

Symposia on *Writing the Book of the World**

THEODORE SIDER

Analysis 73 (2013): 713–15, 751–70

PPR 87 (2013): 706–8, 733–54

Contents

1 Précis of <i>Writing the Book of the World</i>	I
2 Reply to Contessa: structure realism; nonpredicate structure	4
3 Reply to Dorr: epistemology of structure	8
4 Reply to Fine: fundamental terms vs fundamental truths	14
5 Reply to Hirsch: entity-free structure; quantifier variance	22
6 Reply to Merricks: nature of metaphysics; modality	30
7 Reply to Schaffer: metaphysical semantics and multiple realizability	40

1. Précis of *Writing the Book of the World*

Nelson Goodman (1978) said that there is no “ready-made world”. Human physicists speak of “charge”, “mass”, and “distance”, but nothing would be wrong with instead describing the world using cooked-up words that stand to the physicists’ words as Goodman’s ‘grue’ and ‘bleen’ stand to ‘green’ and ‘blue’.

My central thesis is that there *is* a ready-made world. Some concepts are objectively privileged: the “fundamental”, or “joint-carving”, or “structural” ones. A complete description of reality using these concepts—the “book of the world”—gives reality’s fundamental structure. A description using cooked-up concepts can be true, and even in a sense equivalent to the book of the world, but is nevertheless representationally deficient since its structure does not match the world’s structure.

*These are (unabridged versions of) my contributions to two symposia on my book *Writing the Book of the World*. The symposia appeared in *Analysis* (critics: Contessa, Merricks, Schaffer) and *Philosophy and Phenomenological Research* (critics: Dorr, Fine, Hirsch). Thanks to Karen Bennett, Andrew Cortens, and Amie Thomasson, and to my critics, for helpful feedback.

Many of the traditional questions of metaphysics are about the nature of reality's fundamental structure. Is reality "ultimately" just physical, or is there also a mental aspect? This is a question of whether the book of the world mentions mentality—of whether mental concepts are structural.

Realism about structure is akin to realism about natural kinds, particularly as developed by David Lewis (1983*a*, 1986*a*, pp. 59–69) in his theory of natural properties and relations. But structure differs from naturalness in important ways, especially in applying to nonpredicates: one can ask whether modal operators, quantifiers, and even sentential connectives like 'and' or 'not' are structural. This is important because the traditional metaphysical questions are not only about predicate-structure. Is reality ultimately amodal, or is the distinction between necessity and possibility written into the book of the world? This is the question of whether the modal sentential operators \Box and \Diamond carve at the joints. Does reality have a distinguished entity-structure, or are there—as Goodman thought—multiple equally good ways of carving reality into domains of objects? This is the question of whether quantifiers carve at the joints.

Certain questions about objectivity involve structure. Objectivity is often understood as mind-independence (in various senses). But even when we use language to express mind-independent facts, there may yet be a failure of objectivity in a further, and largely neglected, sense. For suppose that we (or our biology or circumstances) have chosen one of many candidate meanings for a word, where the unchosen candidates carve at the joints equally well and would have played a similar semantic role to our actual choice. Then even though we express mind-independent facts with the word, the facts we express are in a sense a projection of the choice of which candidate to mean. Even if 'Europe' is just a name of a certain physical region of the planet, with no historical or political dimension to its meaning, the choice of the particular region signified is a projection of historical and political facts. Thus there can be a certain failure of objectivity in our talk about Europe's borders, say, even if the facts we express are mind-independent. This is not the case with words from physics, for there are no equally joint-carving candidates for physical words that would have played similar semantic roles to our actual semantic choices.

This further way in which objectivity can fail is, I believe, essentially involved in the deepest critiques of contemporary metaphysics. For instance, Eli Hirsch (2011) has been arguing for years (in the spirit of Rudolf Carnap (1950) and Hilary Putnam (1987)) that debates over the ontology of composite and persisting objects (Chisholm versus Lewis versus van Inwagen...) are "merely

verbal” in the sense that the view of each debater comes out true under some possible meaning for the crucial expressions in the debate. Peter van Inwagen (1990) and Lewis (1986*a*, pp. 211–3) think that they substantively disagree over whether there exist tables and chairs exist (in addition to “particles arranged table-wise”). But according to Hirsch’s thesis of “quantifier variance”, quantifiers such as ‘there exist’ can be given multiple meanings. Under some, van Inwagen’s claims come out true; under others, Lewis’s claims come out true; and the only real issue is one of conceptual analysis: which meaning is the ordinary English one? Hirsch stresses that his quantifier variance is compatible with “realism”, but by that he means merely that the facts we express under any particular quantifier meaning are mind-independent. At a deeper level his position is anti-realist, since he denies that the quantificational facts are objective in the further, neglected sense. Quantificational facts on his view are a projection of our conceptual scheme, which is just one amongst many equally good possible schemes.

My own view is that the best hope for a sustainable defense of ontology requires embracing realism about quantificational structure. For if in addition to Hirsch’s multiple meanings for ‘there is’, there is *also* a joint-carving sense of ‘there is’, then there remains a substantive question about ontology: that of what exists in this joint-carving sense.

It can be argued that there is indeed a joint-carving sense of ‘there is’, and thus that ontology is substantive. But whether this is so is itself a substantive question, over which there is room for debate. In other cases, such as certain debates about causation, I myself would stand with the Carnapian side. There is arguably no joint-carving notion of causation—“reality is fundamentally acausal”; this can then be used to argue that certain debates over the nature of causation are merely verbal in something like Hirsch’s sense. But setting aside questions about particular cases, the larger issue is that many questions about meta-metaphysics are about objectivity in the further sense, and ultimately turn on where reality’s joints lie.

I understand the fundamental facts as those involving structural concepts, and I do not define ‘structure’ in modal terms. (In fact, I do not define it at all; structure is primitive—indeed, structural!) Thus I join many recent authors (such as Karen Bennett (2013), Kit Fine (2001, 2012), and Jonathan Schaffer (2009)) in refusing to understand notions in the vicinity of fundamentality in modal terms. But there are differences between me and these other authors. For instance, Fine’s notion of ground is a propositional and comparative notion: entire propositions ground and are grounded, and ground is a relation between

pairs of propositions. Structure on the other hand is subpropositional and absolute: it is parts of propositions (or parts of sentences, to put it linguistically) that are structural, and a single proposition-part, rather than a pair, is structural. These and other subtle differences turn out to matter in surprising ways.

2. 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”. I’m not sure what that means, but I certainly don’t think of “the world’s structure” as some sort of object or fact over and above facts about this or that concept carving at the joints—facts that Contessa apparently accepts.

We do genuinely disagree over applying structure to logical notions. Regarding my defense of Lewis’s “reference magnetism” reply to certain forms of semantic skepticism (section 3.2), Contessa says:

I doubt that reference magnetism is a cure-all for all forms of semantic underdetermination. In particular, I doubt that it can contribute to fixing the meanings of expressions such as ‘exist’ or ‘or’. Ultimately, this is because I cannot see what in the world could possibly act as a reference magnet for those words.

The last sentence makes clear that his concern is with my metaphysics rather than my metasemantics: he doubts my claim that one can speak of logical notions like ‘exist’ and ‘or’—and ‘not’, as he goes on to say—as carving at the joints.

However, his main point is elsewhere (thus he doesn’t address my defense of joint-carving in logic):

...even if I deny the existence of a reference magnet for ‘either ...or ...’ and ‘it is not the case that’, my denial does not necessarily make me a structure anti-realist. I might well be a structure realist who denies that any of the candidate interpretations of ‘either ...or ...’ and ‘it is not the case that’ compatible with the ordinary uses of those expressions carves nature at its joints better than any other. In fact, I might even be a structure realist who thinks that the very question of which interpretations of ‘either ...or ...’ and ‘it is not the case that’ carve nature at its joint is misguided, because such expressions are not meant to carve nature in the first place.

Contessa is certainly right to say that one could buy realism about structure in general without accepting that disjunction and negation carve at the joints. But consider the two reasons he gives. The second we have already discussed: one might reject the legitimacy of applying structure to logical connectives. But the first is different: one might simply hold that negation and disjunction do not in fact carve at the joints. But why is that a problem? My realism about structure is *supposed* to be neutral on what does or doesn’t carve at the joints. It isn’t meant to settle, on its own, particular questions about objectivity, substantivity, or semantic determinacy. Its point is rather to provide a background metaphysics in which theses about these matters may be articulated and defended. Whether there is objectivity (etc.) in any particular domain depends on the substantive question of whether the key concepts in that particular domain carve at the joints. (I myself answer such questions differently in different cases: I defend joints in nature, objectivity, and semantic determinacy in ontology (chapter 9) but not in modality (chapter 12).)

Consider next:¹

...structure (qua structure) is ultimately idle. What does the job in virtually every supposed application of structure is not the world’s structure as a whole but one or more of its components. The world’s structure, for

¹Side point: I don’t think reference magnetism answers Quine’s gavagai problem (see p. 30, n. 17).

example, does not act as a reference magnet for ‘green’, ‘gavagai’, ‘either ...or ...’, or ‘it is not the case that’. Rather, some of its components do, and structure realists can legitimately disagree with one another as to which specific components of the world’s structure (if any!) act as reference magnets for each of those expressions. So, structure realism, in and of itself, does not provide its supporters with the tools to defuse all forms of semantic scepticism; only realism about the specific components of the world’s structure that act as reference magnets for some term that can do so.

In addition to reiterating the point just discussed, this passage claims that since only the “components” of the world’s structure, and not “the world’s structure as a whole”, take part in structure-involving explanations, “structure (qua structure) is ultimately idle”. Now, it’s true that such explanations cite facts about particular concepts carving at the joints (and not just “the world’s structure”), but such explanations still essentially involve the concept of structure: one cannot delete the concept of joint-carving from the explanans “concept *C* carves at the joints”.² Suppose *a* is negatively charged and *b* is positively charged. If you want to explain why *a* and *b* attracted each other, you couldn’t, of course, just cite “positive and negative charge”; you’d need to cite the fact that *a* has negative charge and *b* has positive charge. But the appeal to charge in the explanation is still essential; you couldn’t just cite *a* and *b*.

I’ll close by mentioning one last concern of Contessa’s:

...this talk of the world’s structure is useful only insofar as ‘structure’ is taken to be a placeholder for whatever the world might turn out to contain (at a fundamental level). If that’s the case, then ‘structure’ is likely to turn out to be a term of convenience, a term that picks out a motley assortment of entities with little or nothing in common besides their being collectively referred to as ‘structure’. The referent of ‘structure’ might turn out to be a metaphysical hodgepodge of radically heterogeneous entities, as opposed to something that one can be genuinely realist about

²The issue comes to a head when Contessa speaks of the “components” of the world’s structure doing all the work in explanations. Are the components facts of the form “*C* carves at the joints”, or the individual concepts *C*? Only if they’re the former do they do all the work, but then structure hasn’t been shown to be explanatorily idle.

On pp. 94-96 I discuss (and reject) a view I call “ontologism”, which presupposes a “sparse ontology” and tries to do (some of) the work I do with ‘concept *C* carves at the joints’ instead with ‘concept *C* stands for some entity’ (compare Armstrong (1978*a,b*) on universals). Contessa may have such a view in mind in this passage.

(independently of one's being realist about its specific components, that is).

I attempt to answer something like this concern in chapter 7 by arguing that 'structure' is not a mere disjunction, but rather is itself structural (joint-carving carves at the joints). Contessa quotes and discusses part of this argument, but I think we are talking at cross-purposes. He takes me to be arguing that metaphysics will have an "arbitrarily demarcated object" unless the joint-carving notions are similar to one another; and he objects that metaphysics would not have an arbitrarily demarcated object if, say, a multiple-category ontology were correct. But what I was saying is that the joint-carving notions *aren't* similar in any way *other than being joint-carving* ('*first-order* heterogeneity'), just as Contessa says they might well be. (If logical notions carve at the joints then a multi-category ontology isn't needed to establish this point.) My point was that since joint-carving notions don't share anything *else* in common, being joint-carving itself had better count as a genuine similarity—otherwise explaining something in terms of 'joint-carving' would be like saying that a bereaved person and an Oscar winner both cried because they were both either happy or sad.

Despite being at cross-purposes, though, there is indeed a genuine disagreement here. For Contessa goes on to say that my general notion of structure (which includes both predicate and nonpredicate structure) does not have the explanatory value I claim it has. If true this claim would undercut my argument that structure is structural, and would strike at the heart of my book—nonpredicate joint-carving is one of the most distinctive features of my account, and is required for most of what I say about the metametaphysics of ontology, logic, time, and modality. I don't find Contessa's reasons for the claim compelling, however. I don't agree that only predicate structure is required in the case of inductive skepticism, for inductive skepticism could be based on "grue-ified" logical connectives rather than predicates; I don't think Contessa has shown that 'structure' is explanatorily otiose (as I explained above); and (as mentioned earlier) I think that structure can indeed be applied to nonpredicates like logical connectives, in which case joint-carving can play a role in securing semantic determinacy for such expressions. Moreover, questions about objectivity and substantivity can be raised for nonpredicates like modal operators, quantifiers, and the connectives of propositional logic; and nonpredicate joint-carving has a role to play in understanding and answering such questions.

3. Reply to Cian Dorr

How can we form reasonable beliefs about structure, about where reality's joints lie? Cian Dorr's preferred strategy uses the fact that structure is connected to various subject matters, such as reference, similarity, fundamentality, necessity, laws, explanation, substantivity, objectivity, and reasonable credence; he says: "using these connections, we can exploit our knowledge of those subject matters to gain knowledge about structure". (This strategy is not undermined if structure is more metaphysically fundamental than reference, similarity, and so on. To repeat what has become almost a cliché in recent metaphysics: the order of epistemic priority need not match the order of metaphysical priority.)

I recommended a different strategy:

The familiar Quinean thought is that we search for the best—simplest, etc.—theory that explains our evidence. My addition...is that this search is ideological as well as doctrinal; we search simultaneously for a set of concepts and a theory stated in terms of those concepts. We solve for the best and most explanatory pair $\langle I, T_I \rangle$ of ideology I and theory T_I in terms of that ideology. (p. 13)

I agree that Dorr's strategy is a reasonable one, but I think that mine is too.

As Dorr points out, my "best-pair strategy" is highly schematic. I left it that way because I wanted to remain as neutral as possible, because I thought that existing accounts of explanation, simplicity, and so forth—largely from the philosophy of science—would be applicable, and because I didn't know how to fill in the details. But keeping the account so abstract is admittedly worrisome, and Dorr argues that when we bear down on examples we can see that the account does not work.³ I appreciate Dorr's excellent paper for pushing me on this, and hope now to deepen my account by filling in more details.

At the very least, the "best" ideology-theory pair $\langle I, T_I \rangle$ ought to predict as much of the data (what we observe) as possible, and it ought to contain as simple and powerful laws as possible, where simplicity is measured in terms of I . For instance, syntactically simple generalizations using predicates like 'mass' and 'charge' have been found which, together with background assumptions

³Along the way he raises a number of other important and challenging concerns. Unfortunately, I only have space for ultra-brief replies to Dorr's two objections to my argument that structure is structural. In reply to the first: set-theoretic constructions of meanings are available (but see Sider (2013c)). In reply to the second: reference, similarity, and lawhood are surely not highly disjunctive since they themselves can take part in good explanations.

and perhaps bridge principles, imply much of the data; thus the criterion favors the ideology and theory of physics.⁴ Predictive generalizations using zodiac-theoretic ideology, on the other hand, have not been found; the criterion rules against astrology.

But now consider the first of Dorr’s three specific examples. Dorr compares a “state-of-the-art physical theory” T , containing the predicate ‘mass’ (measured in Planck units), with a theory T_1 that results from T by replacing ‘mass’ everywhere with the intuitively non-joint-carving predicate ‘mass*’, where the mass* of an object is defined as the seventeenth root of its mass. The challenge is to say why T_1 is a worse theory, given that the theories are syntactically identical.

Dorr notes that T_1 is inconsistent with our background knowledge of physics. For instance, the ratio of the mass of the proton to that of the electron is known to be (approximately) 1836; supposing T then to contain the sentence ‘the ratio of the mass of the proton to that of the electron is 1836’, T_1 contains the sentence ‘the ratio of the mass* of the proton to that of the electron is 1836’, which is false since the ratio of the seventeenth roots of the masses of protons and electrons is not 1836. But, Dorr claims, this fact cannot be used to say why T_1 is worse than T , since my best-pair strategy “is surely supposed to explain why it is rational to have [the background physical knowledge] rather than having to treat [this knowledge] as independently given”. If the background knowledge needed to be independently given, then the best-pair strategy would be useless, practically speaking, much like the advice “believe the truth”.

The best-pair strategy, I want to respond, explains *both* the background knowledge and the knowledge of where the joints lie. Thus the background knowledge doesn’t need to be independently given, and the best-pair strategy’s advice isn’t useless.

In order to clarify this—and to address Dorr’s examples in more detail—I will develop the best-pair strategy in a certain way. (I don’t think this development is mandatory, but it makes the arguments particularly clean.) First let us speak of joint-carving using a predicate of abstract entities, ‘ x carves at the joints’, rather than my official operator ‘ \mathcal{S} ’.⁵ Second, let us use theoretical terms, understood in David Lewis’s (1970) way, to denote the abstract entities we are evaluating for joint-carving. And third, when we are seeking “the best

⁴The way in which physics underlies other successful theories is another part of the case for physics.

⁵See (Sider, 2013c) for difficulties with defining this predicate using ‘ \mathcal{S} ’.

and most explanatory pair $\langle I, T_I \rangle$ ”, let the theories T_I be constructed using these Lewisian theoretical terms. Thus when we use the best-pair strategy to decide whether mass carves at the joints, we treat ‘mass’ as a Lewisian theoretical term standing for the property of mass, and we evaluate our theory containing ‘mass’ for simplicity, explanatory power, and so forth, under this semantic assumption.

According to Lewis, when a theoretical term τ is introduced by a theory Γ , τ means *the entity that (together with such-and-such other entities) uniquely realizes F_Γ* , where F_Γ is Γ ’s “realization formula”: the result of stating Γ in a single sentence and then replacing theoretical terms with variables.⁶ Given this account, Γ is definitionally equivalent to its “unique realization sentence”: the sentence saying that F_Γ is uniquely realized. Accordingly, the best-pair strategy as I am now developing it works in the following way. When we ask which theories best (most simply, etc.) explain the data, our question in the first instance is about unique realization sentences. When we discover which unique realization sentences best explain the data, we are entitled, *ceteris paribus*, to believe that they are true, and that the Lewisian theoretical terms defined from them denote entities that carve at the joints.⁷

In sum, to apply the best-pair strategy we introduce Lewisian theoretical terms for the entities to be evaluated for joint-carving, and compare how well their associated unique realization sentences explain the data.

Return now to Dorr’s first example, and let $T(x)$ be the realization formula for T . The following unique realization sentence presumably explains the data quite well:

(1) $T(x)$ is uniquely realized

So if no competing unique realization sentence can be found that explains the data better, the best-pair strategy lets us conclude that (1) is true, and that mass carves at the joints—where ‘mass’ is defined as meaning “the x that uniquely

⁶I follow Lewis in assuming that we have sufficient logical resources to construct the realization formula.

⁷There are some pressing questions of detail about this strategy (thanks to Cian Dorr for these and other observations). Should the term-introducing theories describe how the theoretical entities relate to observational matters, or should they merely describe the laws governing the theoretical entities? And should the quantifiers in the unique realization sentences be restricted to joint-carving entities, if our domain of abstract entities is construed “abundantly” (Lewis, 1986a, p. 59)? Certain answers to these questions may complicate the evaluations of simplicity, and may result in departures from the ordinary meanings of theoretical terms like ‘mass’ (though for the purposes of applying the best-pair strategy, such departures seem unproblematic).

realizes $T(x)$ ". Moreover, given this definition, (1) is equivalent to the theory T itself. Thus if it can be established that (1) explains the data better than competing unique realization sentences, this will support *both* the claim that T is true and the claim that mass carves at the joints; and establishing this fact would seem to be something that we can in fact do.

But now we reach a complication: how is 'mass*' to be understood? Is it a theoretical term tied to T_1 ? Or is it defined by the stipulation that mass* is to be the seventeenth root of mass? Or is it somehow both?

Case 1: suppose we treat 'mass*' as a theoretical term introduced by T_1 alone (thus ignoring the stipulation that mass* is to be the seventeenth root of mass!). Then, assuming 'mass' is a theoretical term defined by T alone, T and T_1 have the same realization formulas, and so 'mass' and 'mass*' are synonymous. We can therefore set case 1 aside.

Case 2: suppose we instead treat 'mass*' as a theoretical term introduced by a theory that includes both T_1 and the claim that mass* is the seventeenth root of mass. That theory would have to speak of 'mass', and thus would presumably also need to contain T as well. Thus the theory would need to incorporate all three of T , T_1 , and the claim that mass* is the seventeenth root of mass; its realization formula would be ' $T(x)$ and $T_1(y)$ and y is the seventeenth root of x ' (where $T_1(y)$ is the realization formula of T_1); and what we would be comparing to (1) would be this unique realization sentence:

(2) ' $T(x)$ and $T_1(y)$ and y is the seventeenth root of x ' is uniquely realized

But (2) is clearly false. No u and v could realize ' $T(x)$ and $T_1(y)$ and y is the seventeenth root of x ': to do so, v would have to be the seventeenth root of u , u would have to satisfy 'the ratio of the x of a proton to the x of an electron is 1836' (this formula is part of $T(x)$), and v would have to satisfy 'the ratio of the y of a proton to the y of an electron is 1836' (this formula is part of $T_1(y)$).

I don't see how else to treat 'mass*' as a theoretical term, so consider next case 3: suppose we treat 'mass*' as a nontheoretical term defined by the stipulation that mass* is the seventeenth root of mass, where 'mass' is a theoretical term defined solely by T . Now, the best-pair strategy as I've developed it does not apply directly to the question of whether mass*—understood as a nontheoretical term—carves at the joints (and I don't want to get into the thorny question of when theoretical and nontheoretical terms denote the same entity). But we can indirectly apply the best-pair strategy to non-theoretical terms, as follows. Suppose one defines a nontheoretical term τ using theoretical terms from one's current theory C , but then notices that C implies

the existence of simple and strong generalizations using τ . One might then consider a new theory N containing new theoretical terms in place of those in C , one of which, τ' , functions analogously to how τ functions in C , in the sense that N 's axioms imply generalizations containing τ' that correspond to the derived generalizations that follow, given the definition of τ , from C . If the unique realization sentence of N explains the data better than does the unique realization sentence of C (and better than other competitors), the best-pair strategy allows one to conclude that it is τ' , rather than any of C 's theoretical terms, that denotes a joint-carving entity.⁸ In the present case, however, this indirect method is unhelpful, since the generalizations implied by T concerning mass* (defined as the seventeenth root of mass) are more complex than the generalizations concerning mass. Suppose we construct a theory T' that includes a theoretical term 'mass**' behaving like T 's defined term 'mass*' (and includes no theoretical term behaving like T 's term 'mass'). The relatively complex mathematical generalizations governing mass* that in T are derived from simpler laws governing mass are mirrored by basic laws governing 'mass**' in T' . The best-pair strategy here applies directly, and tells us that mass rather than mass** carves at the joints.

Dorr describes his two other examples as follows:

- (ii) Let 'Q' abbreviate some true sentence not entailed by T —e.g. 'the moon and the sun look approximately the same size from Earth'. Let me introduce the concept *Q-mass* by stipulating that for any object x and real number n , n is the *Q-mass* of x iff either Q and n is the mass of x , or not- Q and n is zero. Let... T_2 be the [result] of replacing 'mass' in... T with 'Q-mass'.
- (iii) Let me introduce the concept *T-friendliness* by stipulating that for any object x , x is *T-friendly* iff T is true. Let... T_3 be the sentence ' $\exists x(x$ is *T-friendly*)' together with its consequences in first-order logic.

⁸Suppose, for instance, that we began by accepting classical mechanics under a spacetime formulation, but then encountered Jill North's (2009) argument that the Hamiltonian formulation is superior because phase space has less structure than spacetime. We had initially regarded the key terms of the Hamiltonian formulation (involving the geometry of phase space and the Hamiltonian) as nontheoretical terms defined from the theoretical terms of the spacetime formulation; but we can construct a new theory in which those key terms are theoretical terms; and if we buy North's argument we will hold that this new theory's unique realization sentence is explanatorily superior, and that it is the new theory's theoretical terms, rather than those of the spacetime theory, that carve at the joints.

These examples can also be handled by understanding the best-pair strategy as applying, in the first instance, to Lewisian theoretical terms.

Consider first T_2 . Is ‘Q-mass’ a theoretical term defined by T_2 alone (case 1), is it defined by Dorr’s stipulation (case 3), or is it defined by a theory including T , T_2 , and the stipulation (case 2)? Case 1 can be set aside as above. In case 2, we would be comparing (1) with a unique realization sentence saying that there are two unique quantities, each of which realizes a theory structurally similar to T , but where the second quantity yields the same values as the first if Q , but always yields 0 if not- Q . The latter sentence is inferior to (1) since its extra complexity is not outweighed by increased explanatory power.⁹ As for case 3, here we regard ‘Q-mass’ as a defined, nontheoretical term, and ask whether T implies derived generalizations using ‘Q-mass’ that are simpler and more powerful than its underived generalizations using ‘mass’. In fact it does not. Where G is any generalization implied by T , T together with the definition of ‘Q-mass’ does imply ‘if Q then G^Q ’, where G^Q is the result of changing ‘mass’ to ‘Q-mass’ in G ; but this would lead to a weaker and no simpler theory if ‘Q-mass’ were treated as a theoretical term.

Finally consider T_3 . To simplify, let’s redefine ‘is T -friendly’ as meaning ‘is such that $\text{Con}(T)$ ’, where ‘ $\text{Con}(T)$ ’ is the conjunction of the axioms from some axiomatization of T (as in note 6 I assume we have the linguistic resources to formulate this sentence). Case 1: ‘ T -friendly’ is a theoretical term defined solely by T_3 . Since T_3 includes only ‘ $\exists x(x \text{ is } T\text{-friendly})$ ’ (plus its logical consequences), the unique realization sentence for T_3 has very little explanatory power. Case 2: ‘ T -friendly’ is a theoretical term defined from a theory combining T , T_3 , and ‘an object is T -friendly iff $\text{Con}(T)$ ’. The unique realization sentence of this combined theory is more complex than (1) without providing any explanatory gain. Case 3: ‘ T -friendly’ is a nontheoretical term defined as meaning “is such that $\text{Con}(T)$ ”. Does T imply simple and powerful generalizations using this defined term? If T implies that at least one thing exists then it, together with the definition of ‘is T -friendly’, implies the sentence ‘ $\exists x(x \text{ is } T\text{-friendly})$ ’. But this sentence is not powerful on its own, and thus does not point the way to a better new theory in which ‘ T -friendly’ is a theoretical term. It is powerful only when combined with ‘an object is T -friendly iff $\text{Con}(T)$ ’; but in a theory where ‘is T -friendly’ is a theoretical term, this biconditional could not be regarded as a stipulative definition of ‘is T -friendly’, but rather would need to be an

⁹It does have a *little* more explanatory power if T implies that anything has nonzero mass, for in that case it will imply Q .

underived axiom, thus rendering the theory no simpler than T .

4. Reply to Kit Fine

Kit Fine’s rich paper raises important issues about the metaphysics of fundamentality. Fine and I agree that fundamentality-theoretic concepts are not to be defined modally, that they are not to be understood in terms of ordinary meaning, and that some such concepts are to be adopted as conceptually primitive. But we differ over which such concepts to adopt.

Fine’s (2001; 2012) primitive concepts are those of *ground* and *reality*. One proposition grounds another if and only if the first accounts for the second in a certain metaphysically distinctive way; one proposition holds in reality if and only if it is part of “reality’s intrinsic structure” (2001, section 8). My primitive concept, on the other hand, is that of joint-carving.

Ground and reality apply at the level of entire propositions (or sentences, or facts), whereas joint-carving applies at the level of parts of propositions (or sentences, or facts). Mass, existential quantification, negation, and the like, rather than entire propositions, are candidates for joint-carving. Thus my account is more atomistic than Fine’s: the locus of fundamentality for me is sub-propositional whereas for Fine it is propositional. Because of this difference, I said, our approaches differ over whether the notion of a “fundamental truth” obeys the following principle of Combination:

Combination If sentences S_1, S_2, \dots are fundamental truths then any true sentence S stated in the vocabulary of S_1, S_2, \dots is also a fundamental truth

For, I said, it is natural for me to define a fundamental truth as a truth that involves only joint-carving expressions—“fundamental terms” as Fine calls them; and it is natural for Fine to understand a fundamental truth as one that holds in reality. Thus my approach implies Combination whereas Fine’s allows him to deny it (since nothing forces him to hold that his primitive notion of reality obeys Combination). Furthermore, I claimed, this means that Fine’s approach can accommodate certain views that mine cannot. For example—this is Fine’s example—certain logical atomists thought that although $\sim P$ is a fundamental truth, $\sim\sim\sim P$ is not, since it is grounded in $\sim P$.

But as Fine points out, in chapter 7, in addition to the notion of fundamental truth just mentioned (a truth cast in fundamental terms), I introduced a second

notion of fundamental truth: that of *a truth that doesn't hold in virtue of others*. (In-virtue-of is my counterpart of ground.) My atomism about fundamentality commits me to Combination only about fundamental truths in the former sense; and moreover, I never considered Combination about fundamental truths in the latter sense.¹⁰ This leaves my discussion of the issue in disarray.

Fine's complaints (and perplexity) about what I said here are fair; and to be honest, some of them reveal parts of the book that were hazy and underdeveloped. Rather than retracing my steps, I'd like to move forward and reassess the situation.

One of my central principles is:

Completeness Every truth that contains at least one nonfundamental term holds in virtue of a truth that contains only fundamental terms

Completeness is an example of a “principle of underlying”, a principle saying that the fundamental in some sense underlies, is responsible for, everything else. Any theory of fundamentality must surely include some such principle. I accept this particular principle of underlying because of my atomism about fundamentality (which is at the very core of my position). For if I denied Completeness, and said, for instance, that $\sim P$ can hold in virtue of nothing at all even if \sim is not a fundamental term, then the notion of term-fundamentality—the most basic sort of fundamentality, for me—would play no role in my principle of underlying (and one's principle of underlying is central to one's theory of fundamentality).

Fine also accepts a principle of underlying, which says roughly¹¹ this:

Fine-Completeness Every truth that does not hold in reality is grounded in truths that do hold in reality

Thus we each accept that a certain status halts the demand for grounding/in-virtue-of propositions. For me that status is being true and composed of

¹⁰Fine guesses that the omission was because of my opposition to ground. This is partly true, though I must confess that I wrote chapter 8 before I wrote the bits of chapter 7 in which I distinguish two senses of ‘fundamental truth’, and did not properly reconsider chapter 8. But in any case, opposition to ground isn't a good reason to set the issue aside. First, I'm not so much opposed to ground as convinced that metaphysical semantics is a more useful notion, so I must still face the issue. Second, nothing I say about metaphysical semantics rules out the possibility that, e.g., $\sim\sim\sim P$ has $\sim P$ as its metaphysical truth condition, which again requires me to face the issue.

¹¹What he actually accepts is that every *factual* truth that does not hold in reality is grounded by truths that do hold in reality (2001, p. 27).

fundamental terms—the first sense of ‘fundamental truth’ above—whereas for Fine that status is holding in reality. Thus the “halting status” obeys Combination for me but need not for Fine. (The halting status is important to the theory of fundamentality, for as Fine has stressed (2012, section 1.2), inability to provide grounds, or see how they could be provided, is one of our main guides to fundamentality.)

Return now to the logical atomist. Suppose first that she denies that \sim is a fundamental term. In that case I cannot accommodate her position, for her view is surely that $\sim P$ doesn’t hold in virtue of anything, which would violate Completeness. Fine, on the other hand, could say that $\sim P$ holds in reality even though \sim is not a fundamental term, in which case her position would not violate Fine-Completeness.¹²

Suppose, on the other hand, that the logical atomist accepts that \sim is a fundamental term. Then my principle of Completeness does not conflict with her claim that $\sim P$ doesn’t hold in virtue of anything. Now, the logical atomist does say that $\sim\sim\sim P$ holds in virtue of something (namely, $\sim P$), whereas Completeness does not demand saying this (assuming that P contains only fundamental terms). But Completeness does not *prohibit* saying this, so I can allow the logical atomist’s position in this case.¹³ As for Fine, he will presumably say in this case that $\sim P$ holds in reality; so he too can allow the logical atomist to refrain from providing a ground for $\sim P$ while providing a ground for $\sim\sim\sim P$. But notice that since Fine is free to hold that $\sim\sim\sim P$ does not hold in reality (even though $\sim P$ does hold in reality), he is free to *require* the logical atomist to provide a ground for $\sim\sim\sim P$ while not requiring her to provide a ground for $\sim P$. (Here it matters that my halting status but not Fine’s obeys Combination.)

So in sum, although my account isn’t as inflexible in accommodating certain positions as I had claimed, it nevertheless remains less flexible than Fine’s. There of course remains the question of what conclusion to draw from this. I embraced the inflexibility; but Fine suggests that an appropriate dedication

¹²Fine might well not say this, however, since he might hold that containing only fundamental terms is necessary (though not sufficient) for holding in reality. (He would hold this if he defined term-fundamentality analogously to how he defines object-reality in Fine (2009, p. 172).)

¹³Fine’s framework may be a more natural home for this position, however: Fine’s grounding structure is richer than mine, and seems better suited for “discretionary” cases of grounding beyond what is demanded by principles of underlying (cf. his “impure logic of ground” (2012)). I’m not sure whether this difference between our positions is due merely to the fact that I was focused solely on nondiscretionary in-virtue-of relations, or whether it has a deeper source (e.g., in our difference over atomism).

to the “data” of metaphysics requires rejecting my account, and is mystified by my tendency to let theory dictate data. Although my treatment of the dialectic could have been clearer, rather than letting my theory dictate the data, I intended to argue for my theory on the basis of *all* the data. Some parts of the data admittedly appear to go against my theory, but others appear to favor it: my arguments for preferring my theory over Fine’s (see especially section 8.3.1), and the intuitive support that I think my atomistic approach enjoys. Choosing a theory on the basis of all the data requires examining whether the recalcitrant bits can be explained away. Still, there may be underlying methodological differences between Fine and me. Fine’s view appears to be that a metaphysics of fundamentality ought to be “neutral” in the sense of accommodating all positions that are “coherent” and have “some plausibility”; that is why it is so important for an account of fundamentality not to rule out the form of logical atomism under discussion. I on the other hand think that fundamentality is just another metaphysical issue, rather than a “first metaphysics” that lays out a neutral playing field in which all other metaphysical theories can compete (p. 136). So consistency with all coherent and plausible forms of logical atomism is not a datum that simply must be accommodated.¹⁴

Let me turn finally to Fine’s distinction between the E- and the D-projects:

The E-project is concerned with *saying* what can be said in the most fundamental terms, while the D-project is concerned with *describing* what can be described in the most fundamental terms. We can easily bring out the difference between the two projects with the case of disjunction. I can say ‘p or q’ and it is not clear that this can be said except by using disjunction or the like. But suppose now that I correctly describe the world by means of the sentence ‘p or q’. Then the use of ‘or’ is dispensable, since I can alternatively describe the world by means of p or q, depending upon which is true. Thus even though ‘or’, or the like, may be indispensable for saying what we can say, it would not appear to be indispensable for describing what we can describe.

Fine says that sometimes I seem focused on the D-project and sometimes on the E-project.

This distinction strikes me as illuminating, important to understanding different approaches to fundamentality, and worthy of future study. Whether

¹⁴Perhaps there is a further methodological difference over whether judgments about particular cases are invariably more secure than judgments about more general theoretical claims.

one accepts it is, I suspect, crucial on a range of questions, such as whether in-virtue-of should be understood “conditionally” or “biconditionally”¹⁵, whether fundamentality is atomistic or holistic, and whether the fundamental facts involve logical notions.

My approach to fundamentality is at its core based on rejecting the distinction, or at least, on denying that it has the significance Fine thinks it has. Perhaps the opposition can put by saying that I think the D-project should be pursued by following the rules that Fine thinks govern the E-project. My goal is to describe *the world* (and not just talk about what we can say); but I think that doing so requires saying what the distinguished structure of the world is, which requires more than giving grounds.

For example, Fine says that disjunction is dispensible in a description of the world (the D-project) since a complete description of the world can simply include whichever disjuncts are true. My view, on the other hand, is that if ‘or’ is a fundamental term (as it may well be), then a complete description of the world (and not just of what we can say) must include the true disjunctions. The world has fundamental disjunctive structure, and so no complete description can leave out the disjunctions.

Might one accept my theory of fundamentality but deny that ‘or’ and related terms are fundamental? Nothing in my general metaphysics of structure rules this position out, but my epistemology says to regard indispensable ideology as being joint-carving (section 2.3), and we presumably cannot dispense with all notions in the vicinity of ‘or’. Thus it would appear that my metaphysics conflicts with Fine’s distinction only in conjunction with this quasi-Quinean epistemology. Relatedly, Fine thinks that I neglect the distinction in my discussion of scientific theories. In reply to my claim that “good scientific theories... must be cast in [fundamental] terms” (p. 23), he says:

This may be fine under the E-project but not under the D-project, since even the logical constants may not be D-fundamental. Indeed, reference to theories is largely irrelevant to the D-project since theories are concerned to describe the world at a certain level of generality and so require notions which may not be required when the world is described in full specificity.

My view, though, is that when describing fundamental reality—the D-project—we have reason to prefer hypotheses that include simple strong generalizations, and reason to regard terms as fundamental when they are indispensable for

¹⁵See below, and Sider (2013c).

formulating such generalizations. Thus my epistemology of structure involves a sort of identification of the D- and E-projects.

So far I have been stressing the nature and depth of the opposition between my approach to fundamentality and Fine's distinction between the D- and E-projects. The distinction is fruitful in part for this reason. Also it is a natural distinction to make. Nevertheless I think that it is ultimately problematic.

On Fine's view, grounding suffices for (or is identical to) "metaphysical explanation". For example, $P \vee Q$ is metaphysically explained by P (when P is true), since disjunctions are grounded by their (true) disjuncts. The distinction/expression distinction, I take it, reinforces this: even though $P \vee Q$ says something more than P , it doesn't describe anything more; and in a good metaphysical explanation the explanandum doesn't describe anything more than the explanans. Now, metaphysical explanation is peculiarly ambitious: matters which can normally be taken for granted must be made explicit in a metaphysical explanation, since the point is to give a metaphysically complete account of the explanandum. Nomic relations, certainly, cannot be elided: one cannot say that the fact that the ball was dropped (in such-and-such circumstances) metaphysically explains its falling. Given this ambition, one might worry that even logical relations cannot be elided, and thus that P doesn't metaphysically explain $P \vee Q$ after all. But Fine's answer would be that whereas "the ball fell" goes descriptively beyond "the ball was dropped (in such-and-such circumstances)", $P \vee Q$ does not go descriptively beyond P (though it says more).

In what follows I'll argue against the distinction between describing and expressing, and the accompanying claims about metaphysical explanation. I'll do this by considering some examples, beginning with ones that put no pressure on Fine but ending with ones that I believe do. I choose the weak term 'put pressure on' deliberately, since nothing I say will force us to reject Fine's distinction. Rather, my hope is that the examples taken together will help us see that the distinction is illegitimate, and that in the cases where the distinction originally seemed so convincing—especially the logical cases of disjunctive and existentially quantified statements—we were simply neglecting the presence of logical elements in the world's description.

First example: surely Fine will want to hold that it's a fact about reality that a certain electron e is negatively charged. But imagine someone who says: no, the fact that e is negatively charged is grounded in the fact that e has a certain spatiotemporal trajectory. To be sure, 'negative charge' can't be defined in spatiotemporal terms, but that's just about what we can *express* (the E-project).

There is surely something wrong with such a claim; but what is it?

Here Fine has an easy answer: ground implies necessitation. If p grounds q , then necessarily, if p is true then q is true; since propositions about spatiotemporal trajectory don't necessitate propositions about charge, they don't ground them either. (I do worry that our modal belief here is justified by our belief in ground, and not the other way around; but set this aside.)

Next suppose someone says that what can be expressed using the idea of God, or of numbers, can't be expressed in other terms, but that even though God and numbers exist, the world can be completely described using 'God'-less, or nonnumeric, vocabulary. What is wrong with these at least *prima facie* unsatisfying claims?

Here the fact that ground implies necessitation is useless, since true propositions about God and numbers are necessarily true (let us suppose) and hence are trivially necessitated by any propositions whatsoever. So Fine must here rely on the (admittedly intuitive) judgment that propositions about God and numbers are not grounded—are not metaphysically explained by, are something “over and above”, describe the world further than—arbitrarily chosen propositions stated using 'God'-less and nonnumeric vocabulary.

Next example (and here is where I hope the pressure begins): when P is true, will Fine say that the fact that $\Diamond P$ is grounded in the fact that P ? Such a claim is allowed by the rule that ground implies necessitation (since P necessitates $\Diamond P$), and would be somewhat parallel to the claim that disjunctions are grounded in their disjuncts since the implication of $\Diamond P$ by P is sometimes regarded as logical. But it is hard to accept that the fact that $\Diamond P$ is “nothing over and above” the fact that P —that the latter is a complete metaphysical explanation of the former. By my lights, this is because the alleged explanans P contains nothing corresponding to modality; but if Fine agrees that P does not ground $\Diamond P$, this can't be his explanation, since similar reasoning would presumably imply that P does not ground $P \vee Q$.

Final example: are second-order existentially quantified claims grounded in their instances? When a respects b , for example, does this ground the fact that $\exists R R a b$? Analogy with Fine's claim that first-order existential quantifications are grounded in their instances suggests that he would answer yes; and if he does, I venture to guess, he would also say that *each* fact that is expressible in (extensional) second-order logic is grounded in some subclass of the following: i) the true atomic sentences, ii) the true negations of atomic sentences, and iii) a “totality fact” saying that there exist no objects other than a, b, \dots [a list of all objects]. But the latter claim is hard to accept. Surely the *expressive* power that

is added when one adds second-order quantifiers and variables to first-order logic corresponds to added *descriptive* power: the facts stateable in second-order logic are surely “above and beyond” the facts stateable in first-order logic. Metaphysical commitment to second-order logic involves accepting this further realm of facts. Perhaps this intuitive claim can be bolstered by noting that the second-order facts may as a group be more “complex” than the first-order facts. For instance, in a first-order language with symbols for successor and addition but not multiplication, the set of true sentences is decidable (Boolos et al., 2007, chapter 21); but if the second order quantifiers and variables are added to the language then the set of truths becomes undecidable (since multiplication can be explicitly defined from successor using second-order quantifiers and variables (Shapiro, 1991, chapters 4–5)).

What I think is that the last two examples—of possibility and plural quantification—should shake our confidence in Fine’s distinction between description and expression. For they share elements with what are supposed to be paradigm cases of extra expressive content without extra descriptive content—namely the cases of disjunction and existential quantification—and yet intuitively involve extra descriptive content as well. The distinction can seem unproblematic in the alleged paradigm cases perhaps because those cases involve first-order logical consequence, and first-order logical consequences often strike us as metaphysically “trivial”, an attitude I argue against in section 6.5 of my book.¹⁶ I conclude that even in the alleged paradigm cases, the extra expressive content is accompanied by extra descriptive content. Just as modal and higher-order facts are “over and above” nonmodal and first-order facts, so even existentials and disjunctions are “over and above” their instances and disjuncts.

One moral I would like to draw (although I don’t take it to follow from what I’ve said) is that we should take a “biconditional” rather than “conditional” approach to metaphysical explanation. Fine’s approach is a “conditional” one because his statement of metaphysical explanation ‘*P* grounds *Q*’ implies a (necessitated) conditional statement: necessarily, if *P* then *Q*. My “metaphysical semantics” approach, on the other hand, is biconditional because my “metaphysical truth conditions” are biconditionals (section 7.4). (The conditional/biconditional axis cross-cuts the linguistic/objectual axis that divides my approach from Fine’s: one could offer a biconditional form of Finean ground, on which the claim that *P* “biconditionally grounds” *Q* would imply that nec-

¹⁶I don’t mean to suggest that Fine himself is motivated by anything like logical conventionalism.

essarily, P if and only if Q .) If the biconditional approach is correct, that would explain what is wrong with the description/expression distinction. For the biconditional approach says that a metaphysical explanation of A must be a necessary and sufficient condition (of an appropriate strength) for A , and thus implies that P and Fa fail, respectively, to metaphysically explain $P \vee Q$ and $\exists xFx$. The problem with the conditional approach is that it demands too little of a metaphysical explanation.¹⁷

5. Reply to Eli Hirsch

I will focus on two issues from Eli Hirsch's generous and probing comments. The first concerns my "privileged-description claim": that in order to be fully successful, a representation must not only be true, but must also have a structure that "matches" reality's structure. In informal contexts I write as if I understand matching in a Tractarian sense: the representation must stand for a fact whose parts carve at the joints and correspond one-to-one with the representation's parts. But since I ultimately reject the Tractarian conception of matching, Hirsch argues, my privileged-description claim is obscure.

I reject the Tractarian conception of matching because my regimentation of talk about structure is "entity-free". Although in informal contexts I say things like "the predicate 'has unit negative charge' carves at the joints", "the property of unit negative charge carves at the joints", "facts about negative charge are part of the world's fundamental structure", and "there is a joint in nature corresponding to negative charge", ultimately structure is not about such entities as predicates, properties, facts, or "joints in nature" (section 6.3). In fact, it isn't about any entities at all. Thus my most basic claims about structure cannot employ a predicate: 'entity x carves at the joints'. (They must therefore differ from David Lewis's (1983*a*; 1986*a*, pp. 59–69) claims about naturalness, which use the predicate 'is a natural property'.) Instead I employ an "operator", 'struc'.¹⁸ This operator has an odd grammar: it attaches directly to expressions of multiple grammatical categories (to predicates, to sentential operators like 'and', etc.) to form complete sentences. Using it one can say, for example, that struc(has unit negative charge) and not struc(is grue), instead of making claims about predicates, properties, or other entities. Similarly, I would express the informal claim that conjunction carves at the joints by saying

¹⁷Though one may worry that it demands too much—see Sider (2013*c*).

¹⁸This is Hirsch's symbolization; in my book I use the symbol ' \mathcal{S} '.

“struc(and)”. (Not “struc(‘and’)” or “struc(conjunction)” —the claim isn’t about the word ‘and’ or an hypothesized abstract entity that it stands for. ‘Struc’ isn’t a predicate.) So I don’t want to understand matching as involving facts and their joint-carving parts, since I don’t accept such things as joint-carving parts of facts (not at the fundamental level anyway).

Hirsch finds this puzzling:

Here is the best I’m able to make of this. When I say that struc(and) I am in effect saying that the word “and” in some sense contributes to the fundamentality of the fundamental truths in which it figures. The contribution does not consist in there being some joint in the world that “and” corresponds to; it does not consist in there being some fact in the world with the shape of corresponding true sentences of the form “p and q”. What, then, does the contribution consist in? It seems as if Sider is trying to give us a form of Tractarian metaphysics (“language matches the structure of the world”) minus language-shaped facts. That’s a hard trick to pull off. The idea seems to be that truth is not enough, because the structure of our true sentences ought also to conform to “the structure of the world”, even though the world contains no structured items that correspond to the structured sentences. This is, for me, hard going. Sider’s “structure of the world” may seem to be intelligible only as something-we-know-not-what that plays the role of somehow imposing a metaphysical constraint on language beyond truth.

But why is my notion of the structure of the world so unintelligible, a mere “something-we-know-not-what”? When I speak of “the structure of the world”, I am merely saying that struc(and), that struc(has unit negative charge), and so forth. And what do I mean by ‘struc(and)’? Well, I mean struc(and)! This isn’t short for anything about words, so it *doesn’t* mean that the word ‘and’ contributes to the fundamentality of the fundamental truths in which it figures (though that certainly is a consequence of ‘struc(and)’). As I make clear, ‘struc’ is an undefined expression, intended to play a certain specified theoretical role.

In part, Hirsch seems suspicious of ‘struc’ itself. (He’s not the only one.) Perhaps this is because Hirsch doesn’t buy my defense (chapter 2) of the legitimacy of introducing undefined terms like ‘struc’ (which I liken to theoretical posits in physics). But perhaps instead he’s bothered by the “entity-free” nature of ‘struc’. ‘Struc’ is supposed to be about a certain subject-matter (joint-carving), and yet, there is no class of entities that the subject matter concerns. How can the world be structured, Hirsch wants to know, if there are no structured items?

“Nominalists”—in a broad sense of the term—often encounter resistance to their unfamiliar concepts, or unfamiliar interpretations of familiar concepts, which they must introduce in order to avoid unwanted ontological commitments. Thanks to Quine (1948), Chisholm (1957), and Boolos (1984), nowadays one can mostly get away with saying “The sentence is meaningful”, “I’m appeared to redly”, and “The Cheerios are in the bowl” without accepting meanings, sense data, or sets, since Quine, Chisholm, and Boolos’s entity-free concepts are by now entrenched. But where the trail has not yet been blazed, resistance to anything outside of the subject-predicate box remains strong. Arthur Prior put it well:

The fact is that it is difficult for the human mind to get beyond the simple subject-predicate or noun-verb structure, and when a sentence or thought hasn’t that structure but a more complex one we try in various ways to force it into the subject-predicate pattern. (1968, p. 15)

Prior’s example was our tendency to convert conditional statements “If *A* then *B*”, in which sentences *A* and *B* are not referring terms, into predicational statements “‘*A*’ implies ‘*B*’”, in which they are.¹⁹

Primitivists about modality normally express their modal claims using primitive sentential operators rather than primitive predicates of propositions.²⁰ ‘Necessarily, either snow is white or snow is not white’ is not about words, nor is it about propositions, nor is it about any other entity (other than snow); nevertheless it is a perfectly coherent worldly claim. For that matter, ‘Snow is *not* white’ is not about words or propositions or any entities at all (other than snow)—like ‘necessarily’, ‘not’ is an operator rather than a predicate—but nevertheless makes a coherent, worldly claim. Just so for ‘struc(and)’.

To be sure, my regimentation has a downside: general statements about structure present a challenge. A parallel point is familiar from the literature on nominalism proper. Nominalists have an easy time with ‘Ted is sitting’, ‘John is sitting’, ‘Ted is standing’, and so forth, since the singular terms in these sentences name individuals; but generalizations such as ‘Ted and John have something in common’ and ‘Ted is everything that John is’ are more difficult

¹⁹His observation also extends to constructions that are subject-predicate, but are not in the familiar singular mode: people often convert Boolos’s plural existential quantifier “there are some things...” to the singular “there is some plurality...”.

²⁰The best-known articles on the subject (e.g., Plantinga (1976); Stalnaker (1976)) do indeed speak of propositions (or states of affairs), but that is because these articles are primarily concerned with constructing possible worlds.

since they contain quantifiers that seem to range over properties. Options for nominalists include rejecting all such generalizations, regarding the quantifiers in them as being over predicates, and accepting irreducible quantification into predicate position. I similarly have an easy time with talk of structure in single cases, such as “struc(has unit negative charge)” and “struc(and)”, but generalizations, such as “anything that carves at the joints is more likely to be referred to” and “everything that carves at the joints is physical” are more difficult.

One might attempt to meet this challenge by accepting higher-order quantification. Then one could formulate statements like “ $\exists F \text{ struc}(F)$ ”. But whatever one thinks about primitive second-order quantification,²¹ its value in the present context is limited since the second-order variable ‘ F ’ can appear only in (first-level) predicate position, whereas ‘struc’ can attach to expressions of arbitrary grammatical category. For full generality, a sort of primitive higher-order quantification unlike any of the usual sorts would be needed. If an existential sentence “ $\exists \alpha \text{ struc}(\alpha)$ ” is to be implied by each of the following:

struc(has unit negative charge)

struc(and)

struc(there is)

the variable ‘ α ’ must have a bizarre grammar (corresponding to the bizarre grammar of ‘struc’): it must, presumably, be capable of occurring in every grammatical position that can be occupied by an argument of ‘struc’ (in predicate position, quantifier position, sentential connective position, and so forth).

I prefer to meet the challenge differently: by conceding that generalizations about structure cannot be formulated at the fundamental level, but introducing such generalizations at the nonfundamental level. Meanings may be constructed set-theoretically in some way; a predicate of those meanings, ‘carves at the joints’, may be defined in terms of ‘struc’; and then generalizations using quantifiers over meanings and the predicate ‘carves at the joints’ may then be formulated.²²

The issue of generality arises in connection with the notion of “matching”, with which we began. In particular cases, matching can be characterized using ‘struc’: the true sentence ‘There is a quark’ matches the structure of reality if

²¹See my section 9.15.

²²This isn’t all straightforward; see Sider (2013c).

and only if struc(there is) and struc(is a quark). But to define matching in full generality we must bring in quantification over meanings and the predicate ‘carves at the joints’: a true sentence matches the structure of reality if and only if the meanings of its words carve at the joints. On the surface this is pretty much the Tractarian conception, but the underlying entity-free metaphysics, stated in terms of ‘struc’, is quite different. I take it to be in the Tractarian spirit nevertheless, but perhaps Hirsch would disagree.

The second issue on which I will focus is Hirsch’s doctrine of “quantifier variance”. Hirsch’s (2011) defense of this doctrine has been immensely important for recent meta-metaphysics, both for critics and partisans. Its target is ontological disputes, such as that between Lewis (1986a, 211–3) and Peter van Inwagen (1990) over whether there are such things as (nonliving) composite material objects, such as tables and chairs. This debate is “merely verbal”, Hirsch says, because there are possible languages in which van Inwagen’s claims come out true and possible languages in which Lewis’s claims come out true, the difference being over what the quantifier ‘there is’ means. The first language’s quantifier is not supposed to be a restriction of the second’s, nor is ‘there is’ supposed to mean in either language something entirely unrelated to quantification. Rather, the two languages have selected two of the many possible “concepts of existence”.

While certain details of Hirsch’s argument are contestable, I believe that its core presents a deep challenge to ontology. My answer to it has been to concede his multiple quantifier meanings (existence-concepts) but argue that one of them carves at the joints, and thus is metaphysically distinguished. Ontology is about what there is in this distinguished sense.²³

One of my arguments for the existence of a joint-carving quantifier-meaning was that such a meaning is required to give a complete joint-carving description of reality. I expected Hirsch to reply either by denying the metaphysics of joint-carving (at least as applied to quantifiers), or by denying that a complete joint-carving description is necessary, or by arguing that a complete joint-carving description can somehow be given in non-quantificational terms. But Hirsch now suggests a very interesting fourth reply: any of his quantifier-meanings can be used in a complete joint-carving description since *all* of them carve at the joints. This doctrine of “egalitarian quantifier variance” would seem to mesh with the intuitive core of quantifier variance: that no way of

²³Sider (2001, introduction, 2009, 2011, chapter 9). In what follows I will speak in terms of meanings and the predicate ‘carves at the joints’, rather than in terms of ‘struc’.

carving up reality into a domain of objects is better than any other.

Egalitarian quantifier variance implies “redundancy” at the fundamental level, since, intuitively, a complete description of reality can be achieved with any one of Hirsch’s quantifier meanings (plus appropriate other joint-carving notions). But as Hirsch notes, I myself argue that accepting such redundancy is sometimes the right thing to do. Given that ‘not’ carves at the joints, it would be redundant to say that *both* ‘and’ and ‘or’ carve at the joints; but how to choose just one? One might plead ignorance; but instead, in order to avoid “metaphysical arbitrariness”, one might claim that each carves at the joints, thus embracing redundancy (section 10.2).

(The issue of redundancy is one of the thorniest connected with joint-carving. What one *wants* is a metaphysics of joint-carving on which the ‘or’-versus-‘and’ question simply doesn’t arise; but this turns out to be difficult to achieve.)

I will give two arguments against egalitarian quantifier variance. First, it posits far *more* redundancy than the egalitarian position about conjunction and disjunction (even if the latter is extended to the claim that all the one- and two-place truth functions carve at the joints²⁴). For the idea is presumably that “all possible” quantifier meanings carve at the joints, and not just a few (such as the Inwagenite and Lewisian quantifiers). There is, to be sure, a delicate question of what counts as a possible quantifier meaning, but the intuitive idea is that a possible quantifier meaning may countenance any objects that are consistent with certain “given” facts: Hirsch’s (1982, p. 32) incars and outcars may be countenanced, but not ghosts or God. (Incars are objects with the persistence conditions of cars, except that they can exist only inside garages; outcars are similar but can exist only outside of garages.) However the details are spelled out, the view would posit a massive array of redundant joint-carving meanings.

Compare the following issue about quantities. Which properties or relations concerning the quantity of mass (say) carve at the joints? One very unattractive answer is that the sole one is the *mass-in-grams* relation—the relation that holds between a physical object *o* and a real number *x* iff the mass of *o* is *x* grams. This is unattractive mainly because of its arbitrariness: why grams rather than some other unit, such as kilograms, or pounds? The arbitrariness could be avoided by holding that for *each* of the continuum many choices of

²⁴The claim that all finite truth functions carve at the joints, on the other hand, would involve massive redundancy, and for that reason strikes me as implausible.

unit, u , the relation *mass-in-unit- u* between physical objects and real numbers carves at the joints. But surely no one—not even a platonist—would accept this “egalitarian” view about mass. (This is not to deny that all scales are on a par. For the egalitarian view is just one way of securing that truth; a far better one is holding that there is a small number of joint-carving scale-invariant mass relations, such as *physical object o_1 is more massive than physical object o_2* .²⁵ Given appropriate assumptions about the behavior of such relations, one can prove via “representation theorems” that numerical mass values under any chosen scale may be used to represent the facts about the scale-free mass relations.²⁶)

We should reject the egalitarian view of mass because of the extent of the redundancy it implies. For the same reason we should reject egalitarian quantifier variance. One might object that if redundancy is objectionable then its extent does not matter: *any* redundancy would be objectionable, including the redundancy implied by egalitarianism about truth-functions. But this misunderstands the way in which redundancy of joint-carving meanings is “objectionable”. If redundancy were somehow metaphysically incoherent—if, say, ‘joint-carving’ were *defined* in terms of minimal supervenience bases—then the objection would be correct. But redundancy is not metaphysically incoherent; ‘joint-carving’—or rather, ‘struc’—is an undefined primitive notion. The sense in which redundancy is objectionable is instead epistemic: nonredundancy is, *ceteris paribus*, reasonable to believe. This is just a corollary of the view that simple theories are more reasonable to believe, assuming that redundancy counts against simplicity. Epistemic factors such as simplicity-maximization are matters of degree, and defeasible by other epistemic factors such as arbitrariness-avoidance, so the extent of the redundancy *does* matter.

Measuring redundancy is admittedly fraught. As Hirsch points out, “joints in nature” aren’t things that can literally be counted. (I have been speaking as if there are such things as joint-carving meanings, but this is just one of those cases where it’s more natural to speak using nouns; my ultimate claims are to be stated in terms of ‘struc’, and carry no commitment to joint-carving meanings.) Moreover, the complexity of a theory is not a simple function of how many joint-carving terms it contains. Nevertheless, it seems safe to assume that if $\text{struc}(q_1)$ and $\text{struc}(q_2)$ and ... and $\text{struc}(q_{1,000,000})$ (imagining that the quantifiers “ q_i ” from one million of Hirsch’s ontological languages have been imported into our language) then the world is very complex!

²⁵Or higher-order versions of this view, based on the approach of Mundy (1987).

²⁶See Krantz et al. (1971).

My second argument against egalitarian quantifier variance is that some of the quantifier meanings seem intuitively to be bad candidates for carving at the joints. For example, the quantifier from Hirsch's *incar* language, in which one can say truly that "there exist no cars, but there do exist incars and outcars", interacts closely with the nonfundamental subject matter of cars and garages, and does so in a bizarre manner. The quantifier that counts mereological atoms as having a fusion if and only if they are 666 in number seems oddly arbitrary. And for that matter, the quantifier meaning that recognizes all and only the objects of common sense seems too anthropocentric to be a likely candidate for joint-carving.

Hirsch might reply by restricting his claim that all quantifier meanings carve at the joints to the more "objective" and "non-arbitrary" quantifier meanings, such as the one that counts all pluralities as having fusions, the one that counts no pluralities as having fusions, and the one that counts only pluralities that would compose an instance of a natural kind as having fusions. But this would forsake much of the deflationary impact of quantifier variance, and would undermine Hirsch's defense of common sense ontology.

An intriguing alternate reply might be based on Karen Bennett's (2004) discussion of "plenitude", the view that each modal profile consistent with the nonmodal facts is instantiated by some object. To illustrate, here where I am located, sitting and typing, there exist according to plenitude (at least) three things distinct from me that have the same nonmodal properties as I do: something that essentially types and accidentally sits, something that accidentally types and essentially sits, and something that essentially types and essentially sits. Bennett points out that even though plenitude implies that modal properties are "brute" in the sense that the modal properties possessed by an individual thing do not supervene on its nonmodal properties, it nevertheless avoids a certain sort of arbitrariness since *all* consistent modal profiles are filled.

This is an instance of a general strategy: mitigate apparent arbitrariness (or other untowardness) in individual cases by appealing to non-arbitrariness at the level of all cases. The reply on behalf of Hirsch that I have in mind uses this general strategy: it says that even though many of Hirsch's quantifier-meanings seem arbitrary, anthropocentric, or in some other way nonfundamental, the theory as a whole is not arbitrary or anthropocentric because *all* possible quantifier meanings carve at the joints.

When Bennett uses the strategy, a single quantifier meaning is at issue, and arbitrariness at the level of individual objects is mitigated by nonarbitrariness at the level of all objects. In Hirsch's case, on the other hand, the anthropocentric-

ity or arbitrariness is at the level of certain quantifier meanings, and this isn't mitigated at that level; what is said to be nonarbitrary is the totality of quantifier meanings. Since the concern was about those individual anthropocentric or arbitrary quantifier meanings—they don't seem to carve at the joints—it isn't clear that the concern is answered. For comparison: claiming that *all* properties carve at the joints surely wouldn't make it easier to accept that, despite its anthropocentricity, the property *having a beautiful face* carves at the joints (or would it?).

6. Reply to Trenton Merricks

Trenton Merricks's illuminating paper begins by criticizing the first sentence of my book: "Metaphysics, at bottom, is about the fundamental structure of reality" (p. 1). Metaphysics isn't about any single thing "at bottom", Merricks says, and anyway, plenty of it isn't about structure. He's largely right; I got a little carried away with my rhetoric. But I do think that fundamentality is more central to metaphysics than Merricks allows.

Questions about which concepts are fundamental are indeed just one part of metaphysics. But there are also questions that use fundamental concepts, which I also meant to count as being "about the fundamental structure of reality".²⁷ If quantifiers carve at the joints, for instance, then ontological questions are about the fundamental structure of reality (see, e.g., p. viii).

Merricks might argue that ontological questions would be part of metaphysics even if quantifiers didn't carve at the joints. Strictly speaking he's right: they would be part of philosophy, and not part of epistemology, ethics, or any other branch. But they would be quite unlike metaphysics as traditionally conceived. For they arguably should then be addressed as Eli Hirsch (2011) thinks they should: by conceptual analysis. Hirsch and others with similar visions of how we ought to do ontology write *Metaphysics* on their *CV*s, but perhaps they ought to append asterisks.

Discussions of ontology, essence, properties, causation, and so forth have normally presupposed that fundamental reality has something like ontological, essential, property-theoretic, and causal aspects. If that presupposition is false, and if these matters are, as a result, mere projections of our conceptual scheme (in the sense I develop in section 4.4), we *could* continue to agonize, perhaps in

²⁷Not all such questions are metaphysical—think of physics.

a Hirschian, Strawsonian, or Knobian spirit, over whether, for instance, origins are essential; but no one on the contemporary scene has this attitude.^{28 29}

Fundamentality is also the central concern when metaphysicians tackle Frank Jackson's (1998) "location problem": *fitting things into* a given conception of fundamental reality. David Lewis spent many of his metaphysical hours showing how laws of nature, causation, counterfactuals, chance, and various other phenomena could be fit into a "Humean" world consisting of nothing but local qualities instantiated at points of spacetime (1986*b*, pp. ix–xvii). These phenomena are by Lewis's lights nonfundamental, but his discussion of them was animated by concerns about fundamentality since the point was to defend the underlying Humean metaphysics.

Does it matter what counts as metaphysics? Merricks says:

Sider thinks that a great deal hangs on the answer. For he thinks that metaphysics totters if its borders are arbitrary or result from highly disjunctive criteria. Thus he says:

The status of metaphysics itself hangs on [holding that fundamentality is fundamental]. In their loftiest moments, metaphysicians think of themselves as engaged in a profoundly important and foundational intellectual enterprise. But if fundamentality is highly disjunctive, the field of metaphysics itself—which is delineated by its focus on fundamental questions—would be an arbitrarily demarcated one. (140)

I reply that even if metaphysics turned out to be a hodgepodge, it could still be profoundly important. Metaphysics would be important just so long as (enough of) the topics in the hodgepodge were themselves important.

²⁸Kit Fine is, in a way, a fascinating exception. Throughout his career he has advocated a highly "metaphysical" world-view, with a rich structure of persistence, essence, dependence, and so forth; and yet in conversation he says that his writing is not meant to address what holds "in reality" (in the sense of his 2001). For instance, though he has famously defended three-dimensionalism (2006), he says that this claim concerns merely what is true, and that ultimate reality may well be four-dimensional. Our attitudes could be reconciled if facts about ground (see again Fine (2001)) held in reality, for then investigation into the unreal matters of persistence, essence, and the rest would simultaneously be investigation into the real matter of ground; but as I understand it, he does not think that facts about ground hold in reality.

²⁹Consider also free will. The traditional question was whether humans are part of the natural causal order. Although part of the literature remains focused on that question, another part simply assumes that humans are part of the natural causal order and focuses instead on the nature of human freedom given this assumption. The latter seems less centrally metaphysical, and indeed is largely written by those with Value Theory on their CVs.

I agree that if fundamentality isn't fundamental, the resulting arbitrariness of the borders of metaphysics wouldn't diminish the importance of all particular metaphysical questions. But it would diminish the importance of some of them: namely, those questions *about* fundamentality, such as “what exists in the fundamental sense?”, “what are the fundamental properties?”, and “is modality fundamental?”.

Merricks thinks that my misguided fixation on fundamentality leads me to misread David Armstrong (1997) and other defenders of truthmaking as claiming that an inventory of truthmakers is a complete description of fundamental reality. There may be some truth to this, but I do want to stick up a bit for this “reading of the texts”. First, Lewis in effect read Armstrong in this way:

When I first read [Armstrong's demand for truthmakers], the best explanation I could find was that Armstrong demanded that we do away with unanalysed predication. That seemed strange...[because all theories require unanalysed predication.] So what is *really* going on? I suggest that Armstrong has an unfamiliar notion of analysis. Analysis is not, primarily, a quest for definitions. Rather, it is a quest for truth-makers. (1992, pp. 202–3)

Lewis speaks here of analysis, but “what is the analysis of such-and-such?” played for Lewis much the same role as is played in more recent discussions by “what fundamental underlies such-and-such?”. Second, truthmaking plays the same regulative role in Armstrong's philosophy as fundamentality plays for others.³⁰ Armstrong counts the labor of “metaphysical explanation” as being complete when and only when truthmakers are given: phenomenologists who supply no truthmakers for counterfactuals about experience are criticized (Armstrong, 1989), and he himself stops his discussion of modality when he has given truthmakers (Armstrong, 1997, section 10.1).

A second set of Merricks's criticisms involve my theses of “Purity” and “Completeness”. For present purposes, take Purity to be the claim that fundamental facts involve only fundamental concepts, and Completeness to be the claim that nonfundamental facts hold in virtue of fundamental facts. Together they imply that “impure” facts hold in virtue of “pure” ones, where a pure fact is one that involves only fundamental concepts. (A fundamental concept is one that carves at the joints.) To illustrate, these theses together imply that the impure fact that there are cities must hold in virtue of some pure fact—perhaps

³⁰In Lewis that regulative role is played by “Truth supervenes on being”.

the fact that there are *C*s, where *C* is a complex predicate that is a “metaphysical definition” of ‘city’—involving, let us suppose, the subatomic properties and relations of the microphysical parts of its instances.³¹

Merricks’s argument is that I cannot uphold both Purity and Completeness since I cannot supply a candidate to be the pure fact in virtue of which the following impure fact holds:

NFC: There are no fundamental cities

where a “fundamental city” is defined as a city that is a fundamental entity.³² Although there is much of interest in his argument, I think it takes a wrong turn at the very beginning. For the first candidate that Merricks considers, and quickly rejects, is in my view the correct one: the fact that there are no fundamental *C*s. Against it Merricks says:

...fundamental cities are not supposed to be *C*s of any sort, and so they are not supposed to be *C*s that are fundamental, and so the non-existence of fundamental *C*s is irrelevant to whether there are fundamental cities.

Merricks intends ‘irrelevant’ so as to preclude the holding of an in-virtue-of relation; and by saying that fundamental cities “are not supposed to be” *C*s, he means, I take it, that it follows from the *definition* of ‘fundamental city’ that fundamental cities aren’t *C*s. To insure that this does follow from the definition, let’s redefine a “fundamental city” as a city that is a fundamental entity and not a *C*, and rewrite NFC accordingly:

NFC’: Nothing is: a fundamental entity, a city, and not a *C*

In these new terms, what I want to say is that NFC’ holds in virtue of the fact that nothing is: a fundamental entity, a *C*, and not a *C*. What’s the problem?

One might worry that NFC’ is not a logical truth, and yet is said to hold in virtue of something that is a logical truth. But such a situation is unproblematic, given the nature of in-virtue-of. When *P* holds in virtue of *Q*, *Q* is supposed

³¹In fact I suspect that a better metaphysical definition would be a role-theoretic one; see my reply to Schaffer in this symposium.

³²I’ll follow Merricks in using the term ‘fundamental entity’ (and I’ll pretend that it carves at the joints); but in fact it’s important to me to avoid the term (sections 8.5–8.7). The term suggests that fundamentality attaches, in the first instance, to individuals, whereas I think (roughly) that it attaches to concepts; and within my framework anyway, it conflates *existing* nonfundamentally with having a nonfundamental *nature*.

to be the underlying metaphysics of P , and not a matter of meaning. A linguistically competent speaker may express P without having any idea that it holds in virtue of Q , and indeed may be unable to express Q . Seeking in-virtue-of facts is thus nothing like conceptual analysis. So it is unproblematic that NFC' might be “metaphysically trivial” in the sense of holding in virtue of a logical truth, without this triviality being detectable simply by reflecting on NFC' and its meaning.

Thus the “purely metaphysical” nature of in-virtue-of is the key to addressing Merricks’s concern. Perhaps the issue can be brought out in a slightly different way. I accept:

NFM: Modality is not a fundamental concept

But that means that any fact involving modality must hold in virtue of a fact that doesn’t involve modality. (Any fact involving modality is by NFM impure, and so must, given purity and completeness, hold in virtue of some pure fact, which by NFM cannot involve modality.) But NFM itself is a modality-involving fact. What non-modality-involving fact does it hold in virtue of? My answer is: the fact that M is not a fundamental concept, where M is a “metaphysical definition” of modality (M will be a metaphysical definition of my Humean account of modality, as discussed below).³³ But consider a defender of primitive modality, who disagrees with me about NFM. She might protest that the fact that M is not a fundamental concept is not what she means to be rejecting, when she rejects NFM. It’s of course obvious that M is not a fundamental concept, she might say; what she rejects is the claim that *modality* is nonfundamental. But again: when P holds in virtue of Q , Q is not a sort of unveiling of what P means.³⁴ The in-virtue-of-relation is purely metaphysical. (In this respect, though not others, it is similar to supervenience.) So the modalist’s insistence that by denying NFM she does not mean to deny that M is a nonfundamental concept does not speak against NFM holding in virtue of that very fact.

Merricks has a followup argument concerning my distinction between “deflationist” and “nihilist” positions. I’ll illustrate the distinction in the case of causation. Deflationists and nihilists about causation agree that the “ordinary English notion of causation”—that is, the notion expressed by the ordinary English word ‘cause’—fails to carve at the joints. But the deflationist says

³³See p. 137 on denials of joint-carving.

³⁴It may be confusing that my official theory replaces in-virtue-of with metaphysical *semantics*. But a *metaphysical* semantics is quite different from what linguists and philosophers of language construct (section 7.4).

further that *no* notion of causation carves at the joints, whereas the nihilist accepts a joint-carving notion of causation—causation*, let's call it—and claims that nothing causes* anything.³⁵

Merricks's followup argument is that the metaphysical nihilist's claim that 'cause*' carves at the joints is incoherent. He says:

'Causes*' amounts to something like *causes and is fundamental*. By the lights of one who denies that causation carves at the joints—and so by the lights of the causal nihilist—'causes*' is akin to 'fundamental city'.

But since my view implies that 'fundamental city' does not carve at the joints, he says, 'causes*' can't carve at the joints either.

The problem here is the assumption that the causal nihilist will define 'causes*' as meaning "causes and is fundamental".³⁶ For since the definiens contains the ordinary English word 'cause', the causal nihilist would indeed be barred from saying that the definiendum 'cause*' carves at the joints. But my causal nihilist does not define 'cause*' in terms of 'cause', but rather introduces 'cause*' as an undefined predicate with the meaning-fixing stipulation that it is to carve at the joints and to play an inferential role like 'cause' in certain respects.³⁷

Merricks turns, finally, to my "Humean" account of (metaphysical) necessity. According to this account, for a proposition to be necessary is roughly for it to be a logical consequence of a certain class of propositions, the "modal axioms". Modal axioms come in different sorts, including mathematical truths, analytic truths (under a certain conception of analyticity), "laws of metaphysics", and "axioms of a metaphysical semantics". The account is a highly "deflationary" one in that no metaphysically deep condition is given to unite all the modal axioms. They are given by a mere list (mathematical truths, analytic truths, ...), which is selected, so to speak, "by us rather than by the world"—perhaps by linguistic convention.

Merricks's first objection is this:

³⁵There is a delicate question of what 'a notion of causation' means, if it's to include notions that are distinct from the ordinary English notion of causation—i.e., distinct from causation. The rough idea is that the notion plays a similar inferential role.

³⁶The logical form of the latter is not clear, but this doesn't affect the following argument.

³⁷Of course, if the world is as the causal deflationist says it is, those two stipulations cannot be jointly satisfied, and the term will be semantically empty. See section 5.3 and Sider (2013*b*) for more on the introduction of terms like 'cause*'.

Sider's reduction of necessity implies that the question of whether the laws of nature are necessary or instead contingent is not a substantive question about the metaphysics of those laws. But many metaphysicians—including many who think that modality has some reduction or other—will think that this implication is false. And I think that this implication is false. Thus I object that Sider's reduction of necessity is false because it has a false implication.

This is a little too quick. A question is nonsubstantive, to a first approximation, when it has different answers under different “candidate meanings” of the crucial terms in the question (chapter 4). The crucial term in modal questions is ‘necessarily’, which according to the Humean account does indeed have different candidate meanings, corresponding to different possible choices for what counts as a modal axiom. But not just anything counts as a candidate meaning. For instance, no candidate meaning counts a falsehood as a modal axiom—because, I would argue, any language using the term ‘necessarily’ with such a meaning would be “semantically alien” (p. 50). So the Humean account does not *immediately* imply that the question of whether the laws of nature are necessary is nonsubstantive; it must also be shown that the answer to the question turns on which candidate set of modal axioms our language selects.

However, I'm willing to concede that the answer to the question does turn on this selection. Suppose that the “best-system” account of the laws of nature is correct (as I think it is—section 3.1): laws are generalization in the simplest and most powerful deductive system cast in joint-carving terms. The modal axioms that constitute the actual interpretation of ‘necessary’ do not imply best-system laws, but a language that uses ‘necessary’ a little differently, by including best-system laws as additional modal axioms, does not strike me as semantically alien. Speakers of this language use ‘necessary’ to accomplish roughly the same semantic task as we do; they just do it a little differently. (Compare: the decision of whether to count the Pope as a ‘bachelor’ does not affect the semantic task accomplished by that ‘word’.) After all, the laws of metaphysics are modal axioms, and in my view, the laws of nature and the laws of metaphysics aren't so very different. So I concede that the Humean account implies that whether the laws are necessary is not a substantive question. Its answer is an artifact of a somewhat arbitrary semantic choice on our part (like the answer to the question of whether the Pope is a bachelor).

Merricks insists that the modal status of the laws of nature *is* substantive, but doesn't say why. So perhaps we just have a standoff. But let me point out next that the Humean account isn't as harsh here as one might think. In

particular, it doesn't imply that it's nonsubstantive what the right metaphysics of lawhood is. It's only the modal status of lawhood that is nonsubstantive. This is particularly important because positions on the modal status of laws are normally taken because of stances on the general metaphysics of lawhood: defenders of the best-system theory and the DTA theory usually hold that the laws are contingent, whereas defenders of scientific essentialism usually hold them to be necessary. And the question of which of these three metaphysical accounts of laws is not a modal matter, and so the Humean theory does not imply that it is nonsubstantive.³⁸ So if Merricks's insistence on the substantivity of modal status resonates with you, please check that you aren't conflating this with the substantivity of the nature of lawhood.

Merricks's second objection is to my attack on "arguments from possibility"—arguments of the form: Possibly, P ; it's not contingent whether P ; so, P . I claimed that my Humean account undermines such arguments for conclusions in fundamental metaphysics.³⁹ When P concerns fundamental metaphysics, any reason for believing the premise that P is possible, I said, is going to be a direct reason for believing the conclusion that P is actual; but if you had such a reason, the argument from possibility would be superfluous.

Merricks says that my critique would undermine *all* arguments from possibility. I think not; but seeing why requires examining the critique in detail.⁴⁰

Consider an argument from possibility for the proposition, UC, that composition is unrestricted. For a proposition to be possible is for its negation to not be necessary. So according to the Humean account, the argument's premise that UC is possible amounts to the claim that not-UC is not a consequence of the set A of modal axioms. Now, the only modal axioms that might imply not-UC are the laws of metaphysics; and we may assume for simplicity's sake that the only would-be law of metaphysics that might do this is not-UC itself. So for anyone apprised of this setup—of the truth of the Humean account,

³⁸Scientific essentialism is sometimes formulated in modal terms; but it arguably should instead be formulated as the thesis that the *real definitions* of scientific properties somehow involve lawhood, or that the *grounds* of particular matters of fact somehow involve lawhood, or that all *fundamental facts* in which scientific properties are attributed somehow involve lawhood. I explore these issues in forthcoming work. Second, it's natural to regard modally-formulated scientific essentialism as including the claim (or at least the presupposition) that modality *is* fundamental (contrary to what the Humean theory says), and this claim remains substantive even given the Humean theory.

³⁹Pp. 277–8; and see also Sider (2013a, section 10).

⁴⁰I do think my critique generalizes to *some* other cases, such as arguments from possibility for mathematical conclusions.

of the nature of laws of metaphysics, and so forth—the premise that UC is possible amounts to the claim that not-UC is not a law of metaphysics. Moreover, there aren't any reasons available for thinking that not-UC is not a law of metaphysics, other than reasons for thinking that UC is true. For the “laws of metaphysics” are just those propositions of a certain sort that happen to be true, and it's known in this case that not-UC is of the relevant sort; thus it's known that not-UC is a law of metaphysics if it's true. So the only reason one could have for thinking that UC is *not* a law of metaphysics would be a reason to think that not-UC isn't true—i.e. a reason to think that UC *is* true. In particular, the conceivability of UC is no reason to think that not-UC is not a law of metaphysics, and thus is no reason to think that UC is possible.

To see that this reasoning doesn't generalize, note first that it's not always true that the only available reasons for thinking that p is possible are reasons for thinking that p is true. This was so in the previous case because p 's negation (not-UC) was known to be the type of proposition that is a modal axiom (in particular, a law of metaphysics) if it's true.⁴¹ But in other cases one can know that p isn't a modal axiom without having any idea whether p is true. I can know that the proposition that there are an even number of trees in North America is possible without knowing whether it's true, simply by noting that its negation isn't the sort of proposition that could be a modal axiom (nor could it be entailed by modal axioms): it's the sort of proposition that, even if true, wouldn't be a law of metaphysics, or a mathematical truth, or any other sort of modal axiom (nor are there other modal axioms that would imply it).

No argument from possibility could be based on this last example since the number of the trees in North America is contingent. Can the Humean admit that some arguments from possibility are indeed successful? The matter is a little tricky. Two of Merricks's examples of arguments from possibility are arguments against the justified-true-belief theory of knowledge and Utilitarianism; but if those doctrines are taken extensionally then their negations are, respectively, that some person and proposition are such that it's not the case that the person knows the proposition iff she has a justified true belief in the proposition, and that some action is such that it's not the case that it is right iff it maximizes utility; and even Merricks will agree that since it's contingent whether any people or actions exist, these conclusions are contingent and hence can't be established by an argument from possibility. So the doctrines must be understood differently, if they're to be establishable by arguments from

⁴¹And because it was known that no other modal axiom would imply not- p .

possibility.

One might understand them as having modal content. For instance, letting q be the proposition that any person knows iff she has justified true belief, the JTB theory of knowledge could be taken to be $\Box q$ rather than merely q , in which case the Gettier argument Merricks considers would have the form: $\Diamond \sim \Box q$; it's not contingent whether $\Box q$; therefore, $\sim \Box q$. The second premise is now presumably true, so the question is the status of the first premise, which we may simplify to $\sim \Box \Box q$. Now, my book did not extend the Humean account to iterated modalities. But so long as such an extension is possible and vindicates the T-rule that $\Box \phi$ implies ϕ , the Humean can allow the following sort of reasoning in favor of $\sim \Box \Box q$:

The only kinds of modal axioms that might imply q are analytic truths (compare the proposition that any bachelor is male) and axioms in a metaphysical semantics (compare the proposition that any C is a city). But the conceivability of a person who has knowledge but not a justified true belief is evidence that q isn't implied by any such propositions. So q isn't necessary: $\sim \Box q$. And so, by the T-rule, $\sim \Box \Box q$.

This argument appeals in a crude way to a fact about conceivability to establish a conclusion about analyticity and axiomhood in a metaphysical semantics. But the appeal could be made more sophisticated, or replaced by other sources of evidence, such as empirical information about linguistic usage. The crucial point is that there are in-principle sources of evidence available for thinking that q isn't implied by the modal axioms, and these reasons are not of a sort that would render the modal Gettier argument superfluous.

(Merricks also mentions a third argument from possibility: the use of the possibility of Black's spheres to refute the identity of indiscernibles ("II"). Here I want to stick up for my critique, for reasons roughly parallel to those given in the case of UC. The claim that it's possible for Black's spheres to exist ($\Diamond BS$) amounts to the claim that not-BS is not a consequence of the laws of metaphysics; the only relevant would-be law of metaphysics is II itself; so a reason to believe $\Diamond BS$ would need to be a direct reason against II, mooting the argument from possibility against II.)

To sum up: arguments from possibility have a fighting chance in some cases but not others because we have different sorts of access to different sorts of modal axioms. Only evidence for truth or falsity bears directly on whether

certain sentences are laws of metaphysics; but other sources of evidence (such as conceivability) bear directly on whether certain sentences are analytic.

So I don't agree that my critique of arguments from possibility would generalize as much as Merricks says it would. But I do concede that my discussion of this topic was incomplete. It remains unclear exactly which such arguments the Humean account allows, and, more generally, unclear how much of ordinary modal practice can be reconstructed in Humean terms.

7. Reply to Jonathan Schaffer

In addition to characterizing fundamentality itself, a theory of fundamentality must also account for the *connection* between fundamental and nonfundamental. The fundamental “underlies” everything else, but in what sense exactly? According to Jonathan Schaffer's kind yet deeply challenging critique, my account of the connection—metaphysical semantics—cannot handle multiply realizable nonfundamental facts. This is an important criticism, and I'll discuss it in a moment, but first I want to resist Schaffer's description of my account as being “radically eliminativist” and implying that the nonfundamental is “mere talk”.⁴²

Like many others recently, I reject modal accounts of the connection. (Even though mathematics trivially supervenes on anything at all, there remains a nontrivial question of what underlies mathematical truth.) Closer to the truth, I think, is the view that fundamental truths *ground* fundamental ones, where grounding is understood as a relation, not definable in modal terms, that underwrites (or is) a distinctively metaphysical sort of explanation.⁴³

Still closer to the truth, I hope, is my own account of the connection, according to which each language has a “metaphysical semantics”. Like a linguistic semantics, a metaphysical semantics is an explanatory theory of certain facts involving linguistic communication. A metaphysical semantics can take the same form as a linguistic semantics: it may assign truth-conditions to sentences of the language in question, for instance. But unlike a linguistic semantics, a metaphysical semantics has the goal of explaining how the use of the language in question fits into fundamental reality. This results in two main differences. First, truth-conditions assigned by a metaphysical semantics (“metaphysical

⁴²Caveat: though metaphysical semantics per se is not eliminativist, I do provisionally defend eliminativism about composite objects elsewhere (Sider, 2013a). However, the view I ultimately embrace—an ontology of sets as well as spacetime points—is hard to categorize.

⁴³See Fine (2001, 2012); Rosen (2010); Schaffer (2009); see Wilson (2013) for dissent.

truth-conditions”) must be stated in perfectly fundamental (joint-carving) terms. So unlike a linguistic semantics, a metaphysical semantics could not include the biconditional “‘Snow is white’ is true if and only if snow is white”; doing so would assign ‘snow is white’ as a metaphysical truth-condition, and that sentence contains the nonfundamental terms ‘snow’ and ‘white’. Second, a metaphysical semantics does not aspire to explain as many facts about linguistic communication as a linguistic semantics. For instance, when a linguistic semantics assigns a certain truth condition to a given sentence, the truth condition is normally supposed to in some sense encode what the speaker understands by the sentence. But metaphysical truth-conditions are claims about fundamental reality—about fundamental physics, perhaps—beyond the ken of ordinary speakers.⁴⁴

Suppose, just for the moment, that fundamentally speaking, the only entities that exist are subatomic particles. (This is not in fact my own view.) Are English sentences that quantify over further entities, such as atoms of hydrogen, then false? Not necessarily, I said. For a metaphysical semantics might assign metaphysical truth-conditions that count such sentences as true. It might, for instance, assign ‘there exist an electron and a proton bonded to each other’ as the metaphysical truth-condition for the English sentence ‘there exists an atom of hydrogen’ (section 7.7).

Thus my metaphysical semantics approach allows ‘There exists an atom of hydrogen’ to be a true sentence of English even if there are, fundamentally speaking, no atoms of hydrogen. It is for this and related reasons that Schaffer calls my account eliminativist, and says that it renders quantification over atoms of hydrogen, for instance, as being mere talk. But what I say about atoms of hydrogen (under the current suppositions about fundamental ontology and the metaphysical truth-conditions of sentences about atoms of hydrogen) is closely parallel to what grounding theorists like Schaffer himself say. I say that the metaphysical truth-condition of ‘There exists an atom of hydrogen’ is ‘there exist an electron and a proton bonded to each other’; they say that the fact that there exists an atom of hydrogen is grounded in the fact that there exist an electron and a proton bonded to each other. Although there are important differences between our accounts, they don’t amount to mine being more eliminativist.

⁴⁴Exactly what, then, must a metaphysical semantics explain? This is a pressing question, and alas, my book does not contain an adequate answer. The grounding theorists have a counterpart lacuna: what is this “distinctively metaphysical sort of explanation”?

My account isn't eliminativist in the most straightforward sense, anyway, since (under the current suppositions) it says that there are atoms of hydrogen! The English sentence 'There are atoms of hydrogen' is true (since it has a true metaphysical truth-condition), and so, disquoting, there are atoms of hydrogen.⁴⁵ (This paragraph is written in English, remember; its quantifiers are the English ones, not the the joint-carving ones.)

Schaffer might object that the account implies that there aren't *really* atoms of hydrogen. But what would that mean? If the function of 'really' is just to cast aside loose talk, hyperbole, metaphor, and so forth, then the account implies no such thing. The account isn't saying that when an ordinary speaker utters 'there exists an atom of hydrogen', "all she really means" is that some electron is bonded to some proton—metaphysical truth-conditions are no more an account of what one ordinarily means than grounds are. No, she means—really—that there is an atom of hydrogen. And there really is.

If, on the other hand, 'really' means something like 'fundamentally', then my account does indeed imply that there aren't really atoms of hydrogen (under the current suppositions). But the grounding theorists agree! To be sure, there are subtle differences on how we understand "there fundamentally exists an *F*". On my account, this amounts (roughly) to saying that *F*s are in the range of a joint-carving quantifier; according to a grounding theorist like Kit Fine, it amounts to saying that facts about *F*s hold "in reality", by which he means roughly that they are fundamental facts; and according to a grounding theorist like Schaffer, it amounts to saying that *F*s are not grounded in further entities. The differences between these three approaches are important in other contexts (see chapter 8), but in the present context all three seem like variations on a common theme. I did include the grounding approach in a chapter on my "rivals", but my discussion there primarily argued against the view that facts about ground are fundamental facts—in short because fundamental facts such as that *So-and-so grounds the fact that someone smirked* would involve the property of smirking, and surely no fact about smirking, not even a fact about what grounds it, is fundamental. This argument doesn't touch grounding itself, so long as such facts about ground are nonfundamental. And if grounding theorists concede this (as I think they should⁴⁶), then the grounding and metaphysical semantics approaches remain pretty close. Indeed, one can view metaphysical semantics as

⁴⁵Similarly, since the best metaphysical semantics for English presumably includes metaphysical truth-conditions for fact-talk, my account allows that there are nonfundamental facts, such as the fact that there are atoms of hydrogen.

⁴⁶See Bennett (2011) on this issue.

a linguistic account of ground. I do argue that certain subtle (though important) differences between the accounts favor mine, but these differences don't add up to a different "amount of reality" accorded to nonfundamental entities or discourse.⁴⁷

Moving on: Schaffer's main criticism is that my account cannot handle multiple realizability. He points out that many of my examples of metaphysical truth-conditions presuppose a particular physical theory (often, a toy particle physics). While this may be fine for sentences involving natural kind terms (such as 'hydrogen'), he says, it is a problem for multiply realizable sentences. For instance, the metaphysical truth-condition for:

(M) Moore has hands

cannot be:

(P) The subatomic particles that are part of Moore have such-and-such masses and charges and spins and are in such-and-such spatial arrangements

since Moore could have had hands if physics had been field-theoretic rather than particle-theoretic, or if physics had been particle-theoretic but particles had instantiated fundamental properties other than mass, charge, and spin. Schaffer then goes on to argue against various ways I might attempt to respond to this problem, for instance by constructing a metaphysical truth-condition by disjoining all the possible realizers.

⁴⁷The subtle differences I discuss in my book (section 7.9) involve nonfactual discourse. There are other differences, which I hope to explore in future work. For instance, metaphysical semantics is a "biconditional" rather than "conditional" conception of ground (see below); and metaphysical semantics (as I develop it) has no analog of grounding in facts short of fully fundamental ones. But these too are subtleties, and I would not describe my position as "reject[ing] the posit of a metaphysical relation of grounding" (Schaffer's note 2)—unless 'metaphysical relation' is meant to imply 'fundamental relation'—but rather as holding that the relation of ground is a good tool, but not the best tool, for characterizing the relation between fundamental and nonfundamental. One final speculation. On my view, given the sparse ontology we are assuming for the sake of argument, the English quantifier 'there is' cannot mean the same thing as a joint-carving existential quantifier. Could this fact, the fact that I embrace a sort of quantifier variance (albeit quite different from Hirsch's (2011)), be why Schaffer thinks of my view as eliminativist? If so, I would press Schaffer on the status of the quantifier he uses to describe his "hierarchy of being" (2009, p. 354): if he accepts talk of joint-carving quantification at all, he must surely either say that this quantifier does not carve at the joints, or else embrace a larger fundamental ontology than we have been assuming.

There are indeed serious challenges here (though I don't quite agree with Schaffer on their exact nature) to which I do not have entirely happy answers.

For the remainder, let's assume the supersubstantialist and set-theoretic worldview discussed in the final chapter of my book. The fundamental ontology of this view consists of points of spacetime and sets; and the joint-carving notions include a predicate for set membership, predicates of physical geometry, and predicates for physical magnitudes—perhaps charge, mass, and spin (construed as applying to points of spacetime). Given this worldview it is natural to identify a persisting material object with the set of points in its world tube, in which case a metaphysical truth-condition for (M) will have the form:⁴⁸

M has subsets that are H

where ' M ' names the set of spacetime points identified with Moore, and H is a "metaphysical definition" of 'is a hand'. The question is then how to construct H . Or better: whether there are in-principle obstacles to doing so—I argue in section 7.6 that actually coming up with metaphysical truth-conditions for macro language is neither attainable nor necessary.

H cannot be a completely specific description, in microphysical terms, of the spacetime points that actually make up Moore's hands. Since each person's hands are a little different, this approach would treat different occurrences of the predicate 'hand' differently, leading to a poor metaphysical semantics, by which I mean one that fails badly as an explanatory theory by ordinary standards of good explanation. For similar reasons, H cannot be a disjunction of all possible microphysical descriptions of hands: even if our fundamental theory included infinitary notions (which I assume it does not), such massive disjunctiveness would undermine the explanatory value of the theory.

A better approach is for H to be a functional definition of 'is a hand':⁴⁹

FR: has some property P that plays functional role R

This approach both yields a more explanatory theory (since it avoids extreme disjunctiveness) and covers more possible ways in which (M) might be realized (though see below). Now, all the words in a metaphysical truth-condition must carve at the joints. So the next question is whether the schematic FR can be filled in consistently with this requirement.

⁴⁸For simplicity I treat (M) as saying that Moore contains his hands as timeless parts.

⁴⁹I suggest something like this approach on pp. 130–1 and 294–5.

Schaffer considers the possibility of a functional definition, and says that “multiple realizability rearises: [FR] is itself multiply realizable.” But the multiple realizability problem above was a problem for putative metaphysical truth-conditions like (P) which mention actual-world physics (‘particle’, ‘charge’, etc.) in a way that precludes being true if physics had been different, and FR contains no terms from actual world physics. It rather contains *general* terms like ‘property’ and ‘functional role *R*’; and it *quantifies* over properties without mentioning specific ones such as charge or mass. So long as the metaphysical semantics for the expressions occurring in FR is not tied to actual-world physics, metaphysical truth-conditions based on FR can be true under various possibilities for the nature of physics.

But can the expressions in FR indeed be given such a metaphysical semantics? Let’s examine how ‘property’ and ‘plays functional role *R*’ might be “metaphysically defined”, how any terms used in their definitions might themselves be metaphysically defined, and so forth, until finally joint-carving terms are reached. (Since I’m only trying to show that the definitions are in-principle available, during this drill-down process I will merely gesture at styles of available definitions, rather than attempting to actually construct them.)

First, ‘functional role *R*’ will presumably be an abstract specification of a causal or nomic role. So we need a metaphysical semantics for causal and nomic terms. This I would approach in the Lewisian way: causation reduces ultimately to lawhood plus possible worlds, and laws are regularities in the simplest and most powerful system in a language with predicates for natural properties. Worlds I could identify with maximal and modally consistent sets of propositions, and ‘modally consistent’ can be reduced using the “Humean” account of necessity from chapter 12. For ‘proposition’ and ‘property’ I propose a set-theoretic reduction. Define a *simple property* to be a set, and a simple relation to be a set of ‘tuples.’⁵⁰ Then, in order to obtain a more fine-grained space of abstract entities, define structured properties, relations, and propositions in an iterative way, with the simple properties and relations as the basis (compare Lewis (1986*a*, pp. 56–7)). For example, the structured conjunction of properties *P* and *Q* might be defined as the ordered triple ⟨‘and’, *P*, *Q*⟩, and the structured proposition that is the existential generalization of property *P* might be defined as the ordered pair ⟨‘there is’, *P*⟩. Define other property- and

⁵⁰I reject merely possible individuals, and so on my view the members of simple properties and relations are actual entities. However, the definition of ‘simple property’ is not tied to the actual roster of entities or their natures; ‘simple property’ is just defined as meaning ‘set’.

proposition-theoretic notions accordingly (for instance, x *instantiates* a simple property iff x is a member of it; x instantiates \langle 'and', P , Q \rangle iff x instantiates P and x instantiates Q ; the proposition \langle 'there is', P \rangle is true iff something instantiates P). These iterative constructions can include definitions of higher-order properties and propositions, of the sort needed for specifying functional roles. (A simple definition of this sort would identify the property of *being a property such that so-and-so* with the set of sets such that so-and-so.)

Thus a metaphysical semantics for FR will be (fairly) physics-neutral. It will begin by asserting that there exists some property—i.e. set-theoretic construction— P , without specifying whether the elements of P are particles, or fields, or points of spacetime, or something else. It will go on to say that P plays a certain nomic or causal role; and although this claim will involve the general notions of law and cause (or rather their Lewisian reductions), it will not specify what the laws and causes are, but rather will say that there is some network or other of laws and causes in which P plays such-and-such an abstractly-characterized role.

But there remains a serious problem of multiple realizability. Lewis's account of lawhood uses the notion of a natural property. This might seem unproblematic since my notion of structure is very much in the spirit of Lewisian naturalness, and since I hold that "structure is structural" (section 7.13), making the notion of structure acceptable for use in metaphysical truth-conditions. However, my official notion for talking about structure is not something that can be predicated of properties, or of anything else for that matter. It is rather an operator, ' \mathcal{S} ', that attaches directly to predicates (and other expressions). If my official account were more platonist, employing, say, an ontology of universals and a primitive predicate 'is a fundamental universal', then I could say that a simple property P is fundamental if and only if for some fundamental universal, U , the members of P are exactly those things that instantiate U . But given my more nominalist account I cannot say this or anything like it.

Could I say instead something of the form:

- (1) Simple property P is fundamental iff: P 's members are all and only the charged things, or P 's members are all and only the massive things, or...

(where 'charged', 'massive', etc. are the joint-carving predicates)? Here there is a multiple realizability problem, since (1) does not allow for the possibility of fundamental properties other than the actual ones.

Could I say rather:

- (2) Simple property P is fundamental iff: some joint-carving predicate F applies to all and only the members of P

? No. For we would then need to construct a metaphysical truth-condition for the semantic notion ‘applies to’. As we know from Hartry Field (1972), there would be a choice: will this metaphysical truth-condition be list-like (“ F applies to x iff $F =$ ‘is charged’ and x is charged, or $F =$ ‘is a dog’ and x is a dog, or...”), or will it be “substantive”? The former would reinstate the multiple realizability problem, and the latter would be circular in the present context since a substantive metaphysical truth-condition would presumably need to bring in causal relations, perhaps between the linguistic population in question and its environment.

There is another, separate, problem of multiple realizability. Can functional role metaphysical truth-conditions be constructed “neutrally” so as to accommodate all possible realizers? The problem we’ve been discussing so far is that *no* variation in the fundamental properties can be accommodated, since the only available definition of ‘fundamental simple property’ is list-like. But a further problem is whether quite *radical* variation in the possible realizers can be accommodated. It is really this second problem that is Schaffer’s central concern:

...the truth that Moore has hands can presumably be grounded in many different distributions of fundamental fields, arrangements of point particles, vibrations of strings, or whatnot.

A *fairly* wide range of metaphysical underpinnings can be accommodated (provided the problem with defining ‘fundamental simple property’ can be solved), since few assumptions about the nature of fundamental reality were built into the functional-role metaphysical truth-conditions I suggested above. The main notions in such truth conditions, as we saw, are the notion of property and causal/nomic notions. Properties were constructed from sets, and the definitions did not specify what goes into the sets. The definitions allow the members of these sets to be fields, particles, points of spacetime, points of configuration space, or whatever. The Lewisian definition of a law of nature as a regularity in the simplest and strongest system is similarly neutral, since the simplest and strongest system might turn out to be field-theoretic, particle-theoretic, or even based on configuration space rather than spacetime. Causation then inherits this neutrality. So as far as these key notions in functional-role metaphysical truth-conditions are concerned, the range of accommodated possibilities is wide.

But it is arguably not wide enough. Initially I suggested that the metaphysical truth-condition for (M) should take the form “ M has subsets that are H ”. This choice was based on the identification of persisting things with their world-tubes, and thus is not neutral in a crucial way. If the fundamental physical space were configuration space rather than spacetime, for instance, as some “configuration space realists” suppose,⁵¹ then physical objects could not be identified with their world-tubes. Indeed, it is unlikely that particular physical objects could be identified with particular parts of configuration space: physical-object talk would depend on fundamental reality in a more holistic way, as Alyssa Ney (2012) argues. Thus the metaphysical truth-condition of (M) cannot have the form “ M has subsets that are H ” if it is to accommodate this possibility. Relatedly, my suggested metaphysical truth-conditions quantify over sets, and thus cannot accommodate the possibility that sets do not exist. Relatedly, though more far-fetchedly, these truth-conditions employ “thing-talk”, and thus apparently cannot accommodate the possibility that the notion of a thing is not applicable (a world of “stuff”, say).

Some such concerns can be answered by formulating metaphysical truth-conditions more “abstractly”. One might answer the concern about configuration space, for instance, in the following (highly schematic) way. Define the notion of a “fundamental description of the world”. Next define the notion of a given fundamental description “inducing” various sentences in the language of substantival spacetime. When the fundamental description is itself a spacetime description then the induced sentences are just true fundamental sentences; but when the fundamental description is not a spacetime one, the idea is for induced sentences to be nevertheless “appropriate” ones. The world is “as if” the induced sentences are true, except for the fact that the ultimate metaphysical story is not one of spacetime. Leibnizian relationalists, for instance, are familiar with showing how spacetime-theoretic sentences are in some sense close to the truth—close enough to be legitimate for use in physics—even if not fully accurate; in my terminology these sentences are induced by the relationalist’s fundamental description. But I want also to include inducing by fundamental descriptions that differ more drastically from spacetime substantivalism—for instance, by the fundamental description of the configuration space realist. I’m optimistic that a metaphysical semantics for ‘inducing’ can be constructed in highly abstract terms—terms that don’t presuppose much about which fundamental description is actually true—that

⁵¹See Albert (1996).

will apply to a wide range of fundamental descriptions (especially if we're willing to live with some indeterminacy and arbitrariness). We ourselves seem to have the ability to read off, in many cases anyway, the spacetime upshot of non-spacetime theories. We can do this, for instance, with relationalist fundamental theories, and with presentist fundamental theories. This ability presumably involves the application of some sort of algorithm, which algorithm could be encoded in a metaphysical semantics. The algorithm presumably employs such tricks as looking for "natural embeddings" of models of one description in those of another, and "reverse-engineering": defining the induced theory as the one that would "generate" (or be generated by) the given fundamental description by certain rules.⁵² Finally, given the notion of inducing, we can say that the metaphysical truth-condition for (M) is that ' M has subsets that are H ' is induced by the fundamental description of the world.

There are limits to this approach. The more fundamental theories one tries to accommodate, the more abstract the definition of inducing must become; but the definition must nevertheless continue to have enough content to actually work. Eventually these two pressures will conflict too much. (It's hard to see, for instance, how to make the account so abstract as to cover even non-set-theoretic or non-thing-theoretic possibilities.)

So: what to say about these problems? Their source is two elements of my account; and one or both of those elements could be given up. The first element is the one on which Schaffer focuses: my metaphysical semantics approach to the connection between fundamental and nonfundamental. Now, Schaffer says in his concluding section that the problem is with the eliminativist nature of metaphysical semantics; but as I explained above, metaphysical semantics is no more eliminativist than grounding. Schaffer expands on this charge by saying:

If Sider would countenance nonfundamental 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*.

Here Schaffer construes the issue as being whether "the explanation" of the fact that we tend to say that Moore has hands is *Moore has hands*, or whether it is instead some claim about fundamental reality. But I doubt that there is any

⁵²See Sider (2008) for an instance of the latter trick. Inducing is of course a special case of ground.

such thing as *the* explanation of anything; what counts as a good explanation depends on one's explanatory goals. Given the explanatory goals of linguists, cognitive scientists, or ordinary people, the best explanation of the fact that we tend to say that Moore has hands may indeed be that Moore has hands. But metaphysical semantics by definition has a distinctive explanatory goal: fitting communication into fundamental reality in a certain way. Moreover, when Schaffer himself takes up something like this explanatory goal, part of what he will do is cite grounds; and he won't cite the fact that Moore has hands as an "ultimate ground".

The feature of metaphysical semantics that makes multiple realizability problematic has nothing to do with eliminativism; it is rather its "biconditional" nature. An approach to fundamentality supplies a sense in which the fundamental "underlies" the nonfundamental. If the underlying is a sort of necessary and sufficient condition then I call it and the theory biconditional; if the underlying is a mere sufficient condition, I call them conditional. The usual grounding theories are conditional. Fine's notion of ground, for instance, obeys the rule that "*p* grounds *q*" implies " $\Box(p \rightarrow q)$ ". A biconditional sort of Finean ground, on the other hand, would obey instead the rule that "*p* grounds *q*" implies " $\Box(p \leftrightarrow q)$ ";⁵³ and my account is biconditional because my metaphysical truth-conditions are biconditionals. Since Schaffer and Fine's grounding relations are merely conditional, they can say that the ground of (M) is some fact specific to actual-world physics and thus avoid the multiple realizability problem. One solution to the problem, then, would be to retreat from my demand for biconditional grounds.⁵⁴ But I dislike this response. Ground is supposed to constitute or underwrite metaphysical explanation, and merely conditional ground, it seems to me, does not do this.⁵⁵

My account's biconditional nature is not the sole source of the multiple realizability problems, since functional role metaphysical truth-conditions may be constructed. But their construction requires a metaphysical truth-condition for 'is a fundamental simple property', and that was problematic because my basic locution for talking about structure is an operator rather than a predicate. Thus the second source of the problems is my account's nominalistic nature. As we saw, if one adopted a more platonist approach—employing universals,

⁵³In neither case does the modal statement imply the statement of ground.

⁵⁴Given my arguments in section 7.9 for the superiority of a linguistic approach, the best retreat might be a linguistic form of merely conditional ground.

⁵⁵I say a bit more about this in Sider (2013*d*), and hope to say more in future work.

say, at the fundamental level—then the problem would immediately be solved. I discussed two reservations about the platonist approach in my book. First, I said, abstracta don't generally enhance explanations. But some may regard the very issues we are discussing as answering this. Second, the platonist approach makes it hard to achieve one of the central ambitions of my book, namely to “go beyond the predicate”—to speak of joint-carving for, e.g., quantifiers and sentential connectives such as ‘necessarily’, ‘or’, and ‘not’ (chapter 6). But the platonist approach doesn't *strictly* require stopping at the predicate, since a platonist *could* regard various nonpredicates as standing for abstract entities, and could regiment talk of structure using a predicate ‘carves at the joints’ applied to these entities. For instance, despite its apparent artificiality (p. 90), a platonist could regard the existential quantifier as standing for an abstract entity, the property of *being a property that has at least one instance*, and could regard the question of whether reality has quantificational structure as turning on whether this abstract entity carves at the joints.

Though each of these retreats is possible, I myself prefer a different one. I want to argue that the multiple realizability problem is primarily a problem for my account of modality, and not for my account of metaphysical truth-conditions per se, and thus that it can be solved by modifying my theory of modality (though I will not here attempt this modification).

The multiple realizability problem is generated by the assumption that metaphysical truth-conditions must obey a principle of “Fitting”, as Schaffer calls it—that they “must fit with intuitive usage by holding in most conceivable circumstances”. Now, something like Fitting is most certainly implied by my Humean theory of necessity: I say that the “axioms” of a theory of metaphysical truth-conditions are ipso facto necessarily true (section 12.9) (and it's likely that the metaphysical truth-conditions we have been discussing would count as axioms). But I don't think that Fitting is required by my theory of metaphysical truth-conditions itself.

A metaphysical semantics is supposed to explain certain linguistic phenomena, and these phenomena do include our intuitions about the circumstances in which sentences would be true. But intuitions about nonactual possibilities—especially ones very distant from ordinary thought—is not particularly central to the phenomenon of linguistic communication. Moreover, there are conflicting goals in explanation. Although some goals would be best served by attempting to characterize meaning in terms abstract enough to apply in most possible or conceivable circumstances, others are best served by sticking closer to what is actually going on in linguistic populations and their relations to

their environments, which in turn is best achieved by not attempting to have metaphysical truth-conditions apply so widely. (In fact, an actual-physics metaphysical semantics would satisfy the latter goals better than a functional-role one. We need not choose, once and for all, between the two. Since there is not always a best answer to “what best explains such-and-such?”, there is not always a best answer to “what in fundamental reality underlies such-and-such?”⁵⁶)

We still need the functional role approach even if Fitting is denied. (The arguments given above against (M)’s metaphysical truth-conditions specifying the actual physical nature of Moore’s hands or disjoining such specifications had nothing to do with Fitting.) Thus the problem of finding a metaphysical truth-condition for ‘*P* is a fundamental simple property’ remains.

It may seem that (1) could now stand as the truth-condition. For the objection to (1) given above was that it doesn’t allow for the possibility of modal variation in the fundamental properties, and that objection presupposes Fitting. But there is a further problem with (1): its disjunctive nature threatens the explanatory value of any metaphysical semantics containing it.

Instead of (1) I propose:

- (3) Simple property *P* is fundamental iff: \mathcal{S} (charge) and *P*’s members are all and only the charged things, or \mathcal{S} (mass) and *P*’s members are all and only the massive things, or...

(3) is in a sense intermediate between the two proposals considered above, (1) and (2). (1) defines ‘fundamental’ in a purely disjunctive way, without saying anything to unify the disjuncts. (2) attempts to eliminate the disjunctiveness by semantic ascent, but requires the notion ‘applies to’. (3) retains the disjunctiveness, but in a sense unifies the disjuncts by adding a claim about joint-carving to each disjunct—the claim that consists of attaching the operator for joint-carving ‘ \mathcal{S} ’ to the very predicate used in that disjunct. In my nominalist terms (3) is the closest I can get to (2).

One might worry that I cannot specify the disjuncts in (3) without using the notion of a joint-carving predicate. But the worry here concerns only my ability to articulate (3) (since I do not know what the disjuncts are). As mentioned

⁵⁶Schaffer worries in note 11 that given my assumption of “completeness”, pluralism about metaphysical semantics would lead to pluralism about what is fundamental. But this doesn’t follow. Equally good metaphysical semantic theories that agree on what fundamental reality is like could differ over *which* portion of that fundamental reality is to be assigned to a given sentence, or over whether that sentence is true—one might regard it as true while another regards it as false albeit useful.

earlier, we do not need ourselves to come up with metaphysical truth-conditions; the only question is whether such truth-conditions in principle can be given.

Thus my preferred retreat (panicked flight?) goes as follows. Offer a functional-role metaphysical semantics of (M). Accept (3) as the metaphysical truth condition for ‘Simple property *P* is fundamental’. Concede that it does not apply in all conceivable or possible circumstances, and argue that this does not undermine the explanatory status of a metaphysical semantics containing it. Modify the Humean account of necessity in some way, so as not to count all axioms from a metaphysical truth theory as being necessarily true.

References

- Albert, David Z. (1996). “Elementary Quantum Metaphysics.” In J. T. Cushing, A. Fine and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, 277–84. Dordrecht: Kluwer Academic Publishers.
- Armstrong, David M. (1978*a*). *Nominalism and Realism*, volume 1 of *Universals and Scientific Realism*. Cambridge: Cambridge University Press.
- (1978*b*). *A Theory of Universals*, volume 2 of *Universals and Scientific Realism*. Cambridge: Cambridge University Press.
- (1989). “C.B. Martin, Counterfactuals, Causality, and Conditionals.” In John Heil (ed.), *Cause, Mind, and Reality*, 7–15. Dordrecht: Kluwer.
- (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Bennett, Karen (2004). “Spatio-Temporal Coincidence and the Grounding Problem.” *Philosophical Studies* 118: 339–71.
- (2011). “By Our Bootstraps.” *Philosophical Perspectives* 25: 27–41.
- (2013). *Making Things Up*. Oxford: Oxford University Press. Forthcoming.
- Boolos, George (1984). “To Be Is to Be the Value of a Variable (or to Be Some Values of Some Variables).” *Journal of Philosophy* 81: 430–49. Reprinted in Boolos 1998: 54–72.
- (1998). *Logic, Logic, and Logic*. Cambridge, MA: Harvard University Press.

- Boolos, George S., John P. Burgess and Richard C. Jeffrey (2007). *Computability and Logic*. 5th edition. Cambridge: Cambridge University Press.
- Carnap, Rudolf (1950). "Empiricism, Semantics and Ontology." *Revue Internationale de Philosophie* 4: 20–40. Reprinted in *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd edn. University of Chicago Press, Chicago.
- Chalmers, David J., David Manley and Ryan Wasserman (eds.) (2009). *Metametaphysics*. Oxford: Oxford University Press.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*. Ithaca, NY: Cornell University Press.
- Field, Hartry (1972). "Tarski's Theory of Truth." *Journal of Philosophy* 69: 347–75. Reprinted in Field 2001: 3–26.
- (2001). *Truth and the Absence of Fact*. Oxford: Oxford University Press.
- Fine, Kit (2001). "The Question of Realism." *Philosopher's Imprint* 1: 1–30.
- (2006). "In Defense of Three-Dimensionalism." *Journal of Philosophy* 103: 699–714.
- (2009). "The Question of Ontology." In Chalmers et al. (2009), 157–77.
- (2012). "Guide to Ground." In Fabrice Correia and Benjamin Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality*, 37–80. Cambridge: Cambridge University Press.
- Goodman, Nelson (1978). *Ways of Worldmaking*. Indianapolis: Hackett.
- Hirsch, Eli (1982). *The Concept of Identity*. Oxford: Oxford University Press.
- (2011). *Quantifier Variance and Realism: Essays in Metaontology*. New York: Oxford University Press.
- Jackson, Frank (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Krantz, David H., R. Duncan Luce, Patrick Suppes and Amos Tversky (1971). *Foundations of Measurement*. New York: Academic Press.

- Lewis, David (1970). "How to Define Theoretical Terms." *Journal of Philosophy* 67: 427–46. Reprinted in Lewis 1983b: 78–96.
- (1983a). "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–77. Reprinted in Lewis 1999: 8–55.
- (1983b). *Philosophical Papers, Volume 1*. Oxford: Oxford University Press.
- (1986a). *On the Plurality of Worlds*. Oxford: Blackwell.
- (1986b). *Philosophical Papers, Volume 2*. Oxford: Oxford University Press.
- (1992). "Armstrong on Combinatorial Possibility." *Australasian Journal of Philosophy* 70: 211–24. Reprinted in Lewis 1999: 196–214.
- (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Loux, Michael J. (ed.) (1979). *The Possible and the Actual*. Ithaca, NY: Cornell University Press.
- Mundy, Brent (1987). "The Metaphysics of Quantity." *Philosophical Studies* 51: 29–54.
- Ney, Alyssa (2012). "The Status of Our Ordinary Three Dimensions in a Quantum Universe." *Noûs* 46: 525–60.
- North, Jill (2009). "The 'Structure' of Physics: A Case Study." *Journal of Philosophy* 106: 57–88.
- Plantinga, Alvin (1976). "Actualism and Possible Worlds." *Theoria* 42: 139–60. Reprinted in Loux 1979: 253–73.
- Prior, A.N. (1968). "Changes in Events and Changes in Things." In *Papers on Time and Tense*, 7–19. London: Oxford University Press.
- Putnam, Hilary (1987). "Truth and Convention: On Davidson's Refutation of Conceptual Relativism." *Dialectica* 41: 41–67.
- Quine, W. V. O. (1948). "On What There Is." *Review of Metaphysics* 2: 21–38. Reprinted in Quine 1953: 1–19.

- (1953). *From a Logical Point of View*. Cambridge, MA: Harvard University Press.
- Rosen, Gideon (2010). “Metaphysical Dependence: Grounding and Reduction.” In Bob Hale and Aviv Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*, 109–36. Oxford: Oxford University Press.
- Schaffer, Jonathan (2009). “On What Grounds What.” In Chalmers et al. (2009), 347–83.
- Shapiro, Stewart (1991). *Foundations without Foundationalism: A Case for Second-Order Logic*. Oxford: Clarendon Press.
- Sider, Theodore (2001). *Four-Dimensionalism*. Oxford: Clarendon Press.
- (2008). “Monism and Statespace Structure.” In Robin Le Poidevin (ed.), *Being: Developments in Contemporary Metaphysics*, 129–50. Cambridge: Cambridge University Press.
- (2009). “Ontological Realism.” In Chalmers et al. (2009), 384–423.
- (2011). *Writing the Book of the World*. Oxford: Clarendon Press.
- (2013a). “Against Parthood.” In Karen Bennett and Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, volume 8, 237–93. Oxford: Oxford University Press.
- (2013b). “Hirsch’s Attack on Ontologese.” *Noûs* Forthcoming. Available at http://tedsider.org/papers/hirsch_attacks.pdf.
- (2013c). “Reply to Jonathan Schaffer.” *Analysis* 73: 760–70.
- (2013d). “Reply to Kit Fine.” *Philosophy and Phenomenological Research* 88: 738–46.
- Stalnaker, Robert (1976). “Possible Worlds.” *Noûs* 10: 65–75. Reprinted in Loux 1979: 225–34.
- van Inwagen, Peter (1990). *Material Beings*. Ithaca, NY: Cornell University Press.
- Wilson, Jessica (2013). “No Work for a Theory of Grounding.” MS. Available at <http://individual.utoronto.ca/jmwilson/NWfaTG.doc>.