1.3 and 3.2.)

object in the world as long as a exists—or, as a speaker of L might put it: if a exists, then so does its direction*.

Here then is the picture of the relationship between language and the world that compositionality delivers. To render a language meaningful is to decide which ways for the world to be are to be associated with which sentences. The world determines which sentences are true, by determining which ways for the world to be are actualized. But there is no need for the world to be, in some sense, ‘responsive’ to a sentence’s compositional structure in order to make it true. Compositional structure matters to truth only insofar as it determines logical entailments between sentences, and thereby limits the ways in which one’s decision to associate some ways for the world to be with some sentences can coexist with one’s decision to associate other ways for the world to be with other sentences. The issue of whether a singular term refers is determined by the compositional structures that one chose to associate with ways for the world to be that turn out to be actualized. (If one wanted to describe this picture with a slogan, one might say that compositionality is the view that our language only makes contact with the world at the level of *sentences*.³)

In his contribution to the volume, Matti Eklund complains that the full compositionalist picture ‘doesn’t actually seem to be stated in the theses used to officially characterize compositionality’ (p. xx). I’m embarrassed to report that he is right. Although the official statement of compositionality in the book entails a version of [C2] above, I left out [C1]. This obscures the ensuing discussion, which tacitly presupposes that [C1] is in place.

[C1] tells us that the vocabulary of an uninterpreted language can be rendered meaningful in such a way that its sentences get assigned arbitrarily chosen truth-conditions—arbitrarily chosen, except for the requirement that logical entailments be respected. As

³I am not, of course, the first to articulate a version of this view. See Frege, *Die Grundlagen der Arithmetik*; Dummett, *Frege: Philosophy of Language*; Wright, *Frege’s Conception of Numbers*; Rosen, ‘The Refutation of Nominalism (?)’.

I point out in the book, there is a simple technical result that shows that if a theory is internally coherent, then it is possible to find an assignment of truth-conditions to sentences in the language that respects logical entailments and assigns trivially-satisfiable truth-conditions to every sentence of the theory (§8.2). So it follows from [C1] that as long as the theory is initially regarded as uninterpreted, its vocabulary can be rendered meaningful in such a way that each sentence in the theory gets assigned trivially satisfiable truth-conditions.

Eklund objects that the version of this argument I present in the book is unsound: ‘Everyone agrees that consistent (coherent, conservative) theories have models; the question concerns whether a pure mathematical theory’s satisfying the requirement suffices for it to be true.’ (pp xx). I think Eklund is quite right to issue this complaint. The argument is only sound in the presence of [C1], and [C1] was never made explicit in the text.

1.2 Numbers (Hofweber, Sider)

Trivialism is the view that the truths of pure mathematics have *trivial* truth-conditions: truth-conditions whose satisfaction requires *nothing* of the world. Since it is a truth of pure mathematics that numbers exist, trivialism entails that the requirement that the world contain numbers would be satisfied trivially, regardless of the way the world turned out to be.

In the book I motivate Trivialism by making a case for the following claim:

[DINOSAURS]

For the number of the dinosaurs to be zero *just is* for there to be no dinosaurs.

and, more generally,

[NUMBERS]

For the number of the F s to be n *just is* for it to be the case that $\exists!_n x(Fx)$.

[Numbers] helps motivate Trivialism because it not only entails that numbers exist. It entails that they exist *on pain of contradiction*. (*Proof*: Suppose there are no numbers; by [Numbers], the number of numbers is zero; so zero exists; so numbers exist; contradiction!) But if a world with no numbers would be inconsistent—and therefore *absurd*—there is no particular way our world would need to be in order to make it the case that numbers exist. So the demand that the world contain numbers is a demand that will be satisfied trivially.

According to the account of ‘just is’-statements that I develop in chapters 1 and 2 of the book, the way to decide whether to accept a ‘just is’-statement is to perform a cost-benefit analysis. The *benefit* of accepting [Numbers] is that one would be left with no theoretical gap between there being n F s and the number of the F s being n , so one would no longer need to seek a justification for transitioning from the one to the other; the *cost* is that one has fewer theoretical resources to work with, since one loses the distinction between worlds with n F s and worlds in which the number of the F s is n . What I suggest in the book (§§1.3, 3.0) is that the the benefits of accepting [Numbers] far outweigh the costs, and that [Numbers] should therefore be accepted.

As Ted Sider points out, however, this is potentially misleading: there is an important sense in which my ‘account of mathematical truth does not, in fact, rely on a cost-benefit argument for the truth of [Numbers].’ (p xx) The reason is that in the book I also set forth a *semantic theory* for the language of arithmetic, and this theory guarantees that [Numbers] is correct.

My primary goal in advancing the semantic theory—the Trivialist Semantics, as I call it—was to give proper articulation of Mathematical Trivialism. I wanted a precise statement of the truth-conditions that a trivialist thinks should be associated with each sentence in the language of arithmetic (and, more generally, with each sentence in the language of set theory). As Thomas Hofweber points out (p. xx), this is easily done in the case of *pure* mathematics, since the trivialist can give a precise statement of the

proposal by claiming that true sentences have trivial truth-conditions and false sentences have impossible truth-conditions. But the task is considerably more difficult when it comes *applied* mathematics.⁴

Just because one is able to *articulate* a semantic theory for the language of arithmetic, it doesn't follow that the semantic theory is *correct*. One needs some sort of argument for thinking that the assignment of truth-conditions that is delivered by the theory corresponds to the actual truth-conditions of sentences in one's object-language. Since the Trivialist Semantics is a generalization of the idea that what it takes for 'the number of the F s = n ' to be true is for there to be exactly n F s, the cost-benefit analysis of [Numbers] is one such argument.

The reason I nonetheless think that Sider is right when he suggests that the appeal to [Numbers] is inessential is that, as I argue in chapters 3 and 4, the Trivialist Semantics can be used as the basis of an attractive philosophy of mathematics, and it seems to me that this supplies powerful grounds for taking it to be correct, regardless of whether one has an independent argument for [Numbers].

1.3 Outscoping (Sider)

The Trivialist Semantics entails a semantic clause for each sentence in the language of arithmetic, a clause that specifies a suitable 'trivialist' truth-condition. What do such clauses look like? One might have expected to see something along the following lines:

- (1) 'the number of the dinosaurs = 0' is true at $w \leftrightarrow$ [there are no dinosaurs] $_w$

When $[\dots]_w$ is read 'at w , it is the case that ...', (1) tells us that what would be required of world w for 'the number of the dinosaurs = 0' to count as true at w is for w to contain

⁴In the case of applied set-theory, the only way of doing so that I know of is by way of the compositional semantics that I develop in *The Construction of Logical Space*, Section 3.4.1. In the case of applied arithmetic, it can be done either by way of the compositional semantics described in Section 3.3 or by way of the non-compositional procedure described in Section 8.2.3. It seems to me, however, that the compositional semantics is vastly more illuminating than its non-compositional counterpart.

no dinosaurs. We can therefore think of (1) as assigning to ‘the number of the dinosaurs = 0’ the truth-condition that there be no dinosaurs.

Sadly, there are powerful reasons for thinking that no effectively specifiable semantic theory could entail a trivialist clause in the style of (1) for every sentence in the language. The problem is that such a theory would require the existence of a trivialist *paraphrase-function*: a function that assigns to each sentence ϕ of the object-language a metalinguistic sentence that uncontroversially expresses the truth-conditions that a trivialist would wish to associate with ϕ . And there is a formal result that suggests that there can be no such paraphrase-function: arithmetical truth is just too complex.⁵

Fortunately, the trivialist can get around the problem by availing herself of a device I call ‘outscoping’. Instead of specifying truth-conditions by appeal to a trivialist paraphrase-function, as in (1), the trivialist can use clauses of the following kind:

- (2) ‘the number of the dinosaurs = 0’ is true at $w \leftrightarrow$ the number of z such that
 $[x \text{ is a dinosaurs}]_w = 0$.

The crucial thing to notice about (2) is that even though it contains mathematical vocabulary in its right-hand side, the mathematical vocabulary has all been placed outside the scope of $[\dots]_w$: it has all been ‘outscoped’. The result is that (2) ends up delivering the same trivialist specification of truth conditions as (1).⁶ They both assign to ‘the number of dinosaurs = 0’ the truth-condition that there be no dinosaurs.

Moreover, since (2) does not rely on a trivialist paraphrase-function, it does not fall prey to the formal result I mentioned earlier. In fact, it is easy to characterize a compositional semantics that entails suitably outscoped semantic clauses for every sentence of the language of pure and applied arithmetic—and, indeed, pure and applied set-theory—and does so in a way that yields a trivialist specification of truth-conditions.⁷

⁵*The Construction of Logical Space*, Ch. 7.

⁶To keep things simple, I assume that the range of our metalinguistic variables in (2) includes merely possible objects. One can, however, avoid this assumption by appeal to the machinery in *The Construction of Logical Space*, Ch. 6.

⁷*The Construction of Logical Space*, Section 3.3.

That is the Trivialist Semantics.

As Sider points out, outscoping suffers from an important shortcoming: it is limited in its ability to address ‘discourse-threat’ (p. xx). I will illustrate the point by considering three different kinds of discourse-threat.

Consider, first, a mathematical *nihilist*: someone who thinks that mathematical vocabulary is meaningless. An outscoped semantic clause such as (2) would do nothing to convince the nihilist that ‘the number of the dinosaurs = 0’ is, in fact, meaningful. For she would take the clause itself to be meaningless, on the grounds that its right-hand side contains arithmetical vocabulary. Notice, in contrast, that the nihilist *would* be moved by a paraphrase-function that mapped each mathematical statement to a ‘nominalistic’ counterpart.

Next, consider a mathematical *indeterminist*: someone takes mathematical vocabulary to be meaningful and is happy to engage in ordinary mathematical practice but who thinks that certain mathematical sentences are neither true nor false: the Continuum Hypothesis, as it might be. As before, our outscoped semantic clauses would do nothing to convince the indeterminist of the determinacy of a sentence that she previously regarded as indeterminate. The easiest way to see this is to note that when ϕ is a sentence of *pure* mathematics, the trivialist semantics will deliver a semantic clause in which *all* vocabulary has been outscoped:

(3) ϕ is true at $w \leftrightarrow p$ and $[\top]_w$.

(where ‘ \top ’ is replaced by a tautology and ‘ p ’ is replaced by a metalinguistic version of ϕ)

When ϕ is the Continuum Hypothesis, the resulting version of (3) gives us two pieces of information. It tells us that *if* there is no set whose cardinality is strictly between that of the integers and that of the reals, *then* the Continuum Hypothesis has trivially satisfiable truth conditions (because it will be true at w , regardless of how w happens to

be). And it tells us that *if there is* a set whose cardinality is strictly between that of the integers and that of the reals, *then* the Continuum Hypothesis has trivially unsatisfiable truth conditions (because it will be false at w , regardless of how w happens to be). This means that (3) is not entirely inert: it entails that the Continuum Hypothesis has truth-conditions that are either trivially satisfiable or trivially unsatisfiable. But this is no help to our indeterminist. For if she thinks that there is no fact of the matter as to whether there is a set whose cardinality is strictly between that of the integers and that of the reals, (3) will give her no reason to change her mind.

The third kind of discourse-threat I would like to consider is the one that Sider is most interested in. Consider a skeptic who wonders ‘whether reality contains enough to tie down discourse about causation, morality, or mathematics’ (p. xx). She believes, in particular, that reality must somehow ‘underwrite’ mathematical truth, and worries that ‘[none] of the facts that have been proposed as underlying mathematical truth are wholly comfortable to accept: facts about ‘Platonic entities’, infinitely many concrete objects, primitive modal facts, primitive facts about fictional truth, primitive higher-order facts, and so on. And it’s natural to worry that in the absence of all such underwriting facts, there would be massive indeterminacy in mathematical statements.’ (p. xx)

Sider suggests that in a situation like this ‘using outscoping won’t give us what we want: a guarantee that reality adequately ties down the discourse’ (p. xx). I agree that outscoping does not deliver a new class of facts with which to underwrite mathematical discourse: an alternative to platonic entities, or infinitely many concrete objects, or any of the other potentially underwriting facts on Sider’s list. It seems to me, however, that the kind of discourse-threat we are concerned with here is unlike the kinds of discourse-threat we discussed earlier in an important respect. For in this case outscoping *does* deliver a result that is capable of engaging the skeptic—not by giving the skeptic what she wants, but by showing that the her demands are unreasonable.

The reason is straightforward. If the Trivialist Semantics I propose is along the right lines, then pure mathematics is simply not in need of underwriting. If the truths of pure mathematics have trivially satisfiable truth-conditions, then there are no constraints that reality would have to satisfy in order to ‘tie down’ the relevant discourse. The truths of pure mathematics are like the truths of pure logic: they are not the sorts of things that need to be tied down. (For further discussion of this issue, see footnote 20.)

Sider has a nice analogy, which can be used to shed further light on the question of what outscoping can and cannot do:

Since Rayo is inclined to reject talk of metaphysical structure (because of the apparently unanswerable questions such talk raises), he doesn’t indulge in such talk. But some of us do; and Rayo presumably suspects that our discourse about metaphysical structure is massively indeterminate, or projective of our emotions, or is in some other way a failure. However exactly he conceptualizes this failure, he will want to deny that mathematical discourse fails in the same way. He therefore has as much reason as anyone to want an answer to discourse-threat in mathematics. (p. xx)

Sider is right about my suspicions: I worry that our best theorizing about metaphysical structure might fail to determine well-defined truth-conditions for sentences containing talk of metaphysical structure. More importantly, I think Sider is right about the challenge I face. Since I am happy to indulge in talk of mathematical objects, I am committed to thinking that mathematical sentences, unlike metaphysical-structure sentences, do have well-defined truth-conditions (or at least that enough mathematical sentences have truth-conditions that are well-enough defined to make mathematical practice worthwhile).

But it seems to me that the Trivialist Semantics supports just that conclusion. For it can be used to identify a precise assignment of truth-conditions for every theorem of pure

or applied mathematics (and for the negation of every such theorem). It can be used, for example, to identify perfectly determinate truth-conditions to ‘ $2+2=4$ ’; specifically: the trivially satisfiable truth-conditions.

Here I am presupposing that one is prepared to *use* mathematics in the metatheory. In order to establish that ‘ $2+2=4$ ’ has trivial truth-conditions, for example, one has to prove that $2+2=4$ in the metatheory. But—and this is the crucial point—nobody is in any serious doubt about whether $2+2=4$. So nobody should be in any serious doubt about the truth-conditions that the trivialist semantics associates with ‘ $2+2=4$ ’. What the trivialist semantics does, in other words, is transform our ability to engage in mathematical practice into an ability to identify uncontroversially determinate truth conditions for mathematical statements.

In contrast, I know of no procedure that I could use to satisfy myself that sentences about metaphysical structure have well-defined truth-conditions. The difference between the two cases is partly to do to the dialectical situation: whereas no philosopher would seriously see herself as unable to engage in mathematical practice, some of us have genuine doubts about our ability to engage in the practice of metaphysical-structure talk. It is worth keeping in mind, however, that there is also a difference that has nothing to do with the dialectical situation: as far as I know, there is no analogue of the trivialist semantics for the discourse concerning metaphysical structure. There is no procedure that would allow us to transform an ability to engage in the practice into an ability to identify fully determinate truth-conditions for the relevant statements.

1.4 Explanatory Dangers (Linnebo)

Øystein Linnebo has a complaint that is in some ways related to Sider’s idea that arithmetical truth needs to be ‘tied down’ by reality.

He asks us to consider a semantics for the language of arithmetic in which truth-conditions are determined non-compositionally. Instead, they are ‘determined by syntac-

tically characterized calculations that are performed on [arithmetical] expressions' (p xx). In the case of an arithmetical truth such as '2 + 1 = 3', for example, the relevant calculations might produce the sequence: $2 + 1 = 3$, $(2 + 0)' = 2'$, $2' = 1'$, $1' = 0''$, $0'' = 0''$, \top . One can then take the truth-conditions of the original sentence to be the truth-conditions of the last member of the sequence, which in this case is a tautology.

Linnebo then suggests the following diagnosis:

On the envisaged semantics, the assignment of reference to arithmetical singular terms drops out as irrelevant to the truth of arithmetical sentences. If such terms can be said to refer at all (which I doubt), then their referents will play no role in the explanation of the truth of arithmetical sentences in which the terms occur. Numbers would thus at best be left as explanatory danglers, utterly irrelevant to the truth of sentences that appear to be about numbers. A view that so completely undermines the analogy between mathematical and non-mathematical truth does not deserve to be classified as platonist.'

A compositionalist would disagree with Linnebo about whether the relevant terms should be taken to refer. For the compositionalist's condition [C1] (from Section 1.1) guarantees that Linnebo's procedure can, in fact, be used to specify truth-conditions to arithmetical sentences. And, on any reasonable way of spelling out the details,⁸ sentences of the form ' $\exists x(x = k)$ ' will be assigned trivial truth-conditions. So it follows from condition [C2] that numerals are non-empty: they genuinely refer to objects in the world.

What should our compositionalist make of Linnebo's claim that numbers would then be 'explanatory danglers, utterly irrelevant to the truth of sentences that appear to be about numbers'?

An analogy might be helpful. Consider the specification of truth-conditions to direction*-statements that I described in Section 1.1. Nowhere in this specification did I make

⁸See, for instance, *The Construction of Logical Space*, Section 7.4.

use of an assignment of reference to direction*-terms. What I did instead was supply a *syntactic* recipe for transforming each direction*-sentence into a line-sentence—a recipe that is in some ways analogous to Linnebo’s syntactic recipe for transforming arithmetical truths into tautologies—and used the end point of the transformation to specify truth-conditions for the original sentence. I then noted that the compositionalist would claim that, as long as line a exists, ‘the direction* of a ’ will be a genuinely referential singular term (and therefore that if a exists, so will its direction*).

A critic might take her cue from Linnebo, and claim that the resulting Platonism about directions* isn’t worthy of the name, on the grounds that ‘the assignment of reference to direction*-terms drops out as irrelevant to the truth of the direction*-sentences’. But the compositionalist would see this as a mistake. For if language only makes contact with the world at the level of *sentences* (Section 1.1), the adequacy of a semantic theory cannot possibly depend on whether it proceeds by carrying out a syntactic transformation rather than by assigning referents to singular terms. The details of one’s specification of truth-conditions won’t matter, as long as they deliver the right assignment of truth-conditions to sentences.

Notice, moreover, that a speaker of the direction*-language would claim that directions* are *not* irrelevant to the truth of direction*-statements. She would claim, for example, that what is required of the world for ‘ $\exists x (x = \text{the direction* of } a)$ ’ to be true is for the direction* of a to exist, and go on to conclude that how things stand with directions* is *exactly* what matters to the truth of direction*-statements. (She would, of course, *also* claim that what is required of the world for ‘ $\exists x (x = \text{the direction* of } a)$ ’ to be true is for line a to exist. But this is as it should be, since our speaker believes that all it takes for the direction* of a to exist is for line a to exist.)

The lesson of all this is that, from the compositionalist’s point of view, it would be a mistake to conclude anything about whether a singular term refers—or about whether its referent would be relevant to the truth of sentences in which it occurs—from the

manner in which truth-conditions for the relevant sentences were specified. In particular, it would be a mistake to conclude that the referents of arithmetical terms are ‘explanatory danglers’ from the fact that one uses a particular syntactic procedure to specify truth-conditions for arithmetical sentences.

Since we have been seeing things from the perspective of a compositionalist, our discussion is best understood as a defensive maneuver: it is aimed at showing that if you start out as a compositionalist, then you’ll be able to resist Linnebo’s argument. But, of course, one might claim that Linnebo’s critique is best read not as an effort to get the compositionalist to change her ways, but as a warning to the uncommitted: ‘yield to compositionism’—the warning would go—‘and you’ll end up with a version of Platonism unworthy of the name.’ It is not clear to me, however, that we have a neutral standpoint from which to assess such a warning. For it is hard to separate the question of how best to think of Platonism from the question of whether compositionism is correct. If the warning is to have force, it seems to me that it needs to be accompanied by an independent argument against compositionism.⁹

1.5 Identity and Absolute Generality (Eklund, Russell)

Our discussion of compositionism in section 1.1 was focused on the special case of uninterpreted languages. But one could defend a version of the same idea for cases in which one starts with an interpreted language and uses a linguistic stipulation to render meaningful a certain *extension* of the language.

As before, the compositionalist will claim that as long as the proposed specification of truth-conditions respects logical entailments, there is nothing to stop it from being successful. And, as before, she will claim that all it takes for a term t of the extended language to refer is for the world to satisfy the truth-conditions that have been associated

⁹For more on Linnebo’s views on the relationship between our language and the world it represents, see Linnebo ‘Reference by Abstraction’.

with $\lceil \exists x(x = t) \rceil$ (or some inferential analogue). In this case, however, she will insist that the our assignment of truth-conditions be *conservative*: that it respect the truth-conditions that had been previously associated with sentences of the original language.¹⁰

Suppose, for example, that one starts with a language containing no mathematical vocabulary, and that one sets out to enrich it with arithmetical terms. One way to do so would be to think of the extended language as a *two-sorted* language, in which arithmetical predicates are only allowed to take arithmetical terms as arguments, and in which non-arithmetical predicates are only allowed to take non-arithmetical terms as arguments. One can then stipulate that a sentence built from purely arithmetical vocabulary is to have trivially satisfiable truth-conditions if it follows from the Dedekind-Peano Axioms, and trivially unsatisfiable truth-conditions if its negation follows from the axioms. When further details spelled out in the right kind of way,¹¹ this delivers an entailment-preserving specification of truth-conditions which is conservative over sentences of the original language. So from the compositionalist's point of view, one will have succeeded in extending the original language to contain the language of (pure) arithmetic.

And, of course, the procedure is quite general. For any language L and any theory of pure mathematics T , one can extend L by adding the vocabulary of T and treating the terms of T and L as corresponding to different sorts. As before, the compositionalist will think that one can go on to render the extended language meaningful in a way that delivers trivially satisfiable truth-conditions to the axioms of the theory while respecting the truth-conditions of sentences in the original language.¹²

In describing this sort of idea in the book, I usually worked on the assumption that

¹⁰For further discussion, see *The Construction of Logical Space*, Section 8.2.3.

¹¹Here are the details. Let ϕ be an arbitrary sentence of the extended language, let w be a world, and let O_w be the set of sentences containing no new vocabulary that are true at w , let ν be the sentence $\forall x \forall n(x \neq n)$ (where x is a variable of the non-arithmetical sort and n is a variable of the arithmetical sort). Then perform the following stipulations: (1) ϕ is to be counted as true at w just in case it is a logical consequence of $O_w \cup \{\nu\}$ together with the Dedekind-Peano Axioms; (2) ϕ is counted as false at w just in case its negation is a logical consequence of $O_w \cup \{\nu\}$ together with the Dedekind-Peano Axioms; and (3) ϕ is otherwise counted as lacking a well-defined truth-value relative to w .

¹²For a caveat, see *The Construction of Logical Space*, Section 8.2.3.

‘mixed’ identity statements—i.e. identity statements relating terms of different sorts—were best thought of as ill-formed, and I touted the fact that they *could* be treated as ill-formed as a big advantage of compositionism.¹³ As Gillian Russell points out, however, there are cases in which it is natural to treat mixed identities as *false* rather than meaningless (p. xx). Fortunately, the compositionist would see no obstacle to a specification of truth-conditions that yields this result. In fact, when the details are spelled out in the right kind of way (footnote 11), ‘ $\forall x\forall n(x \neq n)$ ’ will turn out to be true whenever ‘ x ’ is a variable of the original sort and ‘ n ’ is a variable of the new sort.

This delivers a striking conclusion. It entails that the compositionist can think of herself as *extending her domain of discourse* whenever she extends her language in the way just described. It is not, of course, that in extending her language she is making the world grow: the world remains unchanged. The effect of extending her language is, rather, that she acquires additional resources to *describe* the world. And since in this case the additional resources deliver a theory of pure mathematics, what happens is that she acquires additional resources for saying *nothing* (or, more precisely, additional resources for forming sentences with trivially satisfiable truth-conditions). But the compositionist will think that it is nonetheless true that in so doing she expands her domain of discourse, since she acquires terms that refer to objects to which she was previously unable to refer.

In the book I claimed that the compositionist has no good reason to think that there is a definite answer to the question of what objects exist, and therefore no good reason to think that there is such a thing as quantification over ‘absolutely everything’ (§1.5). As Eklund points out (p. xx), my arguments were inconclusive. It now seems to me, however, that the compositionist is in a position to say something stronger than what I said in the book. For, as we just saw, she thinks that there is a procedure that would allow one to start with an arbitrary language and acquire terms that refer to objects outside the domain of that language.

¹³*The Construction of Logical Space*, Section 3.2.3

Eklund is, in effect, describing a version of this procedure when he talks about ‘Turner names’ (p. xx).¹⁴ Simplifying things a bit, the introduction of a Turner name is a special case of the procedure described above: the special case in which one extends the language so as to encompass not an interesting mathematical theory such as ordinary arithmetic, but a boring mathematical theory such as modulo 1 arithmetic, where there is a single ‘number’ with no interesting properties to speak of.¹⁵ From a compositionalist’s point of view, the introduction of modulo 1 arithmetic is no more problematic than the introduction of ordinary arithmetic. To the extent that there is something awkward about introducing such a theory, it is simply to do with the fact that it is so uncompromisingly pointless.

If our discussion is along the right lines, compositionism entails that an arbitrary domain of discourse can be extended, and therefore that there can be no such thing as quantification over ‘absolutely everything’. Philosophers who take such quantification seriously are therefore likely to regard compositionism as suspect. Where, exactly, would they want to get off the boat?

¹⁴Eklund borrows the idea from Turner ‘Ontological Pluralism’, which builds on material from Williamson, ‘Everything’.

¹⁵From a compositionalist perspective, the only significant difference between a Turner name and a name for the modulo 1 ‘number’ is to do with the way in which newly introduced terms interact with pre-existing vocabulary. Let P be an atomic predicate of the original language and let t be a newly introduced term. In the case of arithmetic (modulo 1 or otherwise), it is natural to think that $\lceil P(t) \rceil$ should turn out to be false, or (as I prefer) ill-formed. In the case of the Turner name, on the other hand, one stipulates, amongst other things, that $\lceil P(t) \rceil$ is to count as true for any atomic predicate of the original language.

It is tempting to see this as entailing that a Turner object would have all sorts of interesting properties, if it existed. But that would be a mistake. What happens instead is that the additional stipulations extend the meaning of the original predicates in unexpected ways, and that predications involving a Turner name turn out to have unexpected truth-conditions. Suppose, for example, that the original language includes the atomic predicates ‘Bachelor’ and ‘Married’, both of them standardly interpreted. Then, if t is the Turner name, $\lceil \text{Bachelor}(t) \wedge \text{Married}(t) \rceil$ will turn out to be true: indeed, it will turn out to have trivially satisfiable truth-conditions. But that is not because a Turner object would be a married bachelor. It is because we would have extended our language in a way that entails that ‘Bachelor’ and ‘Married’ have unexpected meanings. ‘Bachelor’, for example, comes to mean something like ‘is a bachelor, if a member of the original domain; is self-identical otherwise’.

Eklund emphasizes that there is something jarring about the idea that a Turner name could be genuinely referential. It seems to me that, as in the case of modulo 1 arithmetic, any awkwardness should be traced back to the fact that the relevant extension of our language would be so uncompromisingly pointless.

My guess is that they would want to resist the idea that our language only makes contact with the world at the level of sentences. They would insist—contrary to what I suggested in Section 1.1—that the world must somehow be ‘responsive’ to a sentence’s compositional structure in order to make the sentence true. In the book I made an effort to spell out a version of this opposing view, and labelled the result ‘metaphysicalism’. I did not, however, identify any contemporary examples of metaphysicalists. What I did instead was suggest that philosophers who take ‘absolutely general’ quantification seriously can be expected to embrace some form of metaphysicalism.

Eklund correctly emphasizes that the transition from absolute generality to metaphysicalism is highly defeasible (p. xx). One can be a compositionalist without noticing that there is a tension between compositionism and absolute generality. And one can reject compositionism by adopting some view other than metaphysicalism.

Although these points are well taken, it is important not to lose track of the big picture. What our discussion brings out is that it is hard to separate the metaphysical question of absolute generality from questions about the relationship between our language and the world it represents. So even if it is not always easy to know how to classify the views of particular philosophers, it can be useful to know that their views about metaphysical issues can carry serious commitments when it comes to the relationship between our language and the world.

2 World

2.1 Metaphysical Structure (Eklund, Sider)

For a feature of reality to be metaphysically structured is for it to have a division into constituent parts that is, in some sense, metaphysically distinguished: a division that is rendered salient not by the way in which *we* happen to represent the relevant feature of reality but, somehow, by the world itself. (As Eklund and Sider point out (p xx,

p. xx), one could also use ‘metaphysical structure’ in a difference sense, whereby a property only counts as being carved out by the metaphysical structure of reality if it is an ‘elite’ or ‘perfectly natural’ property in Lewis’s sense.¹⁶ That is not how I understood ‘metaphysical structure’ in the book, and I now wish I had done more to advertise this.)

As I noted in Section 1.3, I have my doubts about whether there is such a thing as metaphysical structure, and indeed, about whether the notion of metaphysical structure makes good sense. But in the book I was officially *neutral* about such issues. I claimed, in particular, that one could be a compositionalist even if one believed that the world was metaphysically structured.

Eklund suggests that my claims of neutrality may be overblown. He argues that a compositionalist who believed in metaphysical structure might end up with a combination of views that ‘seems distinctly odd’:

(i) there are fine-grained facts, (ii) 0 exists, (iii) the property of being a number exists, (iv) ‘0 is a number’ is true, but (v) yet there is no fine-grained fact $\langle 0, \text{number} \rangle$. (p. xx)

It seems to me that this combination of views is not as bad as Eklund makes it out to be, and would like to use the remainder of this section to explain why.

Consider a compositionalist who believes that ‘0 is a number’ has trivially satisfiable truth-conditions: it will count as true regardless of how the world happens to be. What feature of reality should such a sentence be said to describe? One option would say that—like a truth of pure logic—it describes no particular feature of reality. But perhaps our compositionalist thinks of facts in such a way that every true sentence must describe some fact or other. (She will think that if, like me, she believes that for the fact that p to obtain *just is* for it to be the case that p .) She will then wish to say that ‘0 is a number’ describes a ‘trivial’ fact: a fact that would obtain regardless of how the world turned out to be.

¹⁶Lewis, ‘New Work’.

Now suppose that our compositionalist believes that reality is ‘metaphysically structured’. She believes, for example, that the fact that Socrates died has Socrates and the property of dying as its metaphysically distinguished constituents, and therefore that there is a nice correspondence between the syntactic structure of ‘Socrates died’ and the metaphysical structure of the fact that this sentence describes. But our compositionalist also thinks that this kind of correspondence won’t always obtain. She thinks, for example, that the fact described by ‘Socrates’s death took place’ describes is the very same fact as the fact described by ‘Socrates died’, and therefore that the grammatical structure of ‘Socrates’s death took place’ does *not* correspond to the metaphysical structure of the fact described.

What about ‘0 is a number’? Does our compositionalist think that its grammatical structure corresponds to the metaphysical structure of the fact described? Our compositionalist may well think that the trivial fact—unlike the fact that Socrates died—has no metaphysically distinguished constituents. If so, she will think that ‘0 is a number’ is like ‘Socrates’s death took place’, in that they both have a grammatical structures that fail to match the metaphysical structure of the facts they describe.

It is important to note, however, that none of this would stop our compositionalist from thinking that the number 0 exists. For, according to compositionism, all it takes for a term t to have a referent is for the truth-conditions of $\lceil \exists x(x = t) \rceil$ to be satisfied—and they will certainly be satisfied in this case, since ‘ $\exists x(x = 0)$ ’ is a logical consequence of ‘0 is a number’, which describes the trivial fact. So our compositionalist does *not* think that ‘0’ can only refer to an object if that object is carved out by the metaphysical structure of the fact described by ‘0 is a number’. (For similar reasons, our compositionalist might be assumed to believe that the property of being a number exists.)

Our imagined compositionalist will therefore believe all of the following:

- (i) there is metaphysical structure, (ii) 0 exists, (iii) the property of being a

number exists, (*iv*) ‘0 is a number’ is true, and (*v*) ‘0 is a number’ describes the trivial fact, which does not have 0 and the property of being a number as metaphysically distinguished components.

And, of course, if one understands a fine-grained fact as a fact that is metaphysically structured, this is just a restatement of view that Eklund had described as ‘distinctly odd’. Once it is seen in light of the discussion above, however, the view doesn’t strike me as particularly problematic.

2.2 Why-Closure (Linnebo)

The question ‘Why is it the case that ϕ ?’ can be understood in many different ways. A *straight* reading of the question is one that can be paraphrased as follows:

I can see exactly what it takes to satisfy ϕ ’s truth-conditions and I can see that they are indeed satisfied, but I wish to better understand *why* the world is such as to satisfy them.

A straight reading of ‘Why is there water on Mars?’, for example, might be answered by supplying information about the way in which the Solar System was formed. But it cannot be answered by saying ‘because there is H₂O on Mars’. For this second answer would only be appropriate in a context in which one’s interlocutor doesn’t realize that water is H₂O, and therefore doesn’t fully understand what it takes for the truth conditions of ‘there is water on Mars’ to be satisfied.

One of the key ideas of the book is what might be called the Thesis of Why-Closure—the claim that if ϕ is a first-order sentence that follows from the set of true ‘just is’-statements, then ϕ is *why-closed*: there is no sense to be made of a straight reading of ‘Why is it the case that ϕ ?’.

To illustrate the Thesis of Why-Closure, let us suppose that ‘to be composed of water *just is* to be composed of H₂O’ is a true ‘just is’-statement. Then the conditional ‘if

something is composed of water, it is composed of H_2O ' is a consequence of the set of true 'just is'-statements. So the Thesis of Why-Closure tells us that the conditional should count as why-closed: there should be no sense to be made of a straight reading of the question '*why* are things that are composed of water composed of H_2O ?'. And, indeed, it is natural to reject such a question by saying something along the lines of 'What do you mean *why*? To be composed of water *just is* to be composed of H_2O !'.

If Mathematical Trivialism is correct, then every truth of pure mathematics turns out to be a consequence of the set of true 'just is'-statements.¹⁷ So the Thesis of Why-Closure entails that every truth of pure mathematics should be why-closed. There should, for instance, be no sense to be made of a straight reading of the question 'Why is it the case that $2+2=4$?'.

Linnebo points out that 'the following sequence of equations seems to provide a good explanation of why $2 + 2 = 4$ in terms of the recursion equations governing addition: $2 + 2 = 2 + 1' = (2 + 1)' = (2 + 0)' = (2 + 0)' = 2' = 3' = 4.'$ (p. xx) The trivialist is committed to denying is that this is a good answer to the *straight* reading of 'Why is it the case that $2+2=4$?'. For if *nothing* is required of the world for the truth-condition of ' $2 + 2 = 4$ ' to be satisfied, there can be no good answer to the question of *why* the world is such as to satisfy this condition. But the trivialist can grant that Linnebo's suggested explanation is a good way of answering a *non-straight* reading of the question. There is, for example, a non-straight reading of the question that lends itself to the following paraphrase:

Help me see how the basic rules governing addition come together to secure the truth of ' $2 + 2 = 4$ '.

The trivialist could certainly see Linnebo's explanation as a good way of addressing such a demand. As a result of all this, I am very much in agreement with Linnebo's observation

¹⁷*The Construction of Logical Space*, Section 2.2

that a mathematical sentence can count as why-closed, and still admit of explanation in *some* sense of ‘explanation’. Linnebo anticipates this sort of reaction, and suggests that it won’t take us very far: all one gets is the conclusion that ‘the explanatory burden that [Linnebo believes] should be classified as metaphysical will instead be classified as semantic. The explanatory burden itself will remain unchanged.’ (p. xx)

In a sense, I agree: a mathematical trivialist is under no less pressure to understand why mathematical statements are true than her rivals. At the same time, it seems to me that it can make a difference to think of mathematical explanations as ‘semantic’ rather than ‘metaphysical’. To see this, suppose that one is a trivialist. When asked to explain why there are infinitely many primes, one interprets the request ‘semantically’. One sees the request as amounting to something along the following lines:

Derive ‘there are infinitely many primes’ from principles that a mathematician would regard as basic by using a proof that with such-and-such features.

or perhaps

Derive ‘there are infinitely many primes’ from principles I have a good handle on by using a proof that I am able to get my head around.

On either of these ways of cashing out the explanatory request, it is clear enough what one is being asked to do. So, as understood by the trivialist, there is a satisfactory way of modeling the practice of asking for and giving mathematical explanations.

Now suppose one is *not* a trivialist. One thinks that the truth of ‘there are infinitely many primes’ imposes a *non-trivial* requirement on the world. So one might wonder *why* such a requirement is met. In other words: one might demand an answer to the *straight* reading of ‘Why are there infinitely many primes?’. And the problem with such a question is that it is genuinely unclear whether there is a sensible answer to be given. To see this, notice that the question—on its *straight* reading—could be recast as something along the following lines:

Since I am a non-trivialist, I concede that when God created the Sun’s eight planets—and when she made sure that there were only eight of them—there was more to be done. She would have to do something *extra* to ensure that the number of planets was eight. Similarly, it was not *automatic* that there would be infinitely many primes, as the trivialist believes. I want to understand why this *extra* fact came about. I want to understand why we don’t live in a world in which there are planets, but no number of the planets—and no infinitely many primes.

What would a sensible answer to such a question look like? A causal explanation doesn’t seem to hold much promise. And, to my mind at least, it is unclear what a sensible non-causal explanation could look like. An advantage of adopting trivialism is that one is in a position to reject the need to answer such questions, on the grounds that they are based on false presuppositions. So although Linnebo is right to point out that the trivialist exchanges the burden of a ‘metaphysical’ explanation for the burden of a ‘semantic’ explanation, it seems to me that the exchange is well worth the trouble. (I do not mean to suggest, however, that there is no room to maneuver.¹⁸)

2.3 Explanation (Linnebo)

Some philosophers might take the following to be true:

[TABLE EXPLANATION]

There is a table *because* there are some things arranged table-wise.

Linnebo argues (p. xx) that [Table Explanation] is in tension with the Thesis of Why-Closure (described in the preceding section), if one also accepts:

[TABLES]

For there to be a table *just is* for there to be some things arranged table-wise.

¹⁸Linnebo, ‘Reference by abstraction’.

I agree with Linnebo that there is a tension here. But I don't think the Thesis of Why-Closure is needed to generate it. It seems to me that [Table Explanation] is in tension with [Tables], regardless of one's views about the Thesis of Why-Closure.

The reason is simple. [Table Explanation], as Linnebo understands it, is an explanatory claim. It states that the fact that there is a table is *explained* by the fact that there are some things arranged table-wise. But whether or not one accepts the Thesis of Why-Closure, one should take [Tables] to entail that 'there is a table' and 'there are some things arranged table-wise' are descriptions of the very same feature of reality. So—on the plausible assumption that nothing can explain itself—it follows from [Tables] that [Table Explanation] cannot be correct.

Since I am inclined to think that [Tables] is true, Linnebo is right to suggest that I am committed to rejecting [Table Explanation]. Linnebo goes on to suggest, however, that such a rejection is something I am 'deeply committed to' (p. xx), and with this I disagree.

I am deeply committed to the idea that our conception of what needs explaining and what doesn't is determined by the system of 'just is'-statements we accept. I am also deeply committed to a particular methodology for deciding which 'just is'-statements to accept. But I am *not* deeply committed to any particular choice of 'just is'-statements. It is true that there are 'just is'-statements—like [Numbers], from Section 1.2—that I take to be strongly supported by the methodological considerations I articulate in the book. It is also true that there are particular 'just is'-statements I find attractive, and that [Tables] happens to be one of them. But the aim of the book is emphatically *not* to defend 'just is'-statements like [Tables]. Quite the contrary: my hope is that the basic theoretical framework of the book—and, in particular, my claims about the relationship between 'just is'-statements and why-closure—might be used to accommodate a broad range of philosophical outlooks, with differences between these outlooks modeled, in part, as differences in the 'just is'-statements that their proponents would choose to adopt. In

particular, my hope is that the basic theoretical framework of the book might be used to accommodate a friend of [Table Explanation], and that it might be used to model the disagreement between us: since the friend of [Table Explanation] will be forced to reject [Tables], our disagreement can be modeled, in part, as a disagreement about which ‘just is’-statements to adopt.

It seems to me, moreover, that the methodology I set forth in the book for deciding which ‘just is’-statements to adopt can be used to move the debate forward. For I think the methodology can be used to argue that rejecting [Tables] comes at a cost. One is left with a theoretical gap between the existence of tables and the existence of things arranged table-wise, a gap that leaves room for lines of inquiry such as the following:

I can see that there some particles arranged table-wise, and I can see that if it was *also* the case that there was a table, the table’s existence would be *explained* by the existence of the particles arranged table-wise. But why should I believe that there is actually a table?

A friend of [Table Explanation] might think that such questions turn out to have interesting answers: answers that would lead to fruitful lines of inquiry. I am skeptical that this is so. But my skepticism is not a ‘deep’ commitment. Show me an interesting enough theory of the connection between tables and particles arranged table-wise, and I might lean the other way.

2.4 Grounding (Cameron)

As I understand it in the book, the ‘just is’-operator is a no-difference operator: I take ‘for it to be the case that p *just is* for it to be the case that q ’ to be equivalent to ‘there is *no difference* between its being the case that p and its being the case that q ’. This means, in particular, that I take the ‘just is’-operator to be *symmetric*: if ‘ \equiv ’ is the ‘just is’-operator, then ‘ $p \equiv q$ ’ and ‘ $q \equiv p$ ’ will always be equivalent.

There are philosophers who work with a notion of grounding that is taken to be *asymmetric*: if the fact that p is grounded in the fact that q , then it is *not* the case that the fact that q is grounded in the fact that p . Ross Cameron asks an interesting question: ‘Suppose we come to the table thinking of things in terms of grounding. Are we able to use the resources we are happy with to understand Rayo’s symmetric just is statements?’ (p. xx) He then suggests the following answer: ‘the prospects for defining a symmetric notion in terms of grounding that will also let us assent to the kind of just is claims that Rayo assents to are dim’ (p. xx)

It seems to me that it is important to distinguish between two different issues. The first is the question whether one can define the ‘just is’-operator using vocabulary that the friend of grounding is already committed to. The second is the question whether a friend of grounding would have to disagree with me about which ‘just is’-statements to accept.

I would regard a negative answer to the first question as (mildly) disappointing news. For if the ‘just is’-operator turns out to be definable on the basis of grounding-friendly notions, proponents of grounding would have an easy way of coming to understand the ‘just is’-operator. If it isn’t, then proponents of grounding will be in the same boat of the rest of us: they will have no choice but to understand the notion on the basis of its theoretical role.

In contrast, I would *not* regard a negative answer to the second question as disappointing news. For, as I emphasized in the preceding section, my hope is that the basic theoretical framework of the book might be used to accommodate a broad range of philosophical outlooks, with differences between them modeled, in part, as differences in the ‘just is’-statements their proponents choose to adopt.

To see how this might come about in the case at hand, consider the following grounding claim:

[TABLE GROUNDING]

The fact that there is a table is grounded in the fact that there are some things arranged table-wise.

On the assumption that no fact can ground itself, [Table Grounding] entails that the fact that there is a table is *distinct* from the fact that there are some things arranged table-wise. It therefore entails that the following ‘just is’-statement is false:

[TABLES]

For there to be a table *just is* for there to be some things arranged table-wise.

So someone who accepts [Table Grounding] would disagree with me about whether to accept [Tables], and will do so regardless of whether she thinks that the ‘just is’-operator is definable on the basis of grounding-friendly notions. It seems to me, however, that such a result is good news rather than bad, since it gives us a useful way of characterizing my disagreement with the proponent of grounding.

There is a different way of thinking about grounding where we don’t get a conflict with [Tables]. Suppose that instead of accepting [Table Grounding] one were to accept the following:

[TABLE FUNDAMENTALITY]

What it is, fundamentally speaking, for there to be a table is for there to be some things arranged table-wise.

Suppose, moreover, that [Table Fundamentality] is to be understood as the conjunction of two theses: (a) ‘there is a table’ and ‘there are things arranged table-wise’ describe the *same* fact, and (b) at the fundamental level, the relevant fact concerns things arranged table-wise rather than tables.

On this way of seeing things, [Table Fundamentality] is not in conflict with [Tables]: it *entails* [Tables], because of claim (a). So the friend of [Table Fundamentality] agrees with me about [Tables], even though we disagree about claim (b). Correspondingly, the friend

of [Table Fundamentality] disagrees with the friend of [Table Grounding] about [Tables], even though they are both grounders. The lesson of all this is that a ‘just is’-statement like [Tables] can give us a nice way of modeling certain agreements and disagreements in the debate about grounding. More generally, it seems to me that the basic theoretical framework of the book should be available to friends and foes of grounding alike, and that differences between different sorts of views can sometimes be modeled as differences in the ‘just is’-statements one chooses to adopt.

I would now like to return to the first component of Cameron’s line of questioning: the problem of whether is it possible to define the ‘just is’-operator on the basis of vocabulary that the friend of grounding is committed to already.

Some proponents of grounding are committed to fact-talk. For them the task of defining the ‘just is’-operator should be totally straightforward:

$$\lceil p \equiv q \rceil =_{df} \lceil \text{the fact that } \bar{p} = \text{the fact that } \bar{q} \rceil$$

where $\bar{r} = r$ if r is true, and $\bar{r} = \lceil \neg r \rceil$ if r is false. Of course, a proponent of grounding might wish to resist fact-talk—or she might be happy to indulge in fact-talk but deny that $\lceil \text{the fact that } p \rceil$ picks out a fact for every true sentence p .

An alternative strategy is based on the observation that if the proponent of grounding is committed to a notion of truth-conditions, she could take $\lceil p \equiv q \rceil$ to be true just in case p and q have the same truth-conditions. As Cameron points out, Sider’s brand of grounding presupposes a notion of truth-conditions that might be used for such a purpose.¹⁹ One could claim, in particular, that $\lceil p \equiv q \rceil$ is true just in case Sider’s ‘metaphysical semantics’ entails $\lceil ‘p’ \text{ is true} \leftrightarrow ‘q’ \text{ is true} \rceil$.

Cameron worries, however, that such a definition would lead to unwelcome results, on the grounds that different ‘tautologies get different truth-conditions in [Sider’s] metaphysical semantics; and yet it seems like if there are any two claims such that there is ‘no

¹⁹Sider, *Writing the Book of the World*.

difference' between the demands they make on the world, two tautologies should meet that criterion.' (p. xx)

I agree that it would be odd to suggest that different tautologies place different demands on the world, but I don't think the oddness should cast doubt on the idea that the notion of truth-conditions can be used to characterize the 'just is'-operator. If one thinks that tautologies \top_1 and \top_2 have have different truth-conditions, then one should certainly reject $\lceil \top_1 \equiv \top_2 \rceil$. In doing so one would commit oneself to a particularly fine-grained conception of logical space.²⁰ It is not the conception of logical space that I myself would favor, but I think it can be accommodated by the basic theoretical framework of the book.

2.5 Alstonian Symmetry (Cameron)

Let 'brain state B ' be shorthand for a description of a certain brain state in physical terms, and consider the following 'just is'-statement:

[RED]

To experience the sensation of seeing red *just is* to be in brain state B .

Someone who accepts [Red] thinks that the following line of questioning is based on a mistake:

I can see that the subject is in brain state B . But I would like to know how one could ever be justified in believing that a certain *further* fact obtains:
the fact that the subject is experiencing the sensation of seeing red. And if it

²⁰Depending on how the details are spelled out, one might also end up with the conclusion that tautologies fail to have have trivially truth-conditions. So, as Cameron points out, one won't be able to say that ϕ has trivially satisfiable truth-conditions just in case $\lceil \phi \leftrightarrow \top \rceil$ is true, for \top a tautology. From my point of view, this would, again, be an odd result. But, again, I don't think the oddness should cast doubt on the idea that the notion of truth-conditions can be used to characterize the 'just is'-operator. It is simply a consequence of the view that tautologies fail to have trivially satisfiable truth-conditions. (For further discussion, see Section 1.3.)

turns out that this further fact does indeed obtain, I would like to understand *why* it obtains. Why isn't it instead the case that the subject is a zombie?

If [Red] is true, there is no such thing as the supposed *further* fact: for the subject to be in brain state *B* is *already* for her to experience the sensation of seeing red. In contrast, someone who rejects [Red] is committed to thinking that there is no mistake, and that the ensuing explanatory and justificatory demands need to be addressed. How might she go about addressing them? One option would be to postulate 'qualia'—intrinsic, non-representational, properties of sensory experiences—and claim that they are not reducible to physical properties. She could then go on to suggest that what it takes for a subject in brain state *B* to *also* experience the sensation of seeing red is for her to bear the right kind of conscious access to a non-physical 'redness quale'. Finally, she could claim that the reason our subject is not a zombie is that she is lucky enough to bear this relationship to a suitable qualia.

The lesson of all this, I suggest, is that an advantage of accepting [Red] is that it relieves one from the need to countenance non-physical qualia. Drawing inspiration from a famous argument due to William Alston,²¹ Cameron thinks that such a conclusion cannot go through without additional argumentation: 'If I start out worried that when I say 'I am experiencing redness' [...] I am committing myself to an immaterial redness quale, how is my worry eased by coming to learn that to experience redness just is to be in a certain brain state? Why shouldn't I just be worried now that believing in brain states is to commit to immaterial qualia?'

I agree with Cameron that a supplementary argument is needed, but I disagree with him about the shape that the argument should take. My own view is best introduced with an example. Suppose that one starts out believing that to experience the sensation of seeing red is to bear the right kind of relationship to a non-physical 'redness quale'. Accordingly, one believes that whether or not someone is a zombie—and, in particular,

²¹Alston, 'Ontological Commitments'.

whether or not she is experiencing the sensation of seeing red—is not the sort of issue that could be settled by giving a full description of her brain in purely physical terms. At the same time, one thinks that whether or not the subject is in brain state B is the sort of issue that could be settled by giving a full description of her brain in purely physical terms.

At a later time, one comes to accept [Red]: one comes to believe that there is *no difference* between being in brain state B and experiencing the sensation of being red. As a result, one can no longer believe *both* that one can settle whether the subject is in brain state B by giving a full description of her brain in physical terms *and* that one cannot settle whether the subject experiencing the sensation of seeing red by giving a full description of her brain in physical terms. One of these claims must be given up, and Cameron’s questions will be answered in different ways depending on which of the two one chooses to reject.

Suppose, first, that in accepting [Red] one abandons the view that one cannot settle whether the subject is experiencing the sensation of seeing red by giving a full description of her brain in physical terms. Then one must either give up on qualia altogether, or come to think that qualia are ultimately reducible to the physical. Either way, the result of accepting [Red] is that one will be left with no reason to worry about non-physical qualia.

Now suppose that one instead decides to retain the view that to experience the sensation of seeing red is to bear the right kind of relationship to a non-physical ‘redness quale’. Accordingly, one accepts [Red] by abandoning the view that one can settle whether the subject is in brain state B by giving a full description of her brain in physical terms. Then the result of accepting [Red] is that one should worry that brain state B requires the existence of non-physical qualia.

3 Language

3.1 Semantics (Hofweber)

Hofweber suggests that ‘one would expect considerations about the semantics of English to be relevant for this question [of which ‘just is’-statements to adopt]’. (p. xx) I do claim, after all, that for the ‘just is’-statement $\lceil p \equiv q \rceil$ to be true is for p and q to have the same truth-conditions. But, as Hofweber goes on to observe, there is virtually no mention of natural language semantics in the book. Isn’t this an important omission?

It seems to me that it is not. To see what I have in mind, consider the following ‘just is’-statement:

[WATER]

For the glass to be filled with water *just is* for the glass to be filled with H_2O .

A natural language semantics should certainly not be expected to settle the question of whether [Water] is true—at least not on its own. For even if one expects a natural language semantics to deliver an assignment of *meanings* to words like ‘water’ and ‘ H_2O ’, in some sense of ‘meaning’, one shouldn’t expect it to settle the question of whether ‘water’ and ‘ H_2O ’ express the same *property*, since that would require taking a stand on the *chemical* hypothesis that water is H_2O , and natural language semantics shouldn’t be in the business of making chemical hypotheses.

I am not claiming that linguistic theorizing would be entirely useless to the project of determining whether [Water] is true. The point is just that I don’t think that’s where the action is. If one really wants to know whether to accept [Water], one should turn to chemistry rather than linguistics.

One might be tempted to resist this thought by claiming that there is a certain kind of division of labor. For even if one concedes that it is not the job of natural language semantics to identify the property expressed by ‘water’, one might think that

it is the job of natural language semantics to identify a *function* that would output such a property given suitable information about the way the world is. Perhaps one thinks that the semantics would assign a *primary intension* to ‘water’.²² One could then claim that when applied to the actual world, the primary intension of ‘water’ delivers the property of being composed of H₂O, and when applied to a world in which the lakes and rivers are filled with XYZ, the primary intension of ‘water’ delivers the property of being composed of XYZ. This would allow one to see the task of ascertaining whether [Water] is true as divided into two distinct parts. First, one uses a natural language semantics to associate suitable primary intensions with ‘water’ and ‘H₂O’. Second, one inputs information about the chemical world into these primary intensions and gets the properties expressed by ‘water’ and ‘H₂O’. One can then determine whether [Water] is true by ascertaining whether these properties are one and the same.

I myself am skeptical of the hypothesis that a word like ‘water’ has a well-defined primary intension.²³ But even if it does, it should be agreed on all sides that nobody knows how to identify it in practice. And it is worth emphasizing that contemporary linguistics isn’t much help here. For linguistic theory—and even the branch of linguistic theory that is usually referred to as ‘semantics’—has much more to say about *grammar* than about lexical meaning. As far as I know, there is no work in contemporary linguistics aimed at identifying anything like the primary intentions of our words.

If one wants to determine whether [Water] is true, the only sensible course of action I can think of is to set linguistics aside, and use one’s basic linguistic competence, together with one’s knowledge of chemical phenomena, to form hypotheses about the property that might be expressed by ‘water’. Different hypotheses will lead to more or less fruitful theorizing and do so while adhering more or less faithfully to ordinary linguistic usage. So one should select a hypothesis that leads to fruitful theorizing without diverging too

²²Chalmers, *The Conscious Mind*; Jackson, *From Metaphysics to Ethics*.

²³See Block and Stalnaker ‘Conceptual Analysis’.

drastically from ordinary usage, and assess the truth of [Water] with respect to such a hypothesis. (In the special case of [Water], there is a unique hypothesis that stands out in light of our best chemical theorizing—the hypothesis that ‘water’ expresses the property of being composed of H_2O —but in other cases one’s selection may require a certain degree of arbitrariness.²⁴)

Our discussion has so far been focused on the particular example of [Water]. But what about ‘just is’-statements of other kinds? I think Hofweber is right to think that there may be cases in which natural language semantics might help us ascertain whether a given ‘just is’-statement is true. In other work,²⁵ he has argued that ‘There are (exactly) four moons of Jupiter’ and ‘The number of moons of Jupiter is four’ are linguistically equivalent in a way that guarantees that they must have the same truth-conditions. If this is right, we will have purely linguistic grounds for accepting [Numbers] from Section 1.2. Having an additional argument for [Numbers] would certainly constitute welcome news. But I still think that one shouldn’t put too much weight on linguistic considerations in arguing for [Numbers]. For although there is much to be said about Hofweber’s linguistic arguments, they are ultimately controversial,²⁶ and one wouldn’t want one’s philosophy of mathematics held hostage to the outcome of a debate in linguistics.

3.2 Objective Correctness (Hofweber, Russell)

In the book I consider a natural way of thinking about the notion of truth (§2.3). To set forth a statement is to make a *distinction* amongst ways for the world to be, and to single out one side of that distinction; for the statement to be *true* is for the region singled out to include the way the world actually is. (To set forth the statement that snow is white, for example, is to distinguish between white-snow and non-white-snow ways for the world to be, and to suggest that the world falls on the white-snow side of

²⁴For further discussion, see *The Construction of Logical Space*, Section 2.1.

²⁵Hofweber, ‘Innocent Statements’.

²⁶Jackson, ‘Defusing Easy Arguments’; Moltmann, ‘Reference to Numbers’.

this distinction; for the statement to be true is for the way the world actually is to fall on the white-snow side of the distinction—for it to actually be the case that snow is white.)

Natural as it is, this way of thinking about truth has a surprising consequence. The best way to see the point is to consider the matter from the perspective of someone who thinks that to be composed of water *just is* to be composed of H₂O. Our subject believes, in other words, that the distinction between water-containing ways for the world to be, and the rest, is just the distinction between H₂O-containing ways for the world to be, and the rest.

Suppose that our subject sets out to assess the following sentence for truth or falsity:

[No H₂O]

There is water in the world, but no H₂O.

On the conception of truth that is under discussion, she would have to start by distinguishing between ways for the world to be whereby there is water but no H₂O, and the rest, and then consider the question of which side of this distinction includes the way the world actually is. But our subject is committed to thinking that the relevant distinction is *trivial*, since she is committed to thinking that there is no difference between ways for the world to be whereby there is water and ways for the world to be whereby there is H₂O. This means that, from her perspective, it is impossible for the way the world actually is to supply an *independent check* on the question of whether [NO H₂O] true. As soon our subject has laid down her conception of *logical space*—as soon as she has decided upon a set of distinctions with which to discriminate between ways for the world to be—the issue has been settled. There is no need to look at the way the world actually is.

Not so when it comes to the statement that snow is white. In this case, even after our subject has laid down her conception of logical space—even after she has decided upon a set of distinctions with which to discriminate between ways for the world to be—there

is still room for the way the world actually is to supply an independent check on the question of whether the statement is true. For even if her conception of logical space shapes the way in which she sees statements as distinguishing between ways for the world to be, it does *not* settle the question of whether the way the world actually is falls on one side or the other of the distinction between white-snow and non-white-snow ways for the world to be.

Now consider the question of whether our subject's 'just is'-statement—'to be composed of water *just is* to be composed of H₂O'—is objectively correct. What could one mean by 'objective correctness' here? If all one means is *truth*—and if one is thinking of truth in the way I suggested above—then it's not clear how questions of objective correctness could be settled by the way the world is. For although one could certainly point out that the subject's 'just is'-statement is true relative to her own conception of logical space, that simply leads to the question of whether her conception of logical space constitutes an 'objectively correct' framework with which to assess the truth of a given 'just is'-statement. And it is not clear that this is something that can be settled by the way the world is. For it is not clear that in adopting a conception of logical space one commits to one amongst different ways for the world to be, rather than simply adopting a family of *distinctions* with which to discriminate between ways for the world to be.

We seem to be trapped in a circle. We want the way the world is to settle the question of which conception of logical space is objectively correct, but the way the world is can only make such determinations against the background of a *distinction* between ways for the world to be, and therefore against the background of a conception of logical space. Once we adopt a conception of logical space, however, the issue of which conception of logical space is correct gets settled independently of the way the world is. For none of the distinctions one adopts in adopting a conception of logical space will distinguish between that conception of logical space and one of its rivals. So one's conception of logical space will be vindicated regardless of where in one's system of distinctions the way the world

actually happens to be located.

Hofweber objects that ‘none of this shows that there is no objectively correct logical space, or no objectively true ‘just is’-statements.’ (p. xx) I think he is right. The most my argument shows is that—on a certain way of thinking about truth—it is hard to see how the way the world is could have any impact on the truth or falsity of a ‘just is’-statement (and therefore on the objective correctness of a conception of logical space). But this does not entail there is no such thing as objective truth for ‘just is’-statements. All it shows that one cannot do much to elucidate such a notion, while using a particular set of theoretical tools.

It is partly because my arguments were inconclusive that I tried to remain neutral about these issues in the book. I did, however, report a suspicion that there is no such thing as a notion of *objective* truth or falsity for ‘just is’-statements: all one can say is that a given subject should accept or reject particular ‘just is’ statements depending on whether they lead to fruitful theorizing, and that *relative* to the set of ‘just is’ statements one accepts the statements in that set will count as true.²⁷

Hofweber uses ‘Carnapian Rayo’ as label for someone who accepts this position, and argues that Carnapian Rayo must be mistaken. For consider the following ‘just is’-statement:

[TENOR]

To be Thomas Hofweber *just is* to be the greatest tenor ever.

Hofweber asks: ‘is acceptance and rejection of [Tenor] partly a practical question, the result of cost-benefit analysis on how to conceive of logical space? And is its truth determined by these considerations?’ He thinks the answer is clear: ‘I see little hope for Carnapian Rayo here. [Tenor] is simply false, and no cost-benefit analysis will make any [difference] to that.’ (p. xx)

²⁷*The Construction of Logical Space*, Section 2.3.

How should Carnapian Rayo respond? Recall that what is distinctive about Carnapian Rayo's view is that he thinks that there is no sense to be made of an 'objective' notion of truth for 'just is'-statements: all one can say is that a particular subject should accept or reject particular 'just is' statements depending on whether they lead to fruitful theorizing, and that *relative* to the set of 'just is' statements one accepts the statements in that set will count as true (and the rest will count as false). So the way Carnapian Rayo will want to respond to Hofweber's challenge is *not* by conceding that [Tenor] is objectively false—Carnapian Rayo believes that there is no such thing—but by arguing that one should reject [Tenor] regardless of one's theoretical goals.

Let us therefore consider the question how to determine whether to accept or reject a given 'just is'-statement, according to Carnapian Rayo. In the preceding section I suggested that the way to determine whether to accept [Water] is to use one's linguistic competence, together with one's knowledge of chemical phenomena, to form hypotheses about the property expressed by 'water', and then select a hypothesis that leads to fruitful theorizing without diverging too drastically from ordinary linguistic usage. Similarly, I think Carnapian Rayo should say that the way to ascertain whether [Tenor] is true is to use one's linguistic competence, together with one's knowledge of relevant worldly phenomena, to form hypotheses about the object picked out by 'Thomas Hofweber' and the property expressed by 'greatest tenor', and then select a hypothesis that leads to fruitful theorizing without diverging too drastically from ordinary linguistic usage.

There may well turn out to be more than one hypothesis that deserves to be taken seriously. Suppose that on one hypothesis 'Thomas Hofweber' refers to the present Hofweber time-slice, and that, on another hypothesis, it refers to the mereological fusion of Hofweber time-slices. If one of these hypotheses leads to more fruitful theorizing than the other, then there might be a case for accepting it, even if it diverges from ordinary usage to some extent. Taking 'Thomas Hofweber' to refer to anything in the vicinity of Plácido Domingo, on the other hand, would constitute too much of a divergence from ordinary

usage, no matter how theoretically useful. Similarly, taking ‘greatest tenor’ to express anything unrelated to quality singing in the high male register would constitute too much a divergence from ordinary usage, no matter how theoretically useful. So, unless Hofweber has been keeping quiet about his singing abilities, there is no way of getting the material biconditional corresponding to [Tenor] to turn out to be true without doing unacceptable violence to ordinary linguistic usage, regardless of one’s theoretical goals. (Getting [Tenor] itself to turn out to be true would be harder still, since ordinary usage suggests that ‘Thomas Hofweber’ ought to be treated as a rigid designator, unlike ‘the greatest tenor ever’.)

Notice, incidentally, that this way of thinking does not presuppose that different hypotheses about the referent of ‘Thomas Hofweber’ or the property expressed by ‘greatest tenor’ must correspond to different assignments of meanings to our terms, in a pre-theoretic sense of ‘meaning’. That is an issue that will depend on how robust our pre-theoretic notion of meaning turns out to be. If one assumes that the meanings of our words are definite enough to output unique assignments of reference given sufficient information about the way the world is, then at most one referential hypothesis for a given term can respect its meaning. But, to my mind at least, it seems implausible to suppose that one could identify a notion of meaning as robust as that: linguistic usage is just too messy and parochial.

As a contrast case, consider the expression ‘members of the same species’. On one hypothesis, ‘members of the same species’ expresses the property of sharing a lineage of the right kind; on another, it expresses the property of being members of a group capable of generating fertile offspring. On the assumption that neither of these hypotheses does unacceptable violence to ordinary usage, one might turn to considerations of theoretical fruitfulness in deciding which of these hypotheses to accept. And since judgments of theoretical fruitfulness depend, in part, on one’s theoretical goals, it may well turn out that people with different theoretical goals favor different hypotheses. Accordingly, the

decision to accept a ‘just is’-statement like

[SPECIES]

To be members of the same species *just is* to share a lineage of the right kind.

may ultimately depend on one’s theoretical goals.

The lesson of our discussion is that even if there are cases like [Species] in which Carnapian Rayo thinks that the decision to accept a ‘just is’-statement should be sensitive pragmatic considerations, there are also cases like [Tenor] in which he thinks that one should accept (or reject) the relevant ‘just is’-statement regardless of one’s theoretical goals. It is for that reason that Carnapian Rayo can do justice to Hofweber’s intuition that ‘[Tenor] is simply false, and no cost-benefit analysis will make any [difference] to that’.

Russell agrees with Hofweber that Carnapian Rayo’s views have unpalatable consequences. She brings out the problem by considering a puzzle of personal identity over time. Suppose ‘your brain will be gradually repaired and ultimately replaced with mechanical and computer parts, serving the same functions as the parts they replace. Call the resultant person, with the computer parts in her or his head, ‘A’. Meanwhile your original brain parts will be preserved, reconstructed, and plumbed and wired into the head of a new body. Call the resulting person, with your old brain parts, ‘B’. Now you learn that person B is to be tortured, whereas A will not’ (p. xx).

Russell asks us to consider the question of whether the following ‘just is’-statement is true:

[TORTURE]

For person B to be tortured *just is* for you to be tortured.

She then argues that such a statement must be either objectively true or objectively false:

it just seems wrong to say that there is no objective fact of the matter [about whether [Torture] is true], or that whether or not it is true depends on which System of Representation, or Linguistic Framework, or Conception of Logical Space one chooses to accept. If only changing your Conception of Logical Space could prevent you from being tortured! (p. xx)

Here it is worth remembering Honest Abe: ‘How many legs does a dog have if you call the tail a leg? Four. Calling a tail a leg doesn’t make it a leg.’ Even if one agrees with Carnapian Rayo that the decision to accept a ‘just is’-statement like [Species] should be sensitive to pragmatic considerations, one shouldn’t think that one could change any biological facts by changing one’s conception of logical space. In changing one’s conception of logical space one changes the set of *distinctions* one uses to represent biological facts, but the facts themselves remain unchanged. All Carnapian Rayo is committed to is the claim that changing one’s theoretical goals—for example, turning one’s attention from biological lineage to the ability of groups to interbreed—can lead to a change in the property one should take ‘members of the same species’ to express. And this strikes me as just the right thing to say in the case at hand: researchers immersed in different lines of inquiry might find it useful to take different properties to be expressed by ‘members of the same species’. There is no obvious reason to stop them, unless our best biological theorizing turns out to give one of these properties an especially important role.

Similarly, it seems to me that when it comes to [Torture], the key issue is not whether one could keep oneself from being tortured by changing one’s theoretical goals: Carnapian Rayo would agree that one can’t. The key issue is whether a change in one’s theoretical goals could lead to a change in the set of *distinctions* that one should use to theorize about the world—and, in particular, a change in the objects one should take ‘person B’ and ‘you’ (as used in a particular context) to refer to.

I do not deny that Russell’s [Torture] poses a challenge for Carnapian Rayo. But I

think the challenge is similar to the challenge posed by Hofweber's [Tenor]. In both cases, we have the intuition that the truth-value of the relevant 'just is'-statement should be independent of one's theoretical goals, and in both cases the challenge is to explain how Carnapian Rayo could do justice to this intuition. I argued earlier that Carnapian Rayo is in a position to do justice to the intuition in the case of [Tenor], and I see no reason to think that he couldn't use a similar strategy do justice to the intuition in the case of [Torture].

Carnapian Rayo will take the status of [Torture] to depend on what our best theory of personal identity turns out to be. Suppose, for example, that philosophers were to coalesce around the view that there are immortal souls, and that whether or not you are identical to Russell's person B depends entirely on whether your soul will be transferred onto B's body after the relevant procedure. Then Carnapian Rayo would say that one should accept [Torture] if and only if the relevant soul-transfer takes place, and he would add that one should do so independently of one's theoretical goals. So, in accordance with the initial intuition, he will claim that the decision to accept [Torture] should not be sensitive to pragmatic considerations.

More generally, if philosophers were to come up with a good way of settling the question of whether you are person B, then Carnapian Rayo would follow suit by using those very considerations to argue that [Torture] ought to be accepted, regardless of one's theoretical aims. And if philosophers were unable find a good way of settling the question of whether you are person B, then Carnapian Rayo would follow suit by claiming that, even when conjoined with our best philosophical theorizing about personal identity, the facts about ordinary linguistic usage are not enough to settle the question of what 'person B' and 'you' (as used in a given context) should be taken to refer to. Accordingly he would claim that the question of whether or not one ought to accept [Torture] is something that cannot be decided independently of the purposes at hand.

In this latter scenario, Carnapian Rayo will have failed to respect the initial intuition

that the truth value of [Torture] should be independent of one's theoretical goals. But it is important to be clear that this has nothing to do with his controversial views about the objective truth for 'just is'-statements. His conclusions so far are simply a consequence of the assumption that philosophers are unable to make progress on questions of personal identity, together with a plausible-sounding story about how to decide amongst different hypotheses about the referents of our terms.

Notice, moreover, that Carnapian Rayo has made no claims about the *feature of reality* that Russell's transplant story describes. In particular, he has not claimed that there's no fact of the matter about the way the world is when the torturing takes place. What Carnapian Rayo claims is that—on the assumption that philosophers are unable to come to any definite conclusions about personal identity—our best theoretical considerations are insufficient to decide the question of what referents to assign to 'person B' and 'you' (as used in the relevant context), and that they are therefore insufficient to decide the question of what sorts of *distinctions* to use in representing the way the world is when the torturing takes place. To my mind, at least, this is an eminently sensible conclusion. (Semantic ascent can make all the difference: compare the claim that you are different people in different contexts with the claim that 'you' refers to different people in different contexts.)

3.3 Carnapian Rayo v. Mainstream Rayo (Hofweber)

In the preceding section I did my best to defend Carnapian Rayo against a certain kind of challenge. It seems to me, however, that the main theses of the book are independent of whether Carnapian Rayo's views are ultimately sustainable. Hofweber disagrees. He thinks that '[i]f there is an objective logical space, and if 'just is'-statements are just like other statements, then much of what is to come [in the book] doesn't really make much progress over the mainstream.' (p. xx)

To illustrate the point, he introduces Mainstream Rayo. Unlike Carnapian Rayo, who

‘rides a motorcycle with no helmet on’, Mainstream Rayo is so decidedly unadventurous that when he ‘wants to do something far out he listens to an unusual interpretation of Beethoven.’ In an effort not to come across as boring, however, Mainstream Rayo ‘sometimes puts labels on common views that make them seem more radical than they are.’ (p. xx) For instance, he will sometimes restate the perfectly conventional claim that ‘necessarily, the number of the dinosaurs is zero just in case there are no dinosaurs’ as a ‘just is’-statement: ‘for the number of the dinosaurs to be zero *just is* for there to be no dinosaurs’. But for Mainstream Rayo the ‘just is’ locution is just for show: it leads to no interesting philosophical conclusions.

Hofweber thinks that unless one is prepared to interpret me as Carnapian Rayo, it is hard to see how the book could deliver much by way of a novel philosophical position. For unless one follows Carnapian Rayo in rejecting the idea that it makes sense to speak of an objectively correct conception of logical space, my ‘view will likely end up like Mainstream Rayo’s’ (p. xx)

I certainly agree with Hofweber that one gets a much more radical version of the proposal if one interprets me as Carnapian Rayo. But I don’t think one should conclude from this that the only alternative to Carnapian Rayo is Mainstream Rayo. It seems to me that the two central doctrines of the book are weak enough not to require Carnapian Rayo’s views about objective truth but strong enough to go beyond Mainstream Rayo’s unadventurous commitments. The first doctrine is compositionality, which we discussed in Section 1.1. The second doctrine is the claim that one’s views about which explanatory and justificatory demands should be treated as legitimate are shaped by the ‘just is’-statement one accepts.

I offered an example of the second doctrine in Section 2.5. Let me mention another. Suppose that we take the following ‘just is’ statement to be true:

[DINOSAURS]

For the number of the dinosaurs to be zero *just is* for there to be no dinosaurs.

Then, the basic theoretical framework of the book entails that we should dismiss the following line of questioning as illegitimate:

I can see that there are no dinosaurs. But I would like to know how one could ever be justified in believing that a certain *further* fact obtains: the fact that the number of the dinosaurs is zero. And if it turns out that this further fact does indeed obtain, I would like to understand *why* it obtains. Why isn't it instead the case that there are no numbers?

The second doctrine delivers the result that this line of reasoning carries a false presupposition. For it assumes that there is a *further* fact where there is none: the fact that there are no dinosaurs is *already* the fact that the number of the dinosaurs is zero. So there is no need to justify the claim that the supposed further fact obtains, and there is no need to explain why it obtains if it does.

In contrast, it is consistent with Mainstream Rayo's views that there *is* a further fact, and that the line of questioning ought to be taken seriously. It seems to me that this is an unhappy place to end up, because efforts to answer the relevant explanatory and justificatory demands have led to all sorts of intractable problems.²⁸ On the other hand, I think that someone who rejects the demands as illegitimate is in a position to develop an attractive philosophy of mathematics.²⁹

In any case, I might as well reveal that I am, in fact, Carnapian Rayo.

²⁸A canonical statement of these problems is in Benacerraf, 'Mathematical Truth' .

²⁹*The Construction of Logical Space*, Ch. 3 and 4.

References

- Alston, William. 'Ontological Commitments'. *Philosophical Studies*, 9 (1957) 8–17.
- Benacerraf, Paul. 'Mathematical Truth'. *Journal of Philosophy*, 70 (1973) 661–79.
Reprinted in Benacerraf and Putnam (1983).
- Benacerraf, Paul and Hilary Putnam, eds. *Philosophy of Mathematics*. Second ed. (Cambridge: Cambridge University Press, 1983).
- Block, Ned and Robert Stalnaker. 'Conceptual Analysis, Dualism, and the Explanatory Gap'. *Philosophical Review*, 108 (1999) 1–46.
- Chalmers, David. *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press, 1996).
- Dummett, Michael. *Frege: Philosophy of Language*. Second ed. (Cambridge, MA: Harvard, 1981).
- Frege, Gottlob. *Die Grundlagen der Arithmetik* (1884). English Translation by J.L. Austin, *The Foundations of Arithmetic*, Northwestern University Press, Evanston, IL, 1980.
- Hofweber, Thomas. 'Innocent Statements and Their Metaphysically Loaded Counterparts'. *Philosophers' Imprint*, 7 (2007) 1–33.
- Jackson, Brendan Balcerak. 'Defusing Easy Arguments for Numbers'. *Linguistics and Philosophy*, (forthcoming) 1–15.
- Jackson, Frank. *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (Oxford: Oxford University Press, 1998).
- Lewis, David. 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy*, 61 (1983) 343–377. Reprinted in Lewis (1999).

- Lewis, David. *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999).
- Linnebo, Øystein. ‘Reference by Abstraction’. *Proceedings of the Aristotelian Society*, 112 (2012) 45–71.
- Moltmann, Friederike. ‘Reference to Numbers in Natural Language’. *Philosophical Studies*, 162 (2013) 499–536.
- Rosen, Gideon. ‘The Refutation of Nominalism (?)’. *Philosophical Topics*, 21 (1993) 149–186.
- Sider, Theodore. *Writing the Book of the World* (Oxford and New York: Oxford University Press, 2012).
- Turner, Jason. ‘Ontological Pluralism’. *Journal of Philosophy*, 107 (2010) 5–34.
- Williamson, Timothy. ‘Everything’. (2003) 415–465. In Hawthorne, J. and D. Zimmerman, eds. *Philosophical Perspectives 17: Language and Philosophical Linguistics*, Blackwell, Oxford.
- Wright, Crispin. *Frege’s Conception of Numbers as Objects* (Aberdeen: Aberdeen University Press, 1983).