

Epistemicism and the Liar

Forthcoming in *Synthese*

Jamin Asay

University of Hong Kong

asay@hku.hk

Abstract

One well known approach to the soritical paradoxes is epistemicism, the view that propositions involving vague notions have definite truth values, though it is impossible in principle to know what they are. Recently, Paul Horwich has extended this approach to the liar paradox, arguing that the liar proposition has a truth value, though it is impossible to know which one it is. The main virtue of the epistemicist approach is that it need not reject classical logic, and in particular the unrestricted acceptance of the principle of bivalence and law of excluded middle. Regardless of its success in solving the soritical paradoxes, the epistemicist approach faces a number of independent objections when it is applied to the liar paradox. I argue that the approach does not offer a satisfying, stable response to the paradoxes—not in general, and not for a minimalist about truth like Horwich.

1. Introduction

Epistemicism about vagueness is the view that vague predicates like ‘is bald’ and ‘is a heap’ have precise but unknowable extensions (e.g., Sorensen 1988 and Williamson 1994). If Phil is neither clearly bald nor clearly not bald, then the proposition <Phil is bald> cannot be known to be true (or known to be false), though in fact it is either true or false.¹ If epistemicism about vagueness is true, then the soritical paradoxes can be solved. The sorites paradox for baldness depends upon the premise that, for any number n , someone with n hairs is bald if and only if someone with $n + 1$ hairs is bald. Epistemicism rejects that premise, and thereby blocks the

¹ I use propositions as my principal truth bearer because they are the kind that Horwich embraces. The expression ‘< p >’ stands for ‘the proposition that p ’.

soritical argument that purports to show that someone with an indefinitely large number of hairs is bald. One immediate advantage of this view is that it need not adopt truth-value gaps for propositions that employ vague concepts, and so it can preserve classical logic without restriction, in particular the principle of bivalence (that every proposition is either true or false) and the law of excluded middle (for every p , p or it is not the case that p).

One possible extension of epistemicism is to the liar paradox. In particular, one might attempt to solve the paradox by granting that while the liar proposition

(L) <(L) is false>

has a truth value, its truth value is unknowable: it is epistemically indeterminate. This proposal, which has been adopted by Paul Horwich (2005, 2010), is dubbed ‘semantic epistemicism’ by Bradley Armour-Garb and Jc Beall (2005), who are responsible for the most thorough exposition of the view (which they do not endorse).² If successful, the epistemicist approach can boast the advantage of offering a unified solution to both the soritical and liar paradoxes. Plus, it seems to offer a fairly straightforward and simple solution to the paradox. The problem is diagnosed as being fundamentally epistemic, and so it might appear that no controversial semantic or metaphysical assumptions need to be made. However, as I shall argue, the epistemicist response to the liar paradox faces insurmountable difficulties, and is not nearly so plausible as its advocates suggest.

I begin, in the next section, by showing how epistemicists attempt to solve the standard liar paradox. Section 3 digs more deeply into the view, and explores the logical peculiarities of the epistemicist’s approach to paradox. Section 4 commences my critical discussion, and raises several objections for epistemicism when paired with Horwich’s specific brand of minimalism about truth (1998b). Section 5 then addresses some fundamental problems with the epistemicist response itself.

2. The Epistemicist Solution

The standard liar paradox runs as follows. Assume, in accordance with the principle of bivalence, that the liar proposition is either true or false:

² Horwich endorses their exposition (see Armour-Garb and Beall 2005: 85, footnote 2 and Horwich 2005: 82, footnote 10). Armour-Garb 2003 and Restall 2005 also discuss the view.

(1) $\langle \text{L is false} \rangle$ is true or $\langle \text{L is false} \rangle$ is false.

Next assume (L)'s T-biconditional (i.e., its instance of the truth schema ' $\langle p \rangle$ is true if and only if p '):

(T_(L)) $\langle \text{L is false} \rangle$ is true if and only if (L) is false.

Given the identity

(2) (L) = $\langle \text{L is false} \rangle$,

(T_(L)) is identical to

(3) (L) is true if and only if (L) is false.

From (1), (2), and (3) we may derive a contradiction. If $\langle \text{L is false} \rangle$ is true, then (L) is true and so (L) is false. If $\langle \text{L is false} \rangle$ is false, then (L) is false and so (L) is true.

Instead of rejecting liar sentences as meaningless (Grover 1977), rejecting unrestricted excluded middle (Field 2008), or just embracing the contradictory conclusion (Priest 1979), the epistemicist accepts (L) as perfectly meaningful and either true or false, and instead finds fault exclusively with (T_(L)). As a result, epistemicists accept classical logic without restriction; as Horwich notes, "we can and should preserve the full generality of the Law of Excluded Middle and the Principle of Bivalence" (2005: 82, footnote 10). So in order to dodge the contradiction, the epistemicist status of (L) is paired with some form of rejection of (T_(L)), which is the crux of the epistemicist solution. Without (T_(L)), we cannot draw the inferences necessary to generate the paradox.

Rejecting a T-biconditional like (T_(L)) is no trivial matter, especially for a minimalist about truth like Horwich. According to Horwich's minimalist theory of truth (1998b), the notion of truth is exhausted by the set of T-biconditionals, namely, the instances of the truth schema

$\langle p \rangle$ is true if and only if p .

Rather than analyzing truth in terms of correspondence with the facts, coherence, or pragmatic utility, Horwich thinks that we can give a complete theory of truth by way of the T-

biconditionals.³ But if Horwich endorses the schema without restriction, he runs straight into contradiction via $(T_{(L)})$. As a minimalist, Horwich stresses the primacy of the instances of the truth schema, but he must at the same time find some means of reigning them in, on pain of paradox. Hence, Horwich describes $(T_{(L)})$ as being an “unacceptable” instance of the truth schema. (Armour-Garb and Beall describe $(T_{(L)})$ as being an “incorrect” instance of the schema.) A crucial question facing epistemicists is whether or not they can give a plausible explanation as to why $(T_{(L)})$ is unacceptable. As I shall be arguing, the account that epistemicists have offered is unsatisfactory.

3. The Logic of Paradox

In order to fully appreciate the objections I shall raise against the epistemicist solution below, it will be worthwhile to examine some of the ramifications of the epistemicist position that have yet to be explored in any detail. So far, we have seen three main theses that form the core of the epistemicist position:

- (i) (L) is either true or false.
- (ii) It is impossible in principle to know what truth value (L) has.
- (iii) $(T_{(L)})$ is not an acceptable instance of the truth schema.

The first thesis is motivated by the epistemicist’s desire to preserve bivalence (and the legitimacy of (L)). The second thesis is motivated by the thought that we come to know “truth-theoretic” facts (i.e., facts that explicitly involve truth or falsity) by way of the T-biconditionals. For example, we come to know the truth of the truth-theoretic fact

<<Echidnas are mammals> is true>

by seeing how it follows logically from the non-truth-theoretic proposition

<Echidnas are mammals>

together with its T-biconditional

³ Though he later adds the proviso that only propositions are true (1998b: 43).

<<Echidnas are mammals> is true if and only if echidnas are mammals>.

(L), by contrast, is an indispensably truth-theoretic proposition: falsity is an indispensable component of (L). Because (L)'s T-biconditional is not an acceptable instance of the truth schema, there is no T-biconditional that we can use in conjunction with (L) to determine its truth value. So (iii) is utilized to explain (ii).

Turn now to (iii): $(T_{(L)})$ is not an "acceptable" or "correct" instance of the truth schema. These are cagey descriptions. What does it mean to say that $(T_{(L)})$ is not a legitimate instance of the truth schema? The real question of interest concerns $(T_{(L)})$'s truth value. Is it true? False? Neither? It's clear that $(T_{(L)})$ is not a theorem of Horwich's minimalist theory of truth; but the same is true of lots of propositions, such as <Echidnas are mammals>. It's also clear that we are not supposed to know that $(T_{(L)})$ is true, for if we knew it to be true, we *could* use it in whatever inference we wanted. Let us grant that $(T_{(L)})$ is not known to be true, and that it cannot be used as a premise in an argument.

What else can we say about $(T_{(L)})$, besides that it is unknown? It is, recall, the T-biconditional corresponding to (L):

$(T_{(L)})$ <<(L) is false> is true if and only if (L) is false>.

Its left-hand side is identical to the proposition <(L) is true> and its right-hand side is just (L). Horwich admits that the right-hand side is a genuine proposition, and has a truth value: that's thesis (i). The left-hand side does no more than attribute a truth value to (L). According to epistemicists, (L) does have a truth value, and so the left-hand side that says as much is equally legitimate. So both sides of $(T_{(L)})$ are, by the epistemicist's own lights, genuine propositions with truth values, though we can never know what they are. The biconditional itself, then, which is composed of two truth-valued propositions, is a perfectly genuine truth-valued proposition. Besides, given that Horwich accepts bivalence in its full generality, he must grant that $(T_{(L)})$ is either true or false. Let us grant, alongside the epistemicist, that we don't know that it's true. So even though $(T_{(L)})$ is not "acceptable" or "correct" in the sense of belonging to the minimalist theory of truth, or being known to be true, it is a perfectly legitimate proposition with a truth value.

It might be thought that $(T_{(L)})$ is not known to be true because, as with the epistemicist's stance on (L), we cannot know what truth value it has. But that would be a mistake. $(T_{(L)})$ is not known to be true because it's false. For the epistemicist, $(T_{(L)})$ is either true or false. But if it's

true, it leads straight to contradiction. So $(T_{(L)})$ is false, which is a consistent result. Consider also that $(T_{(L)})$ is identical to

$\langle(L) \text{ is true if and only if } (L) \text{ is false}\rangle,$

which is obviously contradictory if we assume alongside the epistemicist that (L) is an ordinary proposition with an ordinary truth value. Given that they uphold classical logic (and its commitments to no proposition being both true and false and no proposition being neither true nor false), epistemicists have no choice but to treat $(T_{(L)})$ as being false.

Furthermore, all of the various T -biconditionals built up from the liar and its iterations are false. In other words, the members of the series of T -biconditionals

$\langle\langle(L) \text{ is false}\rangle \text{ is true if and only if } (L) \text{ is false}\rangle$
 $\langle\langle\langle(L) \text{ is false}\rangle \text{ is true}\rangle \text{ is true if and only if } \langle(L) \text{ is false}\rangle \text{ is true}\rangle$
 $\langle\langle\langle\langle(L) \text{ is false}\rangle \text{ is true}\rangle \text{ is true}\rangle \text{ is true if and only if } \langle\langle(L) \text{ is false}\rangle \text{ is true}\rangle \text{ is true}\rangle$
 ...

are all false. We've seen already why the first in the series, $(T_{(L)})$, is false. The second T -biconditional in the series is identical to

$\langle\langle(L) \text{ is true}\rangle \text{ is true if and only if } \langle(L) \text{ is false}\rangle \text{ is true}\rangle,$

which is equivalent to

$\langle\text{Either } \langle(L) \text{ is true}\rangle \text{ is true and } \langle(L) \text{ is false}\rangle \text{ is true, or it's not the case that } \langle(L) \text{ is true}\rangle \text{ is true and it's not the case that } \langle(L) \text{ is false}\rangle \text{ is true}\rangle.$

$\langle(L) \text{ is true}\rangle$ and $\langle(L) \text{ is false}\rangle$ do not, according to the epistemicist, share the same truth value (since they contradict each other); but one of them is true (we just don't know which). So the second T -biconditional in the series is false; similar reasoning will show that the others in the series are false as well.

Consider now the liar proposition and iterated applications of 'is true' to it:

$\langle(L) \text{ is false}\rangle$

<<(L) is false> is true>
 <<<**(L) is false> is true> is true>
 <<<<(L) is false> is true> is true> is true.>**

The falsity of (L)'s T-biconditional reveals that the first two members of this second list must have opposite truth values. (I have represented this fact by printing only the first in boldface. In what follows, boldface propositions must share the same truth value, which is the opposite truth value shared by all the non-boldface propositions.) The falsity of the second T-biconditional above reveals that the second and third members of this list must have opposite truth values. One observation we can now make is that, as to be expected, 'is true' is not denominalizable for liar propositions and their truth-theoretic iterations. However, it is "doubly demoninalizable". In other words, although we can't denominalize (or renominalize) from one step of the 'is true' ladder to the next, we can denominalize from one step of the ladder to a step two steps away. (Or, more broadly, from a step on the ladder we can move an even number of steps away, but not an odd number of steps away.) Note that this is a feature shared with ordinary falsity: from <<p> is false> is false> we may infer <p> and vice versa. (The significance of this similarity will emerge in the next section.)

For epistemicists, 'is true' is evenly but not oddly denominalizable for liar propositions. A second curious fact is that 'is false' *is* denominalizable, just as truth is for non-liar propositions. Consider the series

<(L) is false>
 <<(L) is false> is false>
 <<<(L) is false> is false> is false>
 <<<<(L) is false> is false> is false> is false>.

The members of this series all share the same truth value. Given that the biconditional between <(L) is false> and <<(L) is false> is true> is false, the biconditional between <(L) is false> and <<(L) is false> is false> is true, and so the first two members of this series share the same truth value. Similar remarks apply to the rest of the series, relying on the continuous falsity of the iterated liar T-biconditionals.

For semantic epistemicists, 'is true' and 'is false' have definite logical features with respect to paradoxical propositions, though they are quite different from the features they possess with respect to non-paradoxical propositions. When it comes to ordinary propositions, 'is true' is

denominalizable and ‘is false’ is evenly but not oddly denominalizable. So ‘is true’ and ‘is false’ work in reverse for epistemicists when it comes to liar propositions. An important but heretofore unnoticed consequence of epistemicism, then, is that while truth and falsity behave differently in paradoxical settings, they nevertheless *behave*: there are consistent and stable logical roles that they play, even with respect to strange propositions like (L). We shall examine some of the consequences of this peculiar logical behavior in the sections to follow.

4. Epistemicism and Horwich’s Minimalism

I turn now to my criticisms of the epistemicist solution, first with respect to its being paired with Horwich’s overall views about truth. Though Horwich’s broader views on truth are independent of his epistemicism about (L), he is a prominent defender of both views, and so it is worthwhile to consider whether or not they form a tenable pair.⁴ In the next section I set Horwich’s views aside and evaluate the independent merits of epistemicism about the liar.

I raise five distinct objections to Horwich’s overall view: (i) his account of the unacceptability of (T_(L)) is unworkable; (ii) his explanation of (L)’s supposed indeterminacy is unprincipled; (iii) the epistemicist solution requires non-minimalist uses of truth and falsity; (iv) the minimalist view requires substantial (and troubling) revision to account for paradoxical propositions, and yet still cannot account for the most basic paradoxical propositions; and (v) the view collapses the distinction between truth and falsity.

4.1. Groundedness and acceptability

Horwich’s minimalism about truth centers on the claim that truth is a mere logical property, one that gives us expressive capabilities, but no independent explanatory power. Because our language possesses a truth predicate, we can accomplish tasks like endorsing unknown propositions (“What Sophia said is true”) or expressing otherwise infinite generalizations (“Each instance of the principle of bivalence is true”). On Horwich’s view, our understanding of ‘is true’ is constituted by our disposition to accept, *a priori*, instances of the truth schema, such as

<<Echidnas are mammals> is true if and only if echidnas are mammals>.

⁴ This pairing is also the focus of Armour-Garb and Beall 2005.

(Parallel remarks apply for how we understand ‘is false’.) As Horwich realizes, our understanding of ‘is true’ cannot require us to accept *every* T-biconditional, for some, like $(T_{(L)})$, give rise to contradiction. So the burden is on Horwich to say which T-biconditionals are relevant to our understanding of truth, and which are not. (Moreover, any semantic epistemicist bears the burden of offering some account of why $(T_{(L)})$ may be dismissed, on pain of offering an *ad hoc* solution to the paradox.)

The problem with $(T_{(L)})$, according to Horwich, is that its constituent propositions are *ungrounded*.⁵ Roughly speaking, propositions are grounded just in case they (or their negations) are entailed either by the “non-truth-theoretic” facts (that is, those facts that make no explicit reference to truth or falsity), or by those facts together with the T-biconditionals that correspond to those non-truth-theoretic facts (Horwich 2005: 81; see also Kripke 1975). For example <Echidnas are mammals> is grounded in the non-truth-theoretic facts (namely, the biological facts), and <<Echidnas are mammals> is true> is also grounded because it follows from the original grounded proposition together with its T-biconditional. If Sophia said that echidnas are mammals, then <What Sophia said is true> is also grounded, for it follows from all the aforementioned grounded propositions, together with the grounded <Sophia said that echidnas are mammals>. The liar proposition, by contrast, is not grounded; neither it nor its negation is entailed by the non-truth-theoretic facts, even together with the grounded instances of the truth schema. Because (L) is ungrounded, its corresponding T-biconditional $(T_{(L)})$ is, in Horwich’s phrase, not “acceptable” (2005: 81). As a result, we cannot assume it as a premise in the reasoning that generates the paradox.

Note that Horwich does not take on the notion of grounding in order to reveal that (L) has no truth value, as on Kripke’s account. Horwich utilizes grounding to instead show why $(T_{(L)})$ does not belong to the theory of truth, and so is unavailable in the reasoning that leads to the liar paradox. On Horwich’s view, because (L) is ungrounded, $(T_{(L)})$ is an unacceptable instance of the truth schema. (Horwich does *not* infer that because (L) is ungrounded, (L) is not true.) This distinction between grounded and ungrounded propositions is the basis of Horwich’s explanation of why $(T_{(L)})$ may be dismissed, and so is an essential component of the epistemicist’s solution to the paradox. However, this account is a problematic view of what makes a T-biconditional unacceptable, as I shall now argue. Because Horwich’s grounding-based explanation is unacceptable, the epistemicist has no principled basis for rejecting $(T_{(L)})$, and thus no tenable solution to the paradox.

⁵ See Simmons 1999 for independent worries about this proposal.

The problem with Horwich’s grounding explanation is that it also renders perfectly acceptable T-biconditionals as being unacceptable. The epistemicist argues that $(T_{(1)})$ is unacceptable, and I demonstrated earlier that it is, in addition, false by the epistemicist’s lights. But not all T-biconditionals formed from ungrounded propositions are false. Some of those T-biconditionals are provably true, and so it would be *ad hoc* to dismiss them as unacceptable and unavailable to use in inferences on the basis of their being constituted by ungrounded propositions. Consider the ungrounded truth-telling proposition and its ungrounded T-biconditional:

- (T) $\langle (T) \text{ is true} \rangle$
 $(T_{(T)})$ $\langle \langle (T) \text{ is true} \rangle \text{ is true if and only if } (T) \text{ is true} \rangle$.

Because $(T) = \langle (T) \text{ is true} \rangle$, $(T_{(T)})$ is identical to

$\langle (T) \text{ is true if and only if } (T) \text{ is true} \rangle$.

This biconditional is a harmless tautology. It is simply a biconditional between a particular truth-valued proposition and itself. The epistemicist must claim not to know which truth value (T) has, but cannot deny that the biconditional between (T) and itself is true. But this biconditional just is (T) ’s T-biconditional. So $(T_{(T)})$ is true, in spite of its constituents being ungrounded.

On Horwich’s view, $(T_{(T)})$ is to be regarded as an unacceptable T-biconditional because its constituents are ungrounded. As we have seen, “unacceptable” T-biconditionals are unavailable for using in inferences, despite their being truth-valued. Yet we may see that $(T_{(T)})$ is provably true and causes no paradox. One wonders, then, on what grounds $(T_{(T)})$ may be considered an incorrect instance of the truth schema. If Horwich simply insists that $(T_{(T)})$ is unacceptable because (T) is ungrounded, then his preferred restriction on the truth schema is *ad hoc*: T-biconditionals with ungrounded constituents are rejected not because they are independently problematic (some are not), but because some (though not all) of them lead to paradox.⁶ Consequently, if indeed certain T-biconditionals like $(T_{(1)})$ are unacceptable, the notion of grounding is not useful in accounting for why.

⁶ Further problem cases for Horwich include (1), $\langle (1) \text{ is neither true nor false} \rangle$, and (2), $\langle (2) \text{ is either true or false} \rangle$. Because of bivalence, we know that (1) is false and (2) is true, though they and their T-biconditionals are ungrounded. The T-biconditionals for (1) and (2) are provably true, given bivalence, so they too constitute a

4.2. Indeterminacy

The truth-teller upsets the plausibility of using groundedness as the criterion to demarcate the acceptable T-biconditionals from the unacceptable ones. This is a significant problem for Horwich—and for any other epistemicist who might also appeal to groundedness—as it undermines his case for offering a principled reason for rejecting (T_(L)). But the truth-teller causes further problems, for it also upsets Horwich’s explanation of (L)’s epistemic indeterminacy. If Horwich cannot offer an adequate explanation as to why (L) is epistemically indeterminate, then the epistemicist solution loses much of its initial plausibility.

With respect to vagueness, Horwich argues that propositions like <Phil is bald>, where Phil is a borderline case of baldness, are epistemically indeterminate because the regularities of use that underlie our practice of wielding ‘bald’ yield no verdict either way with respect to Phil and baldness (1998a: 64). But epistemic indeterminacy for the liar proposition is different, according to Horwich; this second form of indeterminacy exists in cases where the regularities of use give rise to *conflicting* inclinations. The regularities in question here are those that give meaning to our use of ‘true’ and ‘false’, namely, our acceptance of the T-biconditionals (as well as our use of negation). So the regularities that fund (L)’s meaning give rise to conflicting inclinations. Because of the conflict, Horwich says, (L) is epistemically indeterminate.

The problem for Horwich’s explanation of (L)’s indeterminacy is that it fails to explain (T)’s indeterminacy. No conflict is generated by considering the regularities of use that govern (T)’s constituent notions. As we have seen, (T)’s T-biconditional is provably true. Reasoning with (T) leads to no paradox: if it’s true, then it’s true, and if it’s false, then it’s false. But (T) has a truth value, according to the semantic epistemicist. What the epistemicist can’t do is explain (T)’s indeterminacy the way that Horwich explains (L)’s indeterminacy. This is unfortunate, given that it’s generally desirable to handle (L) and (T) similarly. Such a desideratum is defeasible, of course; but given that (T) and (L) share a self-referential truth-theoretic nature, a unified account of their indeterminacy is preferable to a disjunctive one. Horwich, by contrast, is forced to admit that there are two distinct kinds of indeterminacy when it comes to ungrounded, self-referential truth-theoretic propositions (both of which are distinct from the indeterminacy relevant to understanding vagueness).

Because the indeterminacy belonging to the liar is different from the kinds of indeterminacy active with the truth-teller and vagueness, the major worry about Horwich’s

challenge to the integrity of Horwich’s criterion of groundedness. (They also show that Horwich cannot maintain that all ungrounded truth-theoretic propositions are unknowable.)

explanation of (L)'s indeterminacy is that it (like his explanation of its T-biconditional's unacceptability) is *ad hoc*. But in their assessment of Horwich's position, Armour-Garb and Beall argue that, to the contrary, it is a *virtue* of Horwich's explanation of the indeterminacy that it is not *ad hoc* (2005: 91-92). Their argument is that (L)'s indeterminacy falls right out of Horwich's general account of indeterminacy. Because the meaning-giving regularities of use behind (L) conflict, (L) is epistemically indeterminate. But it strikes me that developing a unique form of indeterminacy that is specific to the liar (and doesn't even extend to the truth-teller), and then explaining the indeterminacy of (L) in terms of it does little to remove the suspicion that the maneuver is *ad hoc*. Put another way, Horwich's general account of indeterminacy looks to be a highly disjunctive account that mashes together several disconnected and independent sources of indeterminacy, and a highly disjunctive account is the basis of the charge of being *ad hoc*, not a solution to it.

One way to press this worry is to ask why the response to the tangle caused by the conflicting regularities behind 'true' and 'not' (and, in turn, 'false') is to diagnose (L) as being indeterminate rather than, say, meaningless or contradictory. Horwich thinks that indeterminacy exists in cases of vagueness because "the regularities for the use of a [vague] predicate say nothing and imply nothing about its application to a given object; they can yield no inclination either to apply or to withhold the predicate" (1998a: 64).⁷ But Horwich is explicit that the indeterminacy relevant to (L) is unique, and is due to the fact that the regularities behind (L)'s constituents are not *silent* but *conflicting*. But from the fact that the regularities that govern 'not', 'true', and 'false' conflict we cannot just infer that propositions involving the three have settled but unknowable truth values. We might just as well conclude that (L) is contradictory, or that conflict leads to *lack* of truth value. We need independent reason to infer *epistemic indeterminacy* from the conflict, and not some other status (such as semantic indeterminacy or contradiction).

In fact, there is some pressure on Horwich, given his own views about meaning, to judge (L) to be meaningless, or at least semantically defective in such a way as to leave it without a truth value. For Horwich, the meanings of 'is true' and 'is false' are constituted by the regularities that govern their use; in particular, they are governed by the T-biconditionals (and negation). It's the T-biconditionals that provide the meanings for 'is true' and 'is false'. But because Horwich rejects the T-biconditional corresponding to (L), he cannot appeal to it when explaining the

⁷ This account of indeterminacy *does* seem to provide an account of (I)'s indeterminate status. So to reframe the objection from the previous paragraph, it would be odd if the explanation of (I)'s indeterminacy closely paralleled Horwich's account of indeterminacy for vagueness, but resembled the indeterminacy of (L) not at all; here we have further (yet still defeasible) reason to think that Horwich's explanation of (L)'s indeterminacy is *ad hoc*.

meaning of ‘is false’ as it appears in (L). So a more natural view for one who both takes ‘is false’ to be given its meaning by the T-biconditionals (and negation) and rejects the T-biconditional corresponding to (L) would be that it’s meaningless, since one of its key components is robbed of its typical meaning. There are, of course, independent reasons to reject the view that (L) is meaningless.⁸ But the present worry asks on what grounds *Horwich* can resist the thought that (L) fails to be meaningful, and so why his choice to take it to possess a unique variety of epistemic indeterminacy is not *ad hoc*. Given that the usual grounds for understanding (L) do not account for its meaning, what reason is there to suppose nonetheless that it is meaningful and has a truth value? What is supplying (L) with its meaning, given that the usual channels are ruled out?⁹

4.3. Two kinds of truth and falsity

On a minimalist conception of truth and falsity that accepts bivalence, to say that a proposition is false is to assert its negation (Horwich 1998b: 71-73). We can codify this principle with a “falsity schema” (which is equivalent to the truth schema):

<*p*> is false if and only if it is not the case that *p*.

We understand predications of falsity, so say minimalists, by way of such biconditionals. But this line cannot be extended into the domain of paradox. The instance of the falsity schema corresponding to (L),

<<(L) is false> is false if and only if it is not the case that (L) is false>,

is false, for it is equivalent to $(\Gamma_{(L)})$. As we saw in section 3, regardless of how ‘false’ operates in non-paradoxical settings, it behaves differently when it comes to (L). Falsity with respect to (L) cannot amount to the truth of its negation. But if falsity is doing something else in paradoxical settings, it appears that there are *two* notions of falsity at hand, and not just one. The epistemicist, then, is committed to a non-minimalist conception of falsity.

⁸ Most prominently, it invites the “strengthened” paradox due to (S), <(S) is not true>, discussed below. If (L) is meaningless, so too is (S), in which case it’s not true, and so appears to be true after all.

⁹ Cf. Beall 2001, which argues that deflationists have independent reason to claim that liar sentences are meaningless, given their ineliminable use of ‘true’ and ‘false’.

Armour-Garb and Beall (2005: 93-94) make a parallel observation with respect to the “strengthened liar” proposition:

(S) <(S) is not true>.

The strengthened liar, according to the epistemicist, is either true or false. Consider that first disjunct:

<<(S) is not true> is true>.

Here we have two appearances of ‘true’. Armour-Garb and Beall argue that Horwich’s pairing of minimalism and epistemicism is unstable because it faces the following dilemma. Either both appearances of ‘true’ are minimalist, or not. If not, then minimalism is false, for there is more to truth than what the minimalist theory offers. If both appearances are minimalist, however, then they behave in standard minimalist fashion as devices of denominalization and renominalization (which is to say, their T-biconditionals are acceptable), and so we can derive a contradiction in the familiar way.

Armour-Garb and Beall think that the above dilemma proves Horwich’s package view to be unstable. Horwich definitely must reject the claim that ‘true’ and ‘false’ as they function in the liar and its strengthened cousin operate as they normally do; this response is enshrined in his rejection of their T-biconditionals, for taking ‘true’ and ‘false’ to operate normally just is to accept the relevant T-biconditionals. The tentative suggestion of the previous section was that this admission should not lead us to think that liar propositions are epistemically indeterminate, but rather that they are meaningless. What (L) means depends upon what ‘false’ means, but ‘false’ doesn’t appear to mean *anything* for a minimalist in cases where the truth and falsity schemas don’t apply. So it seems that Armour-Garb and Beall’s dilemma should be expanded: Horwich must either grant that ‘false’ is minimalist in (L), which means embracing contradiction, reject (L) as being meaningless (or, in effect, not being a real proposition), thereby abandoning epistemicism, or take there to be more to ‘false’ and ‘true’ than is given by the truth and falsity schemas, giving the lie to minimalism.¹⁰

¹⁰ Armour-Garb and Beall also consider a further objection to semantic epistemicism, namely, that it faces a revenge paradox via (K): <No one knows that (K) is true>. Since it is epistemically indeterminate, (K) is a consequence of semantic epistemicism. And so, Armour-Garb and Beall claim, (K) is known by semantic epistemicists, resulting in contradiction. However, we cannot take as a premise that semantic epistemicists know that their theory is true, even

4.4. F-biconditionals and the indispensably truth-theoretic

The appearances of ‘false’ in $\langle(L) \text{ is false}\rangle$ and ‘true’ in $\langle(S) \text{ is not true}\rangle$ are not garden variety. The requisite T-biconditionals are not available to give them their typical meaning. I have been suggesting (alongside Beall 2001) that perhaps minimalists should take liar propositions to be meaningless as a result. But there is another option worth exploring, which draws on the investigations from section 3. There, we uncovered the logic of liar propositions for semantic epistemicists. In particular, we saw that ‘is false’ is a device that enables denominalization and renominalization (the same holds for ‘is not true’). As a result, for that class of propositions whose T-biconditionals are false, we can derive true instances of “F-biconditionals” that conform to the following schema:

$\langle p \rangle$ is false if and only if p .

This schema (whose instances are contradictories of the corresponding instances of the T-schema) makes explicit the surprising fact uncovered earlier that when it comes to epistemicists’ take on paradoxical propositions, falsity logically behaves just like truth for ordinary propositions.

We may now appreciate that Horwich’s minimalist theory of truth requires substantial revision, given his commitment to semantic epistemicism. The minimalist’s challenge is to give a full account of how truth operates by way of the T-biconditionals. To avoid paradox, T-biconditionals like $(T_{(L)})$ are excluded. One necessary supplement to the theory, the need for which Horwich has long recognized, is an account of which T-biconditionals are excluded. (I have already argued that (T) undermines the thought that the exclusion can be defined in terms of ungroundedness.) A second necessary supplement, which Horwich has not recognized, is something to account for how ‘true’ and ‘false’ operate with respect to liar propositions. Given that Horwich thinks that (L) and (S) (among countless others) are genuine, meaningful propositions, he owes us an account of how truth and falsity work for them. Without that

if it is; they might fail to know either because the theory is false, or because their belief in it fails to achieve the status of knowledge. The most one can derive from (K) is that if semantic epistemicism is true, it cannot be known. This result is not contradictory; it just shows that those who go in for semantic epistemicism had better not take themselves to know the truth of their view. This response to Armour-Garb and Beall, however, provides little comfort to the epistemicist, as other revenge problems are available. Consider the similarly ungrounded (J): $\langle \text{It is known that (J) is false} \rangle$. Supposing (J), it is known that (J) is false, and so (J) is false. By *reductio*, (J) is false, and knowably so since it’s provably false. So (J) is known to be false, in which case (J) is true. Thanks go to a referee for identifying this further revenge problem.

account, the presence of falsity in (L) and truth in (S) go unexplained, and Horwich's theory has a significant lacuna.¹¹

What Horwich should do in order to supplement his minimalism so as to account for the behavior of liar propositions is add in the corresponding set of F-biconditionals. When it comes to liar propositions, 'false' acts in the manner of 'true' for ordinary propositions, and 'true' functions like 'false' for ordinary propositions. We explain the role of 'true' in <<Echidnas are mammals> is true> by way of the T-biconditional <<Echidnas are mammals> is true if and only if echidnas are mammals>. Similarly, we explain the role of 'false' in <<(L) is false> is false> by pointing to the F-biconditional <<(L) is false> is false if and only if (L) is false>. So every time Horwich strikes a T-biconditional from the minimalist theory, he should replace it with its corresponding F-biconditional.

The resulting supplemented minimalist theory is as follows. The theory is given by a set of biconditionals. Some take the form

<<*p*> is true if and only if *p*>

and others take the form

<<*p*> is false if and only if *p*>.

There are infinitely many instances of each type, and no proposition has a true instance of both schemas. Which propositions are to be paired with which schema? The answer can't just be that the paradox producing propositions instantiate the falsity schema truly, and the non-paradox producing propositions instantiate the truth schema truly, on pain of being artificial and *ad hoc* as a solution to the liar paradox. The answer can't be that grounded propositions truly instantiate the truth schema, and ungrounded propositions truly instantiate the falsity schema, for (T) is a counterexample. Exactly parallel remarks hold for taking epistemic determinacy/indeterminacy

¹¹ Cf. Simmons 1999: 463-464. To reply that propositions like (S) and (L) are "few and far between" (Horwich 1998b: 42, footnote 21) does little to assuage the objection. Liar propositions are not few in number: there are infinitely many of them. Even one truth-theoretic proposition that cannot be explained by minimalism is sufficient to serve as a counterexample to the minimalist's claim of having accounted for all the facts about truth. As with much in philosophy, no tensions arise in normal cases; it's the extraordinary cases that fuel philosophical interest, and that lead us to question the very principles that we think lie behind the ordinary cases. To say that truth works in the ordinary way in ordinary cases does little to remove the worries that arise from observing how the ordinary begins to act extraordinarily in extraordinary cases.

to draw the distinction, as (I) is indeterminate but has a true instance of the truth schema. Moreover, vague propositions, for Horwich, are indeterminate, but still have true instances of the truth schema. Finally, the answer cannot be that propositions whose meaning-conferring regularities of use conflict, for we now have supplied consistent meaning-conferring principles (the F-biconditionals) for liar propositions (otherwise we leave them meaningless and unexplained). It remains to be seen how Horwich can adequately funnel propositions into the schema to which they belong.

Adding an infinity of F-biconditionals to the minimalist theory of truth may be too much to swallow for many. But suppose that the minimalist bites the bullet and accepts this amendment to the view. There remain still further problems. The supplemental F-biconditionals help explain the role that ‘true’ and ‘false’ play in *most* liar propositions. But there are some liar propositions where they still don’t help, namely, (L) and (S). The problem is that these liar propositions are at the base of the liar sequence, and yet still have truth and falsity as a part of them. In ordinary cases, truth-theoretic propositions “bottom out” in non-truth theoretic propositions:

...
 <<<Echidnas are mammals> is true> is true>
 <<Echidnas are mammals> is true>
 <Echidnas are mammals>.

But the same does not hold of propositions like (L) and (S) (and (I) for that matter); truth (or falsity) is still present even at the ground level. This fact is, of course, utterly familiar, and is the key insight behind the idea that the liar and truth-telling propositions are ungrounded. But because the bottom-level ungrounded propositions are truth-theoretic, nothing in the minimalist theory, even when radically supplemented, can be used to explain what they mean.

According to my offering to the minimalist, we can make use of the T- and F-biconditionals for iterated versions of the liar and the truth-teller in order to account for their use of ‘true’ and ‘false’. But this is not so for the original liar and truth-teller themselves. With respect to (I), I have argued that Horwich needs to revise his view so as to include its T-biconditional in the minimal theory. (He unwarrantedly must exclude it as unacceptable because (I) is ungrounded.) But we can’t use $(I_{(I)})$ to explain (I)’s use of ‘true’. $(I_{(I)})$ is the proposition that $\langle(I) \text{ is true} \rangle$ is true if and only if (I) is true. In other words, $(I_{(I)})$ is the proposition that (I) is true if and only if (I) is true. $(I_{(I)})$ takes us nowhere; it just directs us from (I) back to (I)

again. Hence, (T_(T)) doesn't explain the role of 'true' in (T), as ordinary T-biconditionals do for their constituent propositions. Similarly, we cannot explain the role of 'false' in (L) because (L)'s F-biconditional holds that <(L) is false> is false if and only if (L) is false. Or, in other words, (L) is false if and only if (L) is false. In effect, the presence of truth and falsity in bottom-level propositions like (L), (S), and (T) must go unexplained on the minimalist theory. These propositions are indispensably truth-theoretic, and so go unaccounted for by minimalism, even in its radically supplemented version.

4.5. Triviality

As we have seen, central to Horwich's theory of meaning is that words are given their meaning by the existence of regularities that involve them. Applied to truth, the theory holds that "the acceptance property governing our total use of the word "true" is the inclination to accept instances of the schema 'the proposition *that p* is true if and only if *p*'" (1998a: 45). But a further dimension of Horwich's theory that is particularly relevant is his commitment to the view that two words express the same concept in virtue of having the same basic acceptance property (1998a: 46). If two words share the same basic acceptance property, then they have the same meaning. For Horwich, if we accept the instances of the schema 'the proposition *that p* is X if and only if *p*', then 'X' means the same thing as 'true'.

So consider that set of propositions, including but not limited to (L) and (S), the members of which have true F-biconditionals. In section 3, we observed that these propositions behave in a peculiar fashion. In particular, we found that there is a property such that, when they have it, denominalization can result. We saw that there is a property X such that any proposition <*p*> from that set has X if and only if *p*. X is a property that enables denominalization and renominalization. Now, we know that X is falsity. But X behaves exactly like truth does for all the other propositions not in the set. So it appears that the meaning of 'false' in <(L) is false> is identical to the meaning of 'true' in <<Echidnas are mammals> is true>, by Horwich's own lights. Same acceptance property, same meaning.

The conclusion one might now draw is that (L), despite appearances, turns out not to mean that (L) is false, but that (L) is true (since (L) says that it has the denominalizing property, which is what truth is). But when we first defined (L), we didn't define it as the proposition that declares its own special brand of "falsity" that is distinct from normal falsity (and is really just truth). If that were correct, "(L)" would really have just been (T), the truth-teller. But this is not the case: (L), by pure stipulation, is the proposition that declares itself to be *false*. But we now

seem to have shown that (L), which means that (L) is false, means that (L) is true (since the property (L) ascribes to itself is the same acceptance property that ‘true’ expresses in non-paradoxical settings). In other words, for Horwich, truth *is* falsity. Hence the view appears to suffer from triviality. Because of bivalence, every proposition is true or false. And since being true is the same as being false, every proposition is both true and false.

Let’s approach this argument from a slightly different angle. What is it to say that <Echidnas are mammals> is true? For Horwich, in effect, it is to say that the proposition has the denominalizing property, that is, that property D that enables one to move from <<Echidnas are mammals> is D> to <Echidnas are mammals>. Now, what is it to say that (L) is false? Well, (L) is <(L) is false>, and to say that <(L) is false> is false is, as we saw in section 3, to say that <(L) is false> has the denominalizing property D. <(L) is false> is false if and only if (L) is false; the F-biconditional for <(L) is false> is just as acceptable as the T-biconditional for <Echidnas are mammals>. So to say that <(L) is false> is false is to say that <(L) is false> has D, the same property we attribute to <Echidnas are mammals> when we say that it’s true. But to say that (L) is false is also just to express (L) again. And (L) expresses the claim that (L) *lacks* D: that’s why (L) is the *liar*, and not the truth-teller. Hence, to say that <(L) is false> is false is to attribute D to it. But to say that (L) is false is to deny that (L) has D: that’s what makes (L) the (lying) proposition that it is. Of course, (L) just is <(L) is false>, and so the only way to avoid an outright contradiction here is to collapse the distinction between truth and falsity.

5. Problems for Semantic Epistemicism

Epistemicism about the liar is particularly untenable for Horwich. In this section, I set his views about truth and meaning aside and raise two objections to the epistemicist approach to the liar itself.

5.1. Expressive incoherence

Phil and Sophia are epistemicists. For whatever reason, they have found themselves faced with a Pascalian wager with respect to (L): granted, they cannot know the truth value of (L), but they must choose nevertheless. Whether eternity is filled with paradise or damnation depends upon their making the right decision. Phil and Sophia diverge; Phil takes (L) to be true, and Sophia takes it to be false. How their choice will be rewarded is something they cannot know until it’s too late. Time passes, Phil and Sophia die, and now you, the eternal gatekeeper,

are responsible for sending Phil and Sophia to their eternal resting places. To figure out who won the bet, you consult the Big Book of Truth. Skipping down to 'L', you notice the following on the list:

<(L) is false>

<<(L) is false> is false>

Just to be sure, you check the Big Bad Book of Falsity. There you find:

<(L) is true>

<<(L) is false> is true>

After a moment's reflection, you sympathetically shrug your shoulders at Phil and hand over the keys to paradise to Sophia. "Hold on," Phil says, "the book says I'm right! Give me the keys!" "Not so fast," Sophia retorts, "those keys are rightfully mine." Who is right?

Phil is right. Phil bet on (L) being true. In other words, Phil bet on (L) being among the truths in the world. To say that (L) is true is to say that (L) belongs in the complete, exhaustive theory of the world. And there it is, plain as day, in the Big Book of Truth. Sophia's view, meanwhile, is clearly wrong. She said that (L) was false, and yet it appears in the Big Book. So much the worse for Sophia.

Sophia is right. She bet on (L) being false. That was her claim: (L) is false. And lo and behold, there her claim appears, right in the Big Book. The Big Book says that (L) is false, and the Big Book only reports truths. Phil's view, meanwhile, shows up right in the Big Bad Book of Falsity. Phil said that (L) was true, and *that very claim* is reported in the Big Bad Book. So much the worse for Phil.

One of the responses to the story must be correct, but there is no reason to prefer one over the other. Even equipped with the complete list of truths, we still can't figure out the truth value of (L). In fact, the story of Phil and Sophia brings out what I would deign to call a deep expressive incoherence in the epistemicist position: the two contrary positions it creates (taking the liar to be true, and taking it to be false) appear to express the other. According to epistemicism, the liar proposition has a definite truth value, and so we all have a fifty percent chance of guessing it correctly. For those who guess 'true', they will not be able to express that commitment with '(L) is true', which you might have thought was the natural vehicle for committing yourself to the truth of (L). (L) is **<(L) is false>**, which I have printed in boldface to

indicate that I have guessed that it belongs among the other propositions I take to be true, such as **<Echidnas are mammals>**. I cannot explain why I have printed (L) in boldface by saying that it's in boldface because (L) is true, because then I would be explaining my commitment by saying that (L) is true, which by my own lights is false: if (L) is true, then <(L) is true> is false. As a result, if I believe that (L) is true, I can't assert that (L) is true, for then I would be asserting something (namely, that (L) is true) that I take to be false. In short, by severing the equivalence between <*p*> and <<*p*> is true> for liar propositions, epistemicists make it impossible to coherently express one's commitments that one might take with respect to these two propositions.

The epistemicist might respond to this line of reasoning by either accepting its conclusion—claiming that the “expressive incoherence” I am identifying is simply part and parcel of the view—or pointing out that the supposed incoherence reveals that my initial setup was flawed: there can be no “Big Book of Truth” that could correctly tell you the truth value of (L), precisely for the reasons epistemicists have identified. After all, (L)'s truth value is unknowable *in principle*.

As for the first response, it constitutes a substantial concession, and severely complicates the epistemicist's response to the paradox that initially seemed quite simple and straightforward. At the least, it shows that the epistemicist status belonging to the liar is strikingly different from the epistemicist status belonging to vague propositions. The liar is not merely unknowable; in effect, it's inexpressible as well. As for the second response, epistemicists should be careful not to adopt too strong a notion of unknowability for paradoxical propositions. The problem with liar propositions was initially supposed to be purely epistemic; that was part of the initial attraction of the view. But to deny that there could be a Big Book is to extract metaphysical consequences from epistemicism. Denying the Big Book is to deny that there could be a list of all the truths, truths whose existence the epistemicist is committed to. Furthermore, it seems tantamount to denying that God can be omniscient (or that God could tell us about the truth value of (L)), since the Big Book just is a stand-in for an omniscient being. Perhaps the epistemicist might deny that even God knows the truth value of paradoxical propositions, but doing so is again a significant concession.

5.2. Truthmakers for liars

Epistemicists believe that there is a lot of truth in the neighborhood of paradox. They don't know what the truths are, but they believe that they exist. Suppose it's a fact that (L) is true.

The previous section asked, in effect, whether that fact was (L) or <(L) is true>, since it can't be both (and neither seems to have any more claim on being that fact than the other). This section asks what it is that *makes* (L) true, if it is true (or what makes (L)'s negation true, if it's the true one).

The theory of truthmaking, as I understand it, is an attempt to explain the idea that propositions are true in virtue of reality. Because the world is a certain way, certain propositions are true, and certain propositions are false. Exactly how we should flesh out truthmaker theory and its guiding intuition is a matter of some controversy,¹² but for present purposes let us say that to ask what makes a proposition true is to ask what it is in or about the world that accounts for the truth of that proposition. The proposition <Echidnas exist> is true because of the actual echidnas in the world. If those animals didn't exist, the proposition wouldn't be true. Given that the epistemicist thinks that one of <(L) is false> and <(L) is true> is true, it is a fair question to ask: whichever one is true, what is it that makes it true?

I raise the question because it strikes me that there is no good answer to it. Suppose that (L) is the true one. No contingently existing entity in the universe seems to be responsible for (L)'s truth. I cannot imagine what would have had to have gone differently to result in (L)'s not being true. God, it seems, would have no reason to prefer one of the proposition's being true over the other.¹³ The choice would have no bearing on the rest of the world. Epistemicists about vague propositions, by contrast, might point to our patterns of use that govern vague predicates, and argue that those patterns (together with the normal empirical facts) make vague propositions true. (In that case, what explains our ignorance is that the patterns of use are so complicated that we cannot know what truth values they produce.) But even that option is unavailable to epistemicists about the liar. There's nothing about our use of 'true' and 'false' that would settle the matter either way. It's not as if, had we used 'true' in a slightly different way, then (L)'s negation rather than (L) would have turned out to be true. (By contrast, if we used 'is a heap' in a slightly different way, so says the epistemicist, then the privileged number that constitutes a heap for a given context would be slightly different.) The prospects for finding contingent truthmakers for liar propositions are grim; they are even more dire than those facing epistemicists who must offer truthmakers for propositions like <*n* is the precise number of grains of sand that is sufficient to constitute a heap in context C>. One can always posit the

¹² See Armstrong 2004 for an introduction to truthmaker theory.

¹³ Cf. Horgan 1997.

existence of a new entity to serve as a truthmaker, but the posit needs to be independently motivated, lest it be *ad hoc*.¹⁴

If there are no contingent truthmakers available, how about a necessary one? If there were such a thing, then it would seem that liar propositions have their truth values necessarily. If (L) is true, then it's necessarily true, for that which is responsible for its truth necessarily exists. But it is equally unclear what necessary existents, if there are such things, would serve as plausible truthmakers for true paradoxical propositions. (It would also be a most unwelcome consequence for epistemicism that it be committed to such a startlingly contentious ontological commitment from the outset.) Perhaps the most natural answer is that the paradoxical propositions themselves make themselves true. If (L) is true, then it makes itself true. I have no objection to the idea of self-truthmaking; <There are propositions> is an innocuous example. My problem is rather with the lack of any explanatory value in saying that (L) makes itself true. We look to truthmakers to explain the truth of propositions, but nothing about the *existence* of (L) favors its being true rather than false. Saying that (L) makes itself true is no better than saying that nothing explains why it's true, and that its truth is just some brute fact in the universe.

Now, in asking after (L)'s truthmaker (if indeed (L) be true), it might appear that I am presupposing that *every* truth has a truthmaker. In fact, I do not make that assumption, alongside several other truthmaker theorists (e.g., Lewis 2001). But while I do not believe that every truth has a truthmaker, I do believe that there are no absolutely brute truths. An absolutely brute truth is one whose truth can in no way be explained. One way of understanding the brute fact view is by focusing on how it denies the plausible claim that what's true supervenes on what exists, and what properties those things have (see Lewis 2001). This thesis is offered as a more economical way of capturing the core intuition behind truthmaker theory; it can be accepted even by those who, like Horwich (2008), are generally skeptical of the overall truthmaking program. Had a proposition that is true failed to be true, then there must have been some difference in the world regarding what exists, or what properties those existing things had. But if (L) is true (or false), its

¹⁴ For example, supposing that (L) is true, one might posit the existence of an entity called "the fact that (L) is true", which wouldn't have existed had (L) been false. In this case, the "fact" can't just *be* the true proposition itself, but rather some separate entity that *makes* (L) true. But what is this fact composed of? Do we have any reason to believe in its existence independently of a commitment to epistemicism? The challenge for the epistemicist is to offer a truthmaker that is sufficiently explanatory while not being an *ad hoc* ontological posit.

truth (or falsity) doesn't appear to supervene on the rest of reality, thus violating what is otherwise a plausible claim that holds true for seemingly all other truths.¹⁵

Those moved at all by truthmaker theory—including even those who are not convinced that all truths have truthmakers—will be unmoved by the thought that there are some truths whose truth is an absolutely fundamental and brute matter. The suspicion toward brute truths is not due to an overall suspiciousness of taking things as brute. The question is whether the *truth* of certain propositions is ever a plausible candidate for being a brute, foundational feature of the universe. To illustrate, suppose that there are some brute, fundamental sub-atomic particles. To say that those particles are brute is not to say that the truth of the proposition <'Those brute sub-atomic particles exist> is brute; the proposition's truth is not brute because it depends upon the existence of the particles themselves. Those who take (L) to be a brute truth sidestep a genuine explanatory question that uniquely faces epistemicists: why is (L) true rather than false?

It might be thought that explanations have to stop somewhere, and so stopping with (L) is as good a stopping place as any (or, at the least, a potentially defensible stopping place). But remember: truthmaker theorists agree that explanations have to stop somewhere. It's just that stopping at the level of the truth of propositions is not a suitable resting place; we need to explain truths in terms of reality. Explanations stop at the world, not at the truth of propositions. Furthermore, when we do land on a brute stopping point for explanations, we do so only because that stopping point proves to be explanatorily fruitful.¹⁶ But what explanatory virtues follow from taking (L) to be brutally true? Taking (L) to be brute is not analogous to taking there to be, say, basic beliefs in epistemology, beliefs that confer justification without requiring that they be justified by still further beliefs (see, e.g., Bonjour 1978). Basic beliefs are posited by foundationalist theories of justification because they serve to explain what justifies the rest of our beliefs. By contrast, taking the truth of (L) (or its negation) to be a brute fact offers no parallel advantage. It is a consequence of the epistemicist position that must be lived with, not an explanatorily fecund result.

In fact, not only does taking the truth values of liar propositions to be brute facts produce no explanatory value, it also introduces a vast number of explanatory burdens that do not face other views. Supposing (L) to be the true one, taking (L) as a brute fact raises the

¹⁵ This turns out to be Sorensen's stance with respect to (T), the truth-teller (2001). Sorensen is not an epistemicist with respect to (L), however, for he thinks that there are no liar propositions, let alone truth-valued ones (2001: 182). The objections raised in this section concerning (L) apply equally well to Sorensen's view concerning (T).

¹⁶ For example, Lange (2009) posits the existence of brute counterfactual truths, but then puts them to work by showing how we can use them to give satisfying accounts of various notions important to the philosophy of science, such as the laws of nature.

question of why it's true rather than false, and then admits that the question cannot be answered. There are further explanatory quandaries raised. Do liar propositions have the truth value they do contingently, or necessarily? Do all liar propositions have the same truth value? Liar propositions can be expressed at will *ad nauseam*. Here is one,

The proposition expressed by the first indented sentence of section 5.2 of “Epistemicism and the Liar” is false,

and here is another:

The proposition expressed by the second indented sentence of section 5.2 of “Epistemicism and the Liar” is false.

Do these two new liar propositions have the same truth value? Do they have the same truth value as (L)? Are they necessarily the truth value that they have? Semantic epistemicists must face all of these questions, but they have no resources for answering any of them. By taking liar propositions to have truth values, epistemicists burden themselves with all of these questions, which do not necessarily arise for other views. And by taking such propositions to be brutally true (or false), epistemicists disenable themselves from answering any of these questions. Semantic epistemicists take on extra explanatory burdens that do not face other views, and have no greater resources for handling them.¹⁷

Perhaps the epistemicist might respond by rejecting these truthmaking queries. After all, questions about truthmakers are not without controversy. But the metaphysical concerns driving truthmaker theory are motivated independently of the issues surrounding the semantic paradoxes, and so it would be *ad hoc* for epistemicists to reject such inquiries because they raise challenges for epistemicism. After all, it's not at all obvious that the concerns of epistemicists motivate any of the existing worries about truthmaker theory. If epistemicists must reject truthmaker theory (even in its more modest incarnations that do not presuppose truthmakers for every truth) out of hand simply because it conflicts with their view, then they again suffer a serious theoretical blow to their view.

All told, the epistemicist solution to the liar paradox faces several insurmountable difficulties of its own, and even more when paired with Horwich's other views concerning truth

¹⁷ Cf. the problem of “surprising determination facts” for epistemicists about vagueness (Eklund 2011: 357).

and meaning. Even if epistemicism is of use in solving the soritical paradoxes, it provides no comfort in the case of the liar.

Acknowledgments

My thanks go to the referees for the journal, Bradley Armour-Garb, Thomas Hofweber, and especially Keith Simmons for their invaluable comments. A version of this paper was presented to the Pacific Division of the American Philosophical Association in April 2012 in Seattle. I thank Jeff Snapper for his very helpful commentary, and the audience for a productive exchange.

References

- Armour-Garb, Bradley. 2003. Minimalism, the generalization problem and the liar. *Synthese* 139: 491-512.
- Armour-Garb, Bradley, and Jc Beall. 2005. Minimalism, epistemicism, and paradox. In their *Deflationism and Paradox*, 85-96. Oxford: Clarendon Press.
- Armstrong, D. M. 2004. *Truth and Truthmakers*. Cambridge: Cambridge University Press.
- Beall, Jc. 2001. A neglected deflationist approach to the liar. *Analysis* 61: 126-129.
- BonJour, Laurence. 1978. Can empirical knowledge have a foundation? *American Philosophical Quarterly* 15: 1-13.
- Eklund, Matti. 2011. Recent work on vagueness. *Analysis* 71: 352-363.
- Field, Hartry. 2008. *Saving Truth from Paradox*. Oxford: Oxford University Press.
- Grover, Dorothy. 1977. Inheritors and paradox. *Journal of Philosophy* 74: 590-604.
- Horgan, Terry. 1997. Deep ignorance, brute supervenience, and the problem of the many. *Philosophical Issues* 8: 229-236.
- Horwich, Paul. 1998a. *Meaning*. Oxford: Clarendon Press.
- . 1998b. *Truth*. Second edition. Oxford: Clarendon Press.
- . 2005. A minimalist critique of Tarski on truth. In *Deflationism and Paradox*, eds. Jc Beall and Bradley Armour-Garb, 75-84. Oxford: Clarendon Press.
- . 2008. Being and truth. *Midwest Studies in Philosophy* 32: 258-273.
- . 2010. *Truth-Meaning-Reality*. Oxford: Clarendon Press.
- Kripke, Saul A. 1975. Outline of a theory of truth. *Journal of Philosophy* 72: 690-716.
- Lange, Marc. 2009. *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature*. Oxford: Oxford University Press.

- Lewis, David. 2001. Truthmaking and difference-making. *Noûs* 35: 602-615.
- Priest, Graham. 1979. The logic of paradox. *Journal of Philosophical Logic* 8: 219-241.
- Restall, Greg. 2005. Minimalists about truth can (and should) be epistemicists, and it helps if they are revision theorists too. In *Deflationism and Paradox*, eds. Jc Beall and Bradley Armour-Garb, 97-106. Oxford: Clarendon Press.
- Simmons, Keith. 1999. Deflationary truth and the liar. *Journal of Philosophical Logic* 28: 455-488.
- Sorensen, Roy A. 1988. *Blindspots*. Oxford: Clarendon Press.
- . 2001. *Vagueness and Contradiction*. Oxford: Clarendon Press.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.