

Ted Sider's Solution to the Problem of Radical Semantic Skepticism

Alexander Richey

Columbia University 2015

In this paper, I exposit Ted Sider's proposed solution to the problem of radical semantic skepticism, as it is presented in his *Writing the Book of the World*,¹ and argue that it does not succeed. I begin with an exposition of the problem of radical semantic skepticism, then offer one solution to the problem and a subsequent modification to this solution. Next, in a brief interlude, I roughly characterize Sider's notion of "structure" – whose exposition and defense is the primary aim of his book – and then turn to its application to the skeptical problem. There I pose an objection to it, which, I believe, ultimately causes Sider's proposal to fail and compromises the larger agenda of his book.

The problem of radical semantic skepticism concerns the metasemantic interpretation of the reference relation, that is, the relation that connects words to their meanings. Ordinarily, linguists infer the meanings of words from a set of observations about the way a given population uses them. In English, the word "cow," for example, stands simply for *cow*. The radical skeptic challenges this inference. Since no one in this population has uttered the word "cow" before the year 2100 CE, or since no one has used the word "cow" to refer to cows that are not on planet Earth, the skeptic says that the meaning of the word "cow" should not be taken to be *cow, period*; it should be *cows before 2100 CE on planet Earth*. Obviously, this is not what we mean by "cow" – when one says "cow," under ordinary circumstances, she means *cow* without any such constraints. The problem is that, under the prevailing theory of language², according to which the meaning of a word is determined by its use, there is no way of ruling out such problematic interpretations of the reference relation.

This problem cannot be solved by gathering more empirical data. Suppose that in the year 2100 CE, members of the English-speaking population under consideration use the word "cow" to mean *cow*. Although this fact would rule out the skeptic's initial interpretation of the meaning of the

¹ Sider (2011)

² (and by the practice of field linguists and lexicographers)

word “cow,” he could easily revise his interpretation in a similarly problematic way. “Cow,” he might say, means *cows on planet Earth before 3000 CE*. Then, in the year 3000 CE, he could make a similar revision, and so on. Although more data may ameliorate these difficulties to some extent, in that they would rule out some problematic interpretations, they can never eliminate them altogether.

The question is how do we modify our account of the reference relation to subdue the skeptic?

According to one proposal, one is to give a model theoretic interpretation of a set of sentences *S*, such that “our words mean whatever they must in order for the sentences in *S* to come out true.”³ A model theoretic interpretation is an interpretation of a sentence or set of sentences *S* that adds crucial information that is missing about what the words mean. When we add this information, so that truth-values can be determined, we are said to interpret *S*. The added information is referred to as *interpretation J* of *S* and when this information makes *S* true, *S* is said to be *true in J* or that *J* is a *model of S*. For example, someone might say,

(1) He is walking his dog,

and offer an interpretation that “he” is Jones and “his dog” is Fido. This interpretation gives only what objects the expressions of (1) refer to, since (1) is a simple sentence. There are, however, several other items that might be added in interpretations of more complex sentences, such as, for example, the range or ranges of quantification, sets of initial states, and other things. Interpretations that consist of such items are called **structures**⁴ and different **structures** are used in different

³ Ibid, 24

⁴ **Structure** here is emboldened to distinguish it from Sider’s “structure,” with which it is not to be confused.

applications of model theory. Interpretations can also be relativized to time, place, and possible world.⁵

The descriptivist idea is that the meanings of the words from the sentences in *S* will be fixed by the model-theoretic interpretation of *S* on which all of its sentences turn out true. The meanings of other words and the truth-values of other sentences outside of *S* can then be inferred from the fixed meanings.⁶ Given the many formal tools the descriptivist has at her disposal, she should be able to construct a model that rules out the skeptic's problematic interpretations of our language.

There is, however, a problem with the descriptivist's account, at least as it is currently stated. The sentences in *S* do not yet provide enough "semantic glue," so to speak, to rule out the skeptic's problematic interpretations. Let (2) be the following intuitively false sentence:

(2) Some pigs have wings.

The model theoretic interpretation of (2) has a **structure** that will tell us (a) what objects are in the extensions of its nonlogical expressions (i.e. "pigs," "have wings") and (b) what class or classes its quantifiers range over (i.e. "some"). (2) will come out false iff there is no *x* such that it is in the extensions of both "pigs" and "have wings." And, conversely, it will come out true iff there *is* an *x* such that it is in the extensions of both "pigs" and "have wings." The former result is desirable, in that it accurately reflects what we mean by our words, whereas the latter result is adverse, since it does not reflect what we mean by our words and would furthermore constitute a failure of the theory. Is such an adverse interpretation possible?

⁵ Hodges (2009)

⁶ This theory is schematic in that different versions of it specify different sets of sentences. "S might, for example, be taken to include 'definitional sentences', whatever that means exactly" Sider (2011, 24).

All we need to do is assign extensions to predicates so that every sentence in S, plus [(2)] as well, turns out true. We might, for instance, begin by assigning the set of hard-boiled eggs to ‘pig’ and assigning the set of edible things to ‘has wings’. This makes [(2)] true since some hard-boiled eggs are edible.⁷

We can continue assigning meanings in this way to all of the expressions in S so that our interpretation satisfies the appropriate **structure** and gets the meanings of our words intuitively wrong. Thus we do not have enough “semantic glue” to reliably stick the proper meanings of the words “pigs” and “have wings” to them on all interpretations.⁸

One way of evading this difficulty is to modify the descriptivist account of the reference relation by appealing to the concept of Lewisian *naturalness*. We can add the constraint on our model theoretic interpretations that “predicates must stand for *natural* properties and relations in a correct interpretation.”⁹ This will rule out all of the problematic interpretations since they are less *natural* than the intuitively correct interpretations.

This modification, nonetheless, does not succeed either, because of the ambiguity of the doctrine of *naturalness*. As Sider admits:

Some [writers] regard [naturalness] as an ad hoc response to a problem, with no independent backing. Others regard it as occult metaphysics, as the postulation of an irreducible “semantic force.”¹⁰

Sider intends to answer these objections in an interesting way. He claims that he can give an account of *naturalness* in terms of his own concept, which he calls “metaphysical structure.” If this analysis succeeds, then the Lewisian modification to the reference relation would stand, and descriptivism, so modified, would subdue the semantic skeptic.

⁷ Ibid

⁸ The success of this argument “assumes that it is possible to continue selecting meanings for the nonlogical expressions in our language so that every member of S (and [(2)] as well) turns out true,” *ibid*, 25. Sider writes that “This will always be possible, except in special cases” *ibid*, and offers a further argument which I will not go over here.

⁹ *Ibid*, 27

¹⁰ *Ibid*

Up to now, I have explicated the problem of radical semantic skepticism, one potential solution to the problem, and a subsequent modification to this solution. In order to understand the next step in the dialectic, I will have to discuss Sider's notion of "metaphysical structure" in a brief interlude, and then return to its application to the problem of radical semantic skepticism.

According to Sider, the "structural" notions are notions that, metaphorically speaking, "carve reality at its joints." Imagine, for example, a universe that contains two fluids and that is divided into two halves, one of uniformly red fluid and the other of uniformly blue fluid (see fig. 1).¹¹ Now imagine that, in describing this universe, we do not use the concepts red and blue, but rather use *bred* and *rue*, which divide the universe into two halves along an arbitrary diagonal line (see fig. 2). Although we can express the exact same truths as we can with red and blue, *bred* and *rue* do not carve reality at its joints. They do not conform to what Sider calls the "structure of reality." We should use the concepts red and blue, Sider says, since their conceptual structure better matches reality's structure.¹² (Sider uses the adjective "structural" and the phrase "joint-carving" interchangeably to refer to properties of this kind.)

Structure has several other important features. It is perfectly objective, fundamental,



Figure 1




Figure 2

¹¹ Ibid, 1

¹² It is worth mentioning that this example is mostly pedagogical. Colors probably do not carve reality at its joints. "To carve perfectly, one must use the most fundamental concepts, expressing the facets of reality that underlie the colors" ibid, 5.

theoretically primitive, and, in *Writing the Book of the World*, it is used to reframe a panoply of philosophical matters in its own terms – including similarity, metametaphysics, physical geometry, laws, and, most pertinently for the present discussion, explanation.

Since structure is theoretically primitive, it cannot be defined. One is to come to understand structure by way of examples of the sort just explicated. Although the concept of structure remains vague, it will become clearer as we examine its application to the problem of radical semantic skepticism.

Let us now return to structure's application to the problem. Sider intends to give an analysis of *naturalness* in terms of structure that evades the objections enumerated above. This analysis depends on two claims. Sider first states that

(3) the reference relation is *explanatory*,

since one can explain certain facts by citing what words refer to. Since Sider understands explanation in terms of structure, it follows that

(4) the reference relation is a “reasonably joint-carving relation.”¹³

The joint-carvingness of the relation is what makes it explanatory. So, properly speaking, (4) implies (3).

The way in which (4) implies (3) is complex. Since the reference relation is reasonably joint-carving, Sider says, it is ultimately understood in terms of perfectly joint-carving notions that are

¹³ The reference-relation is “not a perfectly joint-carving relation” *ibid*, 29, because it is not fundamental. Only perfectly joint-carving notions, on Sider's view, are metaphysically fundamental.

fundamental.¹⁴ What then makes the reference relation explanatory is this connection to the fundamental. The explanatory power of the reference relation, in other words, *derives from* its connection to the fundamental, perfectly joint-carving notions. For Sider, these notions are those of physics, math, and logic, appropriately construed.

The fact that the reference relation is explanatory then allows Sider to rule out problematic interpretations of it. No reasonably joint-carving relation with a *physical* basis that plays the role of reference, Sider claims, could relate a human population to bizarre meanings because “the bizarre semantic values¹⁵ themselves,” he writes, “have no simple basis in the physical, nor do they stand in any physically simple relations to human populations.”¹⁶ Furthermore, “given any relation that does relate us to bizarre semantic values, there is surely some other relation with a simpler basis in the physical that relates us to non-bizarre semantic values.”¹⁷

The way that the reference relation’s physical basis is to make the relation explanatory might be made clearer by considering an analogous case. Suppose there is an arbitrary correlation between the motions of the planets and the performance of the S&P 500. The relation between these two phenomena is not explanatory because it has the wrong sort of basis. Neither the performance of the stock market nor the motion of the planets is explained by the relation between them. Likewise, if the reference relation has the wrong basis, it too will not be explanatory. Thus explanatory relations must have the “right sort” of bases. In the case of reference, Sider takes this basis to be the physical.¹⁸

To recap, Sider takes the reference relation to be explanatory, on the grounds that certain facts can be explained by citing what words refer to. Since explanation is understood in terms of

¹⁴ Ibid

¹⁵ “Semantic values” and “meanings” can be used interchangeably.

¹⁶ Ibid

¹⁷ Ibid

¹⁸ Although he admits that, for the reference relation, “It’s highly unclear what exactly the ‘right sort’ of basis is” Ibid, 29.

structure, the reference relation is therefore structural or joint-carving. Structural relations are explanatory in virtue of their connections to the fundamental notions, which, for Sider, are those of physics, math, and logic. The physical basis of the reference relation then rules out the problematic interpretations of the semantic skeptic because they are not physically simple. In this way, the application of structure to our account of the reference relation subdues the skeptic.

It is important to note that Sider's application of structure to the reference relation is not tethered to the descriptivist account of reference. Rather, Sider claims that structure can be used to modify various metasemantic accounts, ranging from Williams's explanation based account, to Millikan's evolutionary account. He claims that structure can solve the problem of radical semantic skepticism as it confronts each view. "A metasemantics," Sider writes, "is a metaphysical proposal about the nature of the reference relation; schematically: 'reference is a relation of such-and-such a type.'"¹⁹ For the descriptivist "reference is a relation that assigns values under which sentences in S — 'definitional' sentences, let us say — come out true..."²⁰ For Millikan, "reference is a relation that was chosen by natural selection to achieve a certain goal."²¹ The problem of radical semantic skepticism shows that, for each account of reference, there are several relations of proposed type. As we saw in the descriptivist account, there are several interpretations of the sentences S; and each interpretation yields a different reference relation. The same problem can be reproduced in other accounts. The question is how to rule out the bad relations. "The answer of [*naturalness*]," understood in Sider's way, "is to reinterpret a metasemantic proposal that reference is a relation of type T as the proposal that reference is a reasonably joint-carving relation of type T."²² In this way, faulty relations are ruled out on any metasemantics. Thus Sider's solution to the problem of radical semantic skepticism is independent of the descriptivist account and can be applied generally.

¹⁹ Ibid, 31

²⁰ Ibid

²¹ Ibid

²² Ibid

I will now argue, by means of counterexample, that a relation's being joint-carving does not imply that it is explanatory. If this is so, then Sider cannot rule out problematic relations on the grounds that they are not explanatory (and therefore not joint-carving) and the semantic skeptic need not yield.

Assume, as Sider does, that the fundamental notions are those of physics, math, and logic and consider the notion of weight. Weight is not a perfectly joint-carving notion because it is understood in terms of the more fundamental notions of mass and gravitational acceleration:

$$(5) W = mg$$

where m is mass and g is gravitational acceleration. Thus (5) expresses a reasonably joint-carving concept that is understood ultimately in terms of perfectly joint-carving concepts. Now let us define a comparative relation of weight: $W(x, y)$ ("x weighs more than y") as

$$(6) W(x, y) =_{df} \exists x \exists y [(W_x = n_1) \wedge (W_y = n_2) \wedge (n_1 > n_2)]$$

where W is the predicate for "weighs n " and n is a numerical value and $>$ denotes the "greater than" relation. We now have a reasonably joint-carving relation that is understood in terms of a reasonably joint-carving concept. The problem for Sider is that the fact that $W(x, y)$ is reasonably joint-carving does not prevent its relata from being highly non-joint-carving. Consider the following sentences:

$$(7) \text{James weighs more than Jill or John or Joe or Jack.}$$

(8) Bob weighs more than Betty and Bill and Boon.

Here we have a case of a reasonably joint-carving relation whose relata are non-joint-carving.²³ What's more, even if Sider were to specify a universe of discourse over which $W(x,y)$ is to quantify, whose elements are themselves perfectly or near-perfectly joint-carving, then sentences like (7) and (8) will still be constructible via the ordinary syntactic rules governing "or" and "and." Moreover, it would be inadvisable in such cases for Sider to alter the standard rules of syntax in order to bar us from speaking of sums or disjunctions. For these concepts figure largely in the way we understand the world. We often use the concepts of average and variance, for example, and we sometimes use these concepts to refer intelligibly to highly gerrymandered sets of individuals. Joint-carving relations must be able to have reasonably non-joint-carving relata if we are to express mundane facts about, for example, the distance between my weight and the average adult's weight. Thus we may conclude that

(9) It is possible for a reasonably joint-carving relation to have highly non-joint-carving relata.

If the counterexample of $W(x,y)$ holds – and there does not seem to be any reason to suspect that it does not – then (9) is irresistible.

So far I have argued that a reasonably joint-carving relation need not have joint-carving relata. I will now argue, on the basis of (9), that reasonably joint-carving relations need not be *explanatory*. We may infer from the generality of (9) that

(10) It is possible for a reasonably joint-carving *reference relation* to relate words to highly non-joint-carving semantic values.

²³ (7) and (8) are non-joint-carving because they contain many disjunctions ("or") and could contain many more.

We assume with Sider that

- (11) Any reference relation that relates words to highly non-joint-carving semantic values is not explanatory.

Thus, from (10) and (11), it follows that

- (12) It is possible that there is a reasonably joint-carving reference relation that is not explanatory.

Therefore, the counterexample of (12) implies that

- (4) The reference relation is reasonably joint-carving,

does *not* imply that

- (3) The reference relation is an explanatory relation.

Therefore, bad relations cannot be ruled out on the grounds that they are not explanatory and not all non-explanatory relations are not ruled out by claiming (4). The semantic skeptic is revived.

It is difficult to see how Sider might resist this conclusion. One possible move is to object to my inference from (9) to (10), that is, from the general claim that it is possible for joint-carving relations to have non-joint-carving relata to the specific claim that it is possible that the *reference*

relation have non-joint-carving relata. This objection, however, cannot be sustained. Recall descriptivism:

- (R) $R(w, x) =_{df}$ w is a word and x is the semantic value assigned to w by the interpretation I of the sentences S that maximizes truth assigned to sentences that are definitional in the linguistic population that uses w .²⁴

The problem is that (R) is reasonably joint-carving by Sider's own criteria since it is cast in terms of the highly joint-carving notions of logic and model theory. What's more, it fits Sider's proposal that "reference is a reasonably joint-carving relation of type T."²⁵ The fact that (R) fails indicates that structure, so understood, fails.

Here Sider might further object that (R) is not joint-carving in the relevant sense. Recall *modified* descriptivism: "the modified descriptivist proposal says that reference is a reasonably joint-carving relation that assigns values under which definitional sentences come out true."²⁶ Here is a relation that fits this proposal:

- (S) $S(w, x) =_{df}$ w is a word and x is the semantic value assigned to w by the interpretation I of the sentences S that maximizes the combination of (a) joint-carvingness of assigned semantic values; and (b) truth assigned to the sentences that are definitional in the linguistic population that uses w .²⁷

Unlike (R), (S) does not fail, because its possible relata are constrained by conditions (a) and (b). The condition (a) is doing the heavy lifting, so to speak, so that (S) is essentially guaranteed not to relate words to bizarre semantic values. (S), Sider might claim, is an instance of structure's proper application.

²⁴ Ibid

²⁵ Ibid

²⁶ Ibid, 31

²⁷ Ibid

But this cannot be correct. The problem with (S) is that it contains an illegitimate deployment of structure, on the grounds that condition (a) is not defended by any of Sider's prior argumentation. Indeed, the condition (a) is susceptible to just the sort of objection to which *naturalness* was and which Sider originally set out to evade. There is nothing preventing us from claiming, as we did of *naturalness* earlier, that condition (a) is ad hoc and requires the postulation of a mysterious and irreducible semantic force. The reason this is so is that the dialectic of Sider's defense of *naturalness* proceeded by giving an account of *naturalness* in terms of structure that was to rule out problematic interpretations of the reference relation on the grounds that they are not explanatory. The relation (S) does not work in this way. (S) rules out problematic semantic values by explicitly banning them, which is no different from the way the original Lewisian modification of the reference relation was to work. It is therefore susceptible to the same objections.

I hope I have said enough to show that structure does not solve the problem of radical semantic skepticism. I have given an exposition of Sider's solution to the problem and argued that this solution does not succeed on the grounds that a relation's being joint-carving does not imply that it is explanatory. If this is so, Sider cannot rule out problematic relations since they are not explanatory and the skeptic need not yield.

It is worth mentioning that our reasons for accepting structure in the first place are given by its ability to solve problems. If my argument succeeds, then there is less reason to accept structure as a legitimate conceptual tool. Structure may earn its keep by providing solutions to other philosophical problems, but its application to the problem of radical semantic skepticism should not be taken as an case for its acceptance.

Works Cited

- Hodges, Wilfrid. "Model Theory." *Stanford Encyclopedia of Philosophy*. Stanford University, 20 July 2009. Web. 3 May 2012. <<http://plato.stanford.edu/entries/model-theory/>>.
- Sider, Theodore. *Writing the Book of the World*. Oxford: Clarendon, 2011. Print.