

The University of Virginia
 Charlottesville, VA22903, USA
 Merricks@virginia.edu

References

- Armstrong, D. M. 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Bigelow, J. 1988. *The Reality of Numbers: A Physicalist's Philosophy of Mathematics*. Oxford: Clarendon Press.
- Black, M. 1952. The identity of indiscernibles. *Mind* 61: 153–64.
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis* 23: 121–23.
- Ladyman, J. and D. Ross. 2009. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Merricks, T. 2001. *Objects and Persons*. Oxford: Clarendon Press.
- Merricks, T. 2007. *Truth and Ontology*. Oxford: Clarendon Press.
- Williams, B. 1973. A critique of utilitarianism. In *Utilitarianism: For and Against*, eds. J.J.C. Smart and B. Williams, 77–150. Cambridge: Cambridge University Press.

Metaphysical Semantics Meets Multiple Realizability

JONATHAN SCHAFFER

Sider's *Writing the Book of the World* is an outstanding work. In an earlier review (Schaffer forthcoming) I called it 'masterful' and 'required reading', while worrying about Sider's notion(s) of structure. Here I offer the further praise of further discussion, centred on a separate worry concerning whether Sider's *metaphysical semantics* can handle multiple realizability.

Metaphysical semantics is supposed to connect the non-fundamental to the fundamental in a distinctively *linguistic* way, explaining how non-fundamental truths can be grounded in fundamental facts, and so inducing a radically eliminative vision of the non-fundamental as mere talk. I wonder how the story goes when a single non-fundamental truth can be grounded in many different fundamental facts. For instance, the truth that Moore has hands can presumably be grounded in many different distributions of fields, arrangements of particles, vibrations of strings, or whatnot. I am asking how metaphysical semantics can accommodate this sort of multiple realizability.

Multiple realizability came to prominence in philosophy of mind in the 1970s, raising problems for reductive versions of physicalism. I think that it raises analogous problems for Sider's eliminative view of the non-fundamental. Just as one cannot identify the mental state type *pain* with a specific brain

state, so one cannot equate the non-fundamental truth that Moore has hands with a specific fundamental fact. And just as multiple realizability ultimately favours non-reductive treatments of mental state types, so I think it ultimately favours non-eliminative treatments of the non-fundamental truths.

Overview: In Section 1 I describe metaphysical semantics, and identify three conditions of adequacy. In Section 2 I propose three ways for metaphysical semantics to treat multiple realizability, and in Sections 3–5 I argue that none of these three ways can satisfy more than one adequacy condition. I conclude in Section 6 that the underlying problem lies with Sider's eliminativist attitude to the non-fundamental.

1. *Metaphysical semantics*

Sider's metaphysical semantics is said to connect the fundamental to the non-fundamental, and to support a radically eliminative vision of the non-fundamental as mere talk. I begin in exegetical mode by situating metaphysical semantics within the larger project of *Writing the Book of the World*. Overall, Sider is after an account of *the fundamental*, which he equates with the structural, and with the facts that are cast in perfectly natural terms. Thus Sider (2011: 5)¹ says: 'The heart of metaphysics is the question: what is the world ultimately, or fundamentally, like? And fundamentality is a matter of structure: the fundamental facts are those cast in terms that carve at the joints.'

Fundamental and non-fundamental are not merely correlative but indeed connected notions. For presumably the non-fundamental is not merely that which is not fundamental, but also that that which is *generated from* the fundamental. Accordingly Sider (105) notes that 'the fundamental must in some sense be responsible for everything', and he (106) adopts the following working heuristic (which his metaphysical semantics is designed to make precise): 'Every non-fundamental truth is grounded by some fundamental truth.' For instance, supposing that it is a non-fundamental truth that Moore has hands, and given that our world is fundamentally one in which a palette of physical fields are painted over a spatiotemporal manifold, one should regard the truth that Moore has hands as grounded in facts about how fields colour spacetime.²

Metaphysical semantics is Sider's account of the connection between the fundamental and the non-fundamental. Sider rejects positing a metaphysical relation of grounding to connect fundamental and non-fundamental facts (or

1 Unless otherwise stated, all references are to Sider's *Writing the Book of the World*.

2 I am following Sider's lead in contrasting non-fundamental *truths* with fundamental *facts*. This way of speaking is usefully neutral as to whether the non-fundamental is also a realm of facts (or entities of any sort), thus leaving room for Sider's eliminativism about the non-fundamental.

entities of other sorts).³ Instead Sider (112) prefers a ‘linguistic approach’, saying: ‘The suitable notion, in my view, is that of a *metaphysical semantics* – a certain sort of semantic theory.’ On this linguistic approach there are no non-fundamental facts or entities of any sort (and *a fortiori* nothing to slot into the second place of a metaphysical relation of grounding). There are only *non-fundamental truths*, to be explained in fundamental terms. This conception of the non-fundamental as mere talk is what leads Sider towards the radically eliminativist worldview he sketches in his conclusion (Ch. 13), on which – by the lights of the joint-carving quantifier which we ought to adopt – only spacetime points and sets exist. As Sider (*forthcoming*: §3) puts the point: ‘in the fundamental sense of ‘exist’, there simply do not exist any composite entities.’

Metaphysical semantics, as an account of the grounding connection, is supposed (113) to issue theorems with something like the following schematic form:

TS: Sentence *S* in language *L* is true iff ϕ [with ϕ cast in purely fundamental terms]

TS thus connects the truth of sentence types (including non-fundamental ones) to fundamental facts. The format of *TS* could be modified in various incidental ways, while preserving this connection.⁴ I will work with *TS* as stated, since none of the concerns I raise turn on these details. (I am concerned about what goes in for ‘ ϕ ’ in the right biconjunct when there is multiple realizability.)

Sider then subjects metaphysical semantics to at least three conditions of adequacy. First, metaphysical semantics must fit with intuitive usage in most conceivable circumstances, to succeed as a metaphysical *semantics*. In terms

3 Sider rejects grounding by asking whether the associated connection facts are fundamental or non-fundamental. He rejects fundamental connection facts for violating his purity constraint (106: ‘fundamental truths involve only fundamental notions’), and he (124–27) charges non-fundamental connection facts with being unable to properly distinguish factual from non-factual discourse. For Sider (112), claims about what grounds what are just more non-fundamental talk: ‘We should instead take facts about the connection to be non-fundamental facts, which may ultimately be explained in terms of fundamental facts that do not involve a fundamental notion of connection.’

4 For instance, the left biconjunct should be relativized to contexts to handle context-sensitive discourse. Also the truth predicate might be relativized to \langle world, time \rangle pairs or whatever the truth predicate needs to be relativized to. (The left biconjunct might then look as follows: ‘Sentence *S* in language *L* in context *C* is true at $\langle w, t \rangle$ ’) Also, as Sider (113–14) notes, the left biconjunct can be generalized by replacing ‘true’ with some notion of ‘correct’ that covers truth-apt discourse as well as whatever other forms of discourse there might be (e.g. expressive adequacy for ethical discourse, provability for mathematical discourse, etc.) But the core idea is that one sketches something like truth-conditions for the things we say in purely fundamental terms, thereby explaining how the various things we say can be true (or false) in a world that is fundamentally just fields on spacetime (or particles in the void, etc.).

of *TS*, the biconditionals issued must hold in most conceivable circumstances, as per:

Fitting: The biconditionals must fit with intuitive usage by holding in most conceivable circumstances

After all, metaphysical semantics is in the game of ‘giving meanings in fundamental terms’, and of showing ‘*how what we say fits into fundamental reality*’ (112, emphasis added). An approach that failed to fit with intuitive usage in most conceivable circumstances could hardly claim to characterize the meaning of what we say.⁵

Secondly, metaphysical semantics must characterize intuitive usage in purely fundamental terms, if it is to succeed as a *metaphysical* semantics. In terms of *TS*, the right bijnuncts must be cast in purely fundamental terms, as per:

Sparse: The right bijnunct must connect to the fundamental by being cast in purely fundamental terms

After all, metaphysical semantics is in the game ‘giving meanings in *fundamental terms*’ and of showing ‘*how what we say fits into fundamental reality*’ of (112, emphasis added). An approach that failed to give meanings in fundamental terms could hardly claim to connect to the fundamental.

Thirdly, metaphysical semantics must be *explanatory*, if it is to connect the non-fundamental to the fundamental in the appropriate way. In terms of *TS*, the biconditionals issued must explain what we say, as per:

Explanatory: The biconditionals must explain our linguistic behaviour

After all, the connection between the non-fundamental and the fundamental is an explanatory connection. It is *because* the facts about the physical fields are as they are, that the truths about hands are what they are. Thus Sider (114) says that the account of intuitive usage in fundamental terms ‘must successfully explain the linguistic behaviour of the population in question’.

5 It is unclear to me exactly how much fit with intuitive usage Sider requires. In distinguishing metaphysical semantics from linguistic semantics, he (112–13) notes that metaphysical semantics does not attempt integration with syntax or psychology, says that competent speakers may not understand the metaphysical-semantic values assigned to their expressions, and allows that metaphysical semantics may ignore Frege puzzles. This suggests that some imperfect mesh can be tolerated around the edges. But in introducing the idea of a *toy model*, Sider (117) says that one respect in which these models can be toy is that ‘they needn’t match with intuitively correct usage in absolutely all possible worlds or conceivable circumstances. The mesh need only be approximate (the more mesh, the better)’. This suggests that a *real* metaphysical semantics must perfectly mesh with intuitively correct usage in all conceivable circumstances (imperfect mesh is for toys). But in any case Sider clearly requires at least decent mesh with intuitive usage across conceivable circumstances, as is only appropriate for a semantic proposal that aims to explain what we mean.

An approach that failed to be explanatory could hardly claim to *ground* the non-fundamental in the fundamental.

I agree with Sider that these three constraints – *Fitting*, *Sparse*, and *Explanatory* – are apt constraints given the purpose of metaphysical semantics. The question I am asking is whether metaphysical semantics can satisfy these constraints, given multiple realizability.

Metaphysical semantics can claim various advantages, beyond its primary role of making sense of manifest truisms (e.g. Moorean claims about hands). First, it fits a radically eliminative worldview, with the familiar draw of desert landscapes. Secondly, it allows Sider (116) to characterize the sense in which the fundamental is ‘complete’ and ‘responsible for everything’, via: ‘Every sentence that contains expressions that do not carve at the joints has a metaphysical semantics.’ Thirdly, it takes on the task (116–18) of replacing conceptual analysis. Instead of seeking a conceptual analysis of a notion such as *cause* – which Sider regards as a hopeless and unmotivated task – Sider recommends seeking a metaphysical semantics for causal claims (or at least seeking ‘toy models’ thereof, sufficient to convince us that a real metaphysical semantics exist, complicated though they may be). Such is the methodology of future metaphysics as Sider envisions it.

Putting this together: For Sider, metaphysical semantics is to issue theorems of the form *TS*, which must satisfy the constraints of *Fitting*, *Sparse*, and *Explanatory*. Metaphysical semantics thereby connects the non-fundamental truths to the fundamental facts, showing how what we ordinarily say fits how things fundamentally are, and so supporting an eliminative take on the non-fundamental as mere talk.

2. Multiple realizability

So far I have described Sider’s metaphysical semantics, and identified three conditions for adequacy. It is time for metaphysical semantics to meet multiple realizability. For a single non-fundamental truth can in many cases be grounded in many different fundamental facts. For instance, the truth that Moore has hands can be grounded in many different distributions of fundamental fields, arrangements of point particles, vibrations of strings, or what-not. What might a metaphysical semantics look like – and in particular what might go in for ‘ ϕ ’ in the right biconjunct of *TS* – which accommodates such multiple realizability?

It may help to have a working example. So consider the truth that Moore has hands, and continue to imagine that our world is fundamentally one in which a palette of physical fields are painted over a spatiotemporal manifold. Then by *TS* one needs to fill in the following blank, using only fundamental field-theoretic terms:

TSM: ‘Moore has hands’ in English is true iff —

What goes in the blank?

This might seem like a further exercise in exegesis, since Sider (118–21) himself works through some examples, providing a metaphysical semantics for ‘There exists an atom of hydrogen’ and ‘There exists an atom of helium’ in terms of a toy particle physics. But unfortunately Sider’s examples have two special features that together preclude issues of multiple realizability from arising. First, Sider’s examples involve *natural kind terms* (‘hydrogen’ and ‘helium’) which plausibly are semantically tied to whatever their actual fundamental basis might be. Natural kind terms have the special feature that they cannot be multiply realized via alien kinds of fundamental bases. Secondly, the toy particle physics that Sider uses for his illustrative examples permits only one way to build an atom of hydrogen, and only one way to build an atom of helium. (Siderian chemistry knows no isotopes!) His examples have the special feature that the atoms cannot be multiply realized via actual kinds of fundamental bases.

It is thus not obvious what Sider would insert for ‘ ϕ ’ in cases where there is multiple realizability, whether via alien kinds of fundamental bases or multiple actual kinds of fundamental bases. In that sense it is not obvious what metaphysical semantics says.⁶ I am assuming that the truth that Moore has hands is such a case. (If you think otherwise change the example.) Overall I am asking:

Question: What goes in for ‘ ϕ ’ in the right biconjunct of *TS* when there is multiple realizability?

More generally, when a given truth can be grounded in different sorts of fundamental situations, what form might a metaphysical semantics take? (I fear that there is no answer to *Question* that satisfies the triple constraints of *Fitting*, *Sparse*, and *Explanatory*.)

Since what goes in for ‘ ϕ ’ in the right biconjunct of *TS* (e.g. what fills in the blank in *TSM*) can only be cast in fundamental terms, there seem to me to be three main options to consider. A first and perhaps most obvious option is simply to insert the actual fundamental realizer (in the contextually specified situation at issue). With *TSM* this would be to insert the exact actual field values that ground the truth that Moore has hands (in the relevant situation). But a second more liberal option would be to insert the disjunction of potential actual fundamental realizers. With *TSM* this would be to insert the disjunction of potential exact values of the actual fields that could ground the truth that Moore has hands in circumstances fundamentally akin to ours. And a third even more liberal option would be to insert the disjunction of

6 I thus agree with Goswick (2013: 204), who says that Sider ‘says far less than necessary about metaphysical semantics if we’re to understand how it connects our everyday concepts with joint-carving ones’. This paper should be read as an extended invitation for Sider to say more.

possible fundamental realizers. With *TSM* this would be to insert the disjunction of all possible fundamental bases (not just actual field values, but values of alien fields, as well as particle arrangements, string vibrations, stirrings in the mind of Berkeley's deity, etc.) that could possibly ground the truth that Moore has hands in any conceivable circumstance whatsoever. Schematically, these three main answers to *Question* may be phrased as follows:

- *Actual realizer*: $\phi =$ The actual fundamentals are thus
- *Potential actual realizers*: $\phi =$ The actual fundamentals are thus or such or so or ...
- *Possible realizers*: $\phi =$ The actual fundamentals are thus or such or so or ... , or those possible fundamentals are thus or such or so or ... , or these other possible fundamentals are thus or such or so or ... , or ...

Of course I cannot rule out that there is some fourth answer to *Question* which I have not imagined. The friend of metaphysical semantics is invited to view this as an invitation to exhibit some other answer, and then show how it can satisfy the triple constraints of *Fitting*, *Sparse*, and *Explanatory*.

There is one further answer worth mentioning (suggested by Sider's remarks on 130, and also by Sider, personal communication), if only to show that it cannot really be distinct from one of the three answers just mentioned. This further answer involves a 'two step' strategy, whose first step involves quantification over role properties, as illustrated by:

Role Moore: 'Moore has hands' in English is true iff Moore has properties that play the hands role

Or perhaps the right bijunct should say: 'there is something that plays the Moore role and that has properties that play the hands role.' Or: 'there are some things that collectively play the Moore role and collectively have properties that play the hands role.' It won't matter.

Role Moore does not look to realizers of the hands role, but to the role itself. It is worth noting that *Role Moore* is itself not implausible, and is indeed the sort of thing that the believer in non-fundamental entities such as myself can go in for (though I would simplify the right bijunct with the disquotational rendering: '... iff Moore has hands'). The problem of course – and the reason why *Role Moore* can only be a first step – is that 'properties that play the hands role' is evidently a non-fundamental notion. So the second step is to get down to fundamentals by giving a proper metaphysical semantics for the right-hand side of *Role Moore*:

Realizer Moore: 'Moore has properties that play the hands role' in English is true iff ϕ [with ϕ cast in purely fundamental terms]

But multiple realizability re-arises: the truth that Moore has properties that play the hands role is itself multiply realizable. The problem has only

been deferred. And so it seems that one is back to the three options bulleted above. To reach the fundamental, one must eventually descend from role to realizer. I am trying to understand how the right bijunct of *TS* might look at this stage.

3. *First answer: actual realizer*

I have offered Sider three possible ways to handle multiple realizability in fundamental terms (Section 2). But can any of these yield a metaphysical semantics that can satisfy the triple constraints of *Fitting*, *Sparse*, and *Explanatory* (Section 1)? I begin with the first answer – *Actual realizer* – as illustrated by:

TSM1: ‘Moore has hands’ in English is true iff the actual field values are thus

In *TSM1*, the right bijunct should be read as shorthand for the exact and highly complex actual field values that in fact ground the truth that Moore has hands (in the contextually specified situation at issue).⁷

The good news: *TSM1* satisfies *Sparse*. On the working assumption that our world is fundamentally one in which a palette of physical fields are painted over a spatiotemporal manifold, the right bijunct of *TSM1* is indeed cast in purely fundamental field-theoretic terms.

The bad news: *TSM1* satisfies neither *Fitting* nor *Explanatory*. As to *Fitting*, *TSM1* renders ‘Moore has hands’ in English false in any circumstance whatsoever that does not involve these exact actual field values, and so clearly does not fit with intuitive usage in conceivable circumstances in which this non-fundamental truth is grounded in any other way. (This is a multiple realizability problem.) These circumstances include those in which the truth that Moore has hands is grounded in slightly different actual field values, as well as those in which the truth that Moore has hands is grounded in alien fundamental bases. However things fundamentally are at our world is just a ‘drop in the ocean’ of the many and various conceivable ways things might fundamentally be. The truth that Moore has hands might, for instance, be grounded in various sorts of fields painted over spacetime, point particles swirling in the void, strings vibrating through the ether, or even ideas stirring in the mind of Berkeley’s deity.

To illustrate the failure of *Fitting*, continue to imagine that our world is field-theoretic, and imagine an ‘alien’ particles scenario in which the particles happen to be arranged handwise. Evidently, intuitive usage in this particles-arranged-handwise scenario would still uphold ‘Moore has hands’ in English

7 It helps to read ‘Moore has hands’ as concerning a single specific situation. How *Actual realizer* might be extended to quantified or generic claims is a matter I leave aside, since *Actual realizer* proves manifestly inadequate anyway.

as remaining true. But by *TSM1*, ‘Moore has hands’ in English is true iff the actual field values are thus. At a particles scenario no distribution of actual field values obtains. Fundamental reality would have a very different character, and the fundamental facts would need to be cast in entirely distinct terms. So by *TSM1*, ‘Moore has hands’ in English turns out *false* in any particles scenario whatsoever, no matter how the particles are arranged (whether arranged handwise or not). *TSM1* thus systematically departs from intuitive usage in all but the small drop in the ocean of conceivable circumstances that happen to be fundamentally exact duplicates of ours in the situation at issue.

Turning to *Explanatory*, *TSM1* does not explain what we say since it is overly specific, includes irrelevant details, and misses the generalizations. This is a point familiar from the multiple realizability literature, perhaps most memorably from Putnam’s (1975: 295–97) discussion of why a square peg cannot fit through a round hole. Putnam argues that the reductive microphysical story is not explanatory. An apt explanation – which includes only the relevant details, and captures the underlying generalizations – needs to be couched in terms of the higher level geometric structure:

[I]f you are not ‘hipped’ on the idea that *the* explanation must be at the level of the ultimate constituents, and that in fact the explanation might have the property that *the ultimate constituents don’t matter*, that *only the higher level structure matters*, then there is a very simple explanation here. The explanation is that the board is rigid, the peg is rigid, and as a matter of geometrical fact, the round hole is smaller than the peg, the square hole is bigger than the cross-section of the peg. The peg passes through the hole that is large enough to take its cross-section, and does not pass through the hole that is too small to take its cross-section. That is a correct explanation whether the peg consists of molecules, or continuous rigid substance, or whatever.

Likewise, when it comes to explaining linguistic behaviour, a connection drawn at the type level between the truth of a sentence in a language and an exact fundamental situation is overly specific and laden with irrelevant details. An apt explanation for why English speakers say ‘Moore has hands’ involves something like the presence of *hands*. This is an apt level of explanation, with the appropriate specificity, relevance of information and amenability to generalization. Descending to a single actual realizer – as per *Actual realizer* – is a step in the wrong direction, as far as explanation is concerned.⁸

8 The failure of *Actual realizer* to satisfy *Explanatory* is partly due to the fact that TS draws a *type level* connection between sentential truth and the fundamental. I would not object to the idea that a given token of ‘Moore has hands’ might be true in virtue of its actual realizer in the situation at issue. (The dialectic here is somewhat analogous to the way in which multiple realizability draw people away from type physicalism while still accepting token physicalism.) See footnote 12 for some further discussion.

4. *Second answer: potential actual realizers*

The first strategy for handling multiple realizability – via *Actual realizer* – looks to fail both *Fitting* and *Explanatory*. So I turn to the second strategy, involving the answer of *Potential actual realizers*, which disjoins the potential actual fundamental realizers in its right biconditional. This may be illustrated by:

TSM2: ‘Moore has hands’ in English is true iff the actual field values are thus or such or so or . . .

In TSM2, the right biconditional should be read as shorthand for the disjunction of the various exact and highly complex actual field values that could ground the truth that Moore has hands (on the continued working assumption that our world is field-theoretic).

It is not obvious that TSM2 satisfies *Sparse*. After all, TSM2 includes disjunctions on the right, and it is a substantive claim that disjunction is a fundamental term. Though Sider (218–19) ultimately endorses this claim about disjunction, he does so tentatively, through substantive considerations about the logical connectives. I doubt he would want metaphysical semantics to presuppose fundamental disjunction from the start.

But leave that aside: TSM2 does no better than TSM1 with respect to *Fitting* or *Explanatory*, for essentially the same reasons. (Little progress has been made.) As to *Fitting*, while TSM2 at least manages to capture intuitive usage within the small droplet of scenarios with actual kinds of fundamental bases (e.g. scenarios that retain the same fundamental fields as ours), it does no better in capturing intuitive usage across the vast ocean of scenarios with alien kinds of fundamental bases. It delivers a blanket verdict of ‘false’ for ‘Moore has hands’ in English in all scenarios with alien fundamentals, regardless of how these alien fundamentals are (e.g. even in scenarios in which the particles are arranged handwise). There remains a systematic departure from ordinary usage in a wide range of conceivable circumstances.

As to *Explanatory*, TSM2 remains as laden with overly specific and irrelevant fundamental details as TSM1, and – to make matters worse – adds on an open-ended infinitary disjunctivity. Presumably there are infinitely many field distributions that could ground the truth that Moore has hands, which are a disunified lot from the perspective of the fundamental physics. Explanations are supposed to unify the phenomena. Because the right biconditional of TSM2 contains an open-ended infinitary disjunction, the biconditional issued fails to unify our behaviour.

5. *Third answer: possible realizers*

The first two strategies for handling multiple realizability – via *Actual realizer* or *Potential actual realizers* – each looks to fail both *Fitting* and *Explanatory*.

So I turn to the third and final strategy, involving the answer of *Possible realizers*, which disjoins all possible realizers in its right bijnunct. This may be illustrated by:

TSM3: ‘Moore has hands’ in English is true iff the actual field values are thus or such or so, or the particles are arranged thus-wise or such-wise or so-wise, or the strings are vibrating as thus or as such or as so, or Berkeley’s deity is dreaming of thus or of such or of so, or . . .

In TSM3, the right bijnunct should be read as shorthand for the disjunction of all possible fundamental bases (not just actual field values, but values of alien fields, as well as particle arrangements, string vibrations, etc.) that could potentially ground the truth that Moore has hands.

With TSM3 there is finally an answer that approaches *Fitting*. For TSM3 can match intuitive usage in conceivable circumstances both with actual kinds of fundamental bases (e.g. the various ways in which the actual fundamental fields might ground the truth that Moore has hands) and with alien kinds of fundamental bases (e.g. scenarios in which particles are arranged handwise, strings are vibrating handfully, or Berkeley’s deity is dreaming handishly). That said, it is not yet obvious that TSM3 fully satisfies *Fitting*, for there are also conceivable circumstances in which nothing is fundamental, but instead there is a limitless descent of every deeper structure. It is conceivable that some of these limitlessly descending scenarios still ground the truth that Moore has hands. (Gunk may be arranged handwise.) Indeed it is at least epistemically possible that our world is such a scenario (Sider 1993; Schaffer 2003).

But leave that aside: perhaps match with intuitive usage in conceivable circumstances with a fundamental level is match enough for *Fitting*, or perhaps circumstances without a fundamental level are for some subtle reason only superficially conceivable, or perhaps circumstances without a fundamental level should be ignored for some other reason.⁹ In any case, TSM3 satisfies neither *Sparse* nor *Explanatory*. With respect to *Sparse*, leaving aside any worries about including disjunction on the right (which TSM3 inherits from TSM2), the real worry comes from the use of terms for merely conceivably fundamental properties. (This is the flipside of TSM3’s ability to satisfy *Fitting*.) The use of such terms obviously requires some vast infinity of terms: one for every property that is conceivably fundamental. But the terms in the right bijnunct – by *Sparse* – are required to be fundamental joint-carving terms. Assuming that our world is fundamentally field-theoretic, the terms

9 Sider discusses the possibility of limitless ideological descent (133–36), speculates that it is metaphysically impossible, but still grants that it is epistemically possible. I take it that merely epistemic possibilities still count as *conceivable* circumstances with respect to *Fitting*, but will not insist on the point.

of particle theory, string theory, Berkeleyan theology, etc. fail to carve at the joints.¹⁰

Indeed, *TSM3* would require that there be infinitely many joint-carving terms, which conflicts with Sider's (14) guiding motivation of *ideological parsimony*. The use of merely possible fundamental terms would require a veritable ideological explosion, one on which there are as many joints in nature as there are predicates that are even *conceivably* fundamental.¹¹

What emerges is something of an underlying tension between *Fitting* and *Sparse*. In order to satisfy *Sparse*, a metaphysical semantics must restrict what can appear in the right biconjunct of *TS* to actual fundamental terms. But in order to satisfy *Fitting*, a metaphysical semantics must fit intuitive usage in conceivable circumstances with alien fundamentals. But how can a metaphysical semantics fit intuitive usage in conceivable circumstances with alien fundamentals, if its truth-conditions can only be cast in terms of actual fundamentals?

Anyone familiar with linguistic semantics will be aware that this enterprise is pursued in unapologetically possibilist terms. The tension emerging for metaphysical semantics arises from considering whether it should follow suit. Without terms for merely conceivably fundamental properties metaphysical semantics might satisfy *Sparse* but cannot satisfy *Fitting*; with such terms metaphysical semantics might satisfy *Fitting* but cannot satisfy *Sparse*.¹²

As to *Explanatory*, *TSM3* fares no better than *TSM2*. It remains laden with overly specific and irrelevant details about the fundamentals, and it retains an open-ended infinitary disjunction on the right.

10 Sider's overall usage clearly indicates that the 'joint-carving terms' are only those that *actually* carve at the joints, not merely those that conceivably might carve at the joints. For instance, Sider (113) – clarifying that competent speakers need not understand the metaphysical-semantic values assigned to their expressions – says that metaphysical semantics might 'assign to an ordinary sentence about ordinary macroscopic objects a meaning that makes reference to the fundamental physical states of subatomic particles'. And in his concluding worldview, Sider (292) includes 'predicates adequate for fundamental physics' but there is no mention of predicates for rival but contingently false fundamental physical theories.

11 Sider (201) notes the following objection to Lewis's modal realism: '[S]ince Lewis's pluriverse contains myriad 'alien' natural properties and relations, modal realism sins against Ockham's razor in a big way.' On the *Possible realizers* approach, metaphysical semantics would turn out to be just as sinful as modal realism.

12 Sider (118) speaks of his proposal as following in the Armstrongian spirit of invoking 'a kind of purely metaphysical analysis'. But most articulations of metaphysical analysis proceed in possibilist terms, if only because it is hard to see how a claim could be in any sense 'analytic' but yet so narrowly contingently restricted. For instance, Dorr (2005: 261–62) invokes a notion of metaphysical analysis built around locutions like 'to be a square is to be a quadrilateral with equal sides and angles', explaining (2005: 263) that a metaphysically analytic sentence 'expresses the same fact as a logical truth', just in a more perspicuous (and not always transparent) way.

What emerges is something of an underlying moral about *Explanatory*. At least, for anyone moved by Putnam's discussion of what explains (and what does not explain) why a square peg cannot fit through a round hole, an apt explanation for our linguistic behaviour is going to look more like *Role Moore*, which explains our usage of 'Moore has hands' in English via the presence of something with properties that play the hands role (or better: via the presence of hands). This explanation avoids overly specific and irrelevant details, captures the relevant generalization, and unifies the phenomena precisely by staying at the derivative level where talk of hands remains available. As Putnam (1975: 296) puts the point: 'The same explanation will go in any world (whatever the microstructure) in which those *higher level structural features* are present. In that sense *this explanation is autonomous*.'

How can a metaphysical semantics satisfy *Explanatory*, if a metaphysical semantics is supposed to connect our non-fundamental talk to the fundamental, but apt explanations for our non-fundamental talk are autonomous higher level explanations? Anyone familiar with linguistic semantics will also be aware that this enterprise is pursued in unapologetically non-fundamental terms. The problem for metaphysical semantics arises from considering how it can reject such terms yet retain the relevance and generality of real semantic explanation.¹³

Putting these considerations, none of the three answers to *Question* can satisfy more than one of the three adequacy conditions of *Fitting*, *Sparse*, and *Explanatory*. The overall scoreboard reads:

Answer	Score
<i>Actual realizer</i>	Satisfies <i>Sparse</i> but violates <i>Fitting and Explanatory</i>
<i>Potential actual realizers</i>	Satisfies <i>Sparse</i> (assuming that disjunction is fundamental) but violates <i>Fitting and Explanatory</i>
<i>Possible realizers</i>	Satisfies <i>Fitting</i> (ignoring conceivable circumstances of limitless descent) but violates <i>Sparse</i> and <i>Explanatory</i>

Two underlying problems have emerged:

- If a metaphysical semantics invokes only actual fundamental terms then it violates *Fitting*, but if it invokes merely conceivably fundamental terms then it violates *Sparse*

13 Point of clarification: I do not mean to deny that there is an explanatory path from the non-fundamental to the fundamental. My own view (ignoring the contrastivity of explanation) is that one explains why it is true that Moore has hands in terms of Moore's having hands. This is a semantic level explanation, which satisfies *Fitting* and *Explanatory* but makes no effort to satisfy *Sparse*. One can then explain why Moore has hands in a given token situation in terms of the actual realizer in that situation, e.g. the distribution of fields on spacetime. This is a metaphysical explanation which satisfies *Sparse* but makes no effort to satisfy *Fitting* or *Explanatory*.

- An apt explanation of our linguistic behaviour should be an autonomous higher-level explanation

At the root of both of these problems lies multiple realizability. It is multiple realizability that explains why both *Actual realizer* and *Potential actual realizers* violate *Fitting*, and why all three proposals violate *Explanatory*. Just as we cannot identify the mental state type *pain* with any specific brain state, so we cannot equate the non-fundamental truth that Moore has hands with any specific fundamental fact.

6. Concluding remarks

I have asked how metaphysical semantics can accommodate multiple realizability, and argued that the three most natural answers are inadequate. The friend of metaphysical semantics is invited to view this as an invitation to do better, by displaying a treatment that can satisfy *Fitting*, *Sparse*, and *Explanatory*.

I will consider one last move – suggested by Sider (*personal communication*) – which is to go pluralist about the ‘correct’ metaphysical semantics. After all, for Sider (115), metaphysical semantics is itself ‘a high-level notion . . . posited in the course of explaining a high-level phenomenon, the phenomenon of language use by flesh and blood people’. And thus: ‘We should take a non-fundamental, no doubt vague, and perhaps not even uniquely correct, approach to the question of the relationship between the fundamental and the non-fundamental.’ So perhaps Sider can recommend a plurality of approaches including *Actual realizer*, *Potential actual realizers*, *Possible realizers* and even *Role Moore*, as being decent enough for different purposes. He need not make a once-and-for-all choice.

I am not convinced that pluralism is a stable position for Sider, but in any case I do not see how it can help.¹⁴ If this were a situation in which one was being presented with some individually decent rivals that could each perform in complementary roles, then I would see how pluralism could help. But multiple realizability is not presenting one with a situation with some decent rivals that can each be embraced as partially viable approaches useful in different explanatory contexts. Rather multiple realizability is presenting one with a situation with no viable approach whatsoever. No form of metaphysical semantics – at least from among the three considered – has been able to satisfy more than one out of Sider’s own three adequacy constraints.

14 Is pluralism a stable position for Sider? Recall that the fundamental and the non-fundamental are not just correlative but connected notions, and that Sider uses metaphysical semantics to characterize the completeness requirement on the fundamental (Section 1). If the non-fundamental is a vague realm permitting a plurality of partially acceptable theories, how does the fundamental avoid being infected by the vagaries of its correlate, and the pluralism within its completeness requirement?

In general, pluralism may help when an approach permits many decent complementary articulations; it cannot help when an approach permits none.

If metaphysical semantics cannot accommodate multiple realizability, then Sider's linguistic account of the connection between the fundamental and the non-fundamental would fail. Further consequences would include a failure of Sider's account of the fundamental in terms of completeness (which he understands in terms of metaphysical semantics), and a failure of Sider's preferred methodology for doing metaphysics (for which he proposes toy models of metaphysical semantics).

The underlying problem, by my lights, lies with Sider's eliminativist attitude to the non-fundamental. If Sider would countenance non-fundamental entities (plus a metaphysical relation of grounding to connect these entities to the fundamental), he could then offer a unified, relevant and general semantic explanation for why we tend to say things like 'Moore has hands' in English. The explanation is: *because Moore has hands*. In this sort of explanation a non-fundamental fact appears on the right. This allows relevant explanatory features to appear, and unifying generalizations to be captured.

Moreover, it is by appealing to non-fundamental entities that one can say what 'Moore has hands' and 'Moore has feet' have in common (both concern *Moore*), and what 'Moore has hands' and 'Russell has hands' have in common (both concern *hands*). There are also more general patterns in the explanations of what we say that are lost if one looks only to the fundamental realizers (be they actual or possible). Non-fundamental entities are explanatorily fruitful posits.

Sider's *Writing the Book of the World* mixes an eliminativist ontology with a sparse ideology, painting a radically desertified landscape. The image is haunting. (And the reasoning is brilliant.) But I fear that there is just not enough left to ground the non-fundamental. Deserts may be beautiful, but they can also be barren.¹⁵

Rutgers University
New Brunswick, NJ 08901-1411
jonathan.schaffer@rutgers.edu

References

Dorr, C. 2005. What we disagree about when we disagree about ontology. In *Fictionalism in Metaphysics*, ed. M. E. Kalderon, 234–86. Oxford: Clarendon Press.

15 Thanks to Marco Dees, Janelle Derstine, Jeffrey King, Michaela McSweeney, Ted Sider, Jennifer Wang and Tobias Wilsch, and to audiences at Rutgers University, Syracuse University and Western Michigan University, where this material was included in a larger talk entitled 'New Work for the Nonfundamental'.

- Goswick, D. 2013. Review of *Writing the Book of the World*. *Australasian Journal of Philosophy* 91: 202–05.
- Putnam, H. 1975. Philosophy and our mental life. *Mind, Language, and Reality: Philosophical Papers* 2: 291–303.
- Schaffer, J. 2003. Is there a fundamental level? *Nous* 37: 498–517.
- Schaffer, J. forthcoming. Review of Theodore Sider's *Writing the Book of the World*. *Philosophical Review*.
- Sider, T. 1993. Van Inwagen and the possibility of gunk. *Analysis* 53: 285–89.
- Sider, T. 2011. *Writing the Book of the World*. Oxford: Clarendon Press.
- Sider, T. forthcoming. Against parthood. In *Oxford Studies in Metaphysics*, vol. 8, eds. K. Bennett and D. Zimmerman, Oxford: Oxford University Press.

*Symposium on Writing the Book of the World**

THEODORE SIDER

1. Reply to Gabriele Contessa

Gabriele Contessa concludes his fine paper by saying that although he rejects my strong form of realism about structure, he accepts a weak form:

Strong structure realism is what I take Sider to be advocating in *Writing the Book of the World* – it is the view that we do not just need to be realist about what the individual components of the world's fundamental structure are (which might include, depending on one's metaphysical views, universals or bare particulars or what-have-you) but also about the world's structure itself. *Weak structure realism*, on the other hand, is simply the thesis that the world has one structure (as opposed to none or many), a thesis whose denial leads to various forms of metaphysical pluralism or metaphysical anti-realism. Weak structure realism takes 'structure' to be just a placeholder for whatever the world contains at the fundamental metaphysical level (as opposed to something one can be genuinely realist about).

But in fact, with one exception, Contessa and I disagree on little.

For instance, Contessa's definitions above of *both* weak and strong structure realism use the terms 'structure' and 'fundamental', apparently presupposing that such terms are in good standing. But establishing this was in effect the main point of my book. Further, the strong structure realism that Contessa attributes to me is distinguished by its realism 'about structure itself'; but I don't think of 'the world's structure' as some sort of object or fact

* Thanks to Andrew Cortens, Amie Thomasson, and my critics.