

A Puzzle for Constructivism and How to Solve It[†]

Dale Dorsey

Department of Philosophy
University of Kansas
1445 Jayhawk Boulevard
Wescoe Hall, rm. 3090
Lawrence, KS 66045
ddorsey@ku.edu

1. *Constructivism*

The topic of this essay is the metaethics of practical reasons, with special focus on a view I call “constructivism”. Because this term is rather promiscuous, it seems best for me to characterize the view I’m talking about right off the bat. As I understand it, constructivism is characterized by three central tenets. The first of which runs as follows:

Truth-Apt: normative judgments are truth-apt.

For constructivism, normative judgments can be adequately described as true or false. The second feature of constructivism is:

Relational: that which makes a particular normative judgment n_j true is n_j ’s bearing of a favored relation to other normative judgements.

Relational offers a substantive truth condition for normative judgments: a normative judgment n_j is *made true* by the fact that n_j bears a favored relation to other normative judgments, or some suitably specified subset thereof. Third, constructivism as I use the term accepts:

Unrestricted: *Relational* applies to all normative judgments.

Unrestricted holds that all normative judgments permit of the same truth conditions, viz., the bearing of a favored relationship to other normative judgments.

Constructivism, so defined, permits of a number of different specifications. One can identify the proper “relation” in many different ways. One

[†]I would like to thank James Lenman, Cory Wright, Jack Bricke, and an anonymous reader for helpful and challenging feedback on an earlier draft of this paper.

can insist that normative judgments are true if and only if they bear the proper relation to other normative judgments after some suitably identified round of cognitive idealization, for instance, perhaps articulated counterfactually. Constructivism as a class, however, contrasts with a number of alternative metaethical categories, most importantly realism and expressivism. The chief argument between constructivists and their rivals will be over *Relational*. Realists will claim that normative judgments are made true by judgment-independent normative facts.¹ Whether normative judgments bear such a favored relationship to other normative judgments is, for realism, irrelevant to their truth. Expressivists—if they accept *Truth Apt* at all—will insist that the truth of normative judgments is strictly minimal, and does not permit of substantive truth conditions of the sort endorsed by *Relational*.²

The best-articulated, but certainly not only, account of constructivism as I understand it is offered by Sharon Street. Street’s account runs as follows: “the fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A ’s other judgments about reasons.”³ For Street, r is a reason for x to ϕ if and only if the judgment that “ r is a reason to x to ϕ ” “withstands scrutiny” from the perspective of x ’s other judgments about reasons. Hence judgments about reasons will permit of relational truth conditions: whether they are true will be determined by whether they withstand the scrutiny of x ’s other normative judgments. For the purposes of the current paper, I will assume Street’s account; those who prefer an account stated in a different way (e.g., with an alternative characterization of the favored relationship) are free to make the proper translations.

2. Normative Judgments

Reasons, for constructivism, are a product of the normative judgments people make. However, a number of questions arise when it comes to understanding the nature of these judgments. Such questions include, but are not limited to, the nature of judgments in comparison to other propositional attitudes like beliefs and desires, the nature of judgments in comparison to

¹See, for instance, Russ Shafer-Landau, *Moral Realism: A Defence* (Oxford: Oxford University Press, 2003), 16.

²See, importantly, Allan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2003), 42n3.

³Sharon Street, “Constructivism about Reasons” in *Oxford Studies in Metaethics* v. 3 (2008), 223.

other truth-bearers like sentences or propositions, the distinguishing characteristics of non-normative judgments in comparison to normative judgments, etc. Answering these questions is of the first importance for any constructivist view. But I lay discussion of them aside to focus on a further question: what, according to the constructivist, is the semantic content of a normative judgment? What, in other words, do normative judgments *mean*?

Constructivism cannot hold that normative judgments mean just any old thing. They cannot, for instance, ascribe natural or non-natural judgment-independent normative properties (in Roland Dworkin’s terminology, “morons”, or, given our focus on normativity, “normons”⁴). They also cannot be taken to refer to individual perceptual experiences, or conative or affective states.

But why not? The answer, of course, starts with a straightforward claim about *truth*. Generally speaking, we hold that a truth-bearer is true if and only if that truth-bearer’s meaning bears “the right relation” to that bit of the world, or state of affairs, that would make it true (i.e., its *truth condition*).⁵ (“The right relation” can be given any number of interpretations, such as correspondence, satisfaction, mirroring, picturing, etc., etc. I won’t commit to any one of these here.) For the sake of brevity, I will call any theory of truth that accepts this claim a “semantic” theory of truth.⁶ Thus the semantic analysis of any judgment, or any other truth-bearer for that matter, has to properly “match” its truth condition, sufficient to render it the case that whenever the truth condition holds, the judgment’s meaning “bears the right relation” to it. On this general thought, Alston writes that for a given bit of extra-linguistic reality (or “fact”) to make a given proposition (or other truth-bearer) true, “the proposition and the fact that makes it true share the same propositional content. What the fact is a fact *that*, is the same as what the proposition is a proposition *that*.”⁷ Given a semantic

⁴See Dworkin, *Justice for Hedgehogs* (Cambridge, MA: Harvard University Press, 2011), ch. 1.

⁵In referring to states of affairs, “bits” of the world, or, later on, “facts”, I do not intend to be committing myself to a controversial metaphysical position with respect to the existence of such states of affairs, or facts, or their efficacy in a theory of truth. Rather, I use these terms as efficacious shorthand in identifying the extra-linguistic truth conditions of individual truth-bearers.

⁶I do not intend a “semantic theory of truth” to be identical to Tarski’s so-called “semantic conception of truth”, as spelled out by Tarski in “The Semantic Conception of Truth” in *Truth*, ed. Blackburn and Simmons (Oxford: Oxford University Press, 1999). Rather, I simply mean any theory of truth that holds that, to be true, a given judgment’s meaning must bear the right relation to its truth condition. Though Tarski’s view is, as I understand it, a semantic theory, it needn’t be the only one.

⁷William P. Alston, *A Realist Theory of Truth* (Ithaca, NY: Cornell University Press,

theory of truth, then, if normative judgments refer to normons, “ r is a reason for x to ϕ ” would be made true not by its scrutiny-withstandingness, but rather by facts about normons. If normative judgments refer to perceptual experiences, or affective or conative states, “ r is a reason for x to ϕ ” would be made true not by its scrutiny-withstandingness, but rather by facts about perceptual experiences or affective or conative states. And so on.

So far, so platitudinous. However, a problem arises here for constructivism. Insofar as constructivism accepts *Relational*, constructivists must offer a semantic analysis of normative judgments that properly “matches” the bit of the world that makes them true, viz., the fact that they withstand scrutiny. But it would seem that this can occur only if the semantic content of a normative judgment is specified in terms of its scrutiny-withstandingness. To put this in other words, any semantic analysis of “ r is a reason for x to ϕ ” must be put in terms of that judgment itself, i.e., that *that very judgment* withstands scrutiny. And here we find a problem. The semantic analysis of a normative judgment must be elucidated in terms of that which is to be analyzed: that very judgment itself. Any semantic analysis of normative judgments, on a constructivist view, must be circular.

Take an example. Imagine that I judge the following:

1. That “Jumpin’ Jack Flash” is on the radio is a reason for me to pump up the volume.

Given a semantic theory of truth, we must offer a semantic analysis of (1) that suitably matches its truth condition. Given constructivism, however, the truth condition of (1) is the scrutiny-withstandingness of (1) itself. Hence, (1) must be analyzed, at least in part, as follows: “that “Jumpin’ Jack Flash” is on the radio is a reason for me to pump up the volume’ withstands scrutiny from the standpoint of my other normative judgments”. But this analysis is clearly problematic. We were asking what the semantic content (1) is. We were given an answer in terms of (1) itself. To properly understand the propositional content of (1), one must again ask what the propositional content of (1) is. But, quite obviously, a regress now looms. The only additional information we seem able to convey is that (1) should be now be analyzed as: ““that ‘Jumpin’ Jack Flash’ is on the radio is a reason for me to pump up the volume” withstands scrutiny from the standpoint

1996), 38. Alston defines this as a “realist” theory of truth, but his conception of a realist theory is identical to what I have identified as a “semantic” theory. Alston claims that realist theories hold that: “A statement (proposition, belief. . .) is true if and only if what the statement says to be the case actually is the case,” (Alston, 5).

of my other normative judgments' withstands scrutiny from the standpoint of my other normative judgments". And on and on. So far no informative semantics of this judgment has been given, because the semantic analysis of (1) is circular: the semantic analysis of (1) presupposes a prior semantic analysis of (1).

A circular semantic analysis of normative judgments is a problem for any constructivism of the sort I mention. To see this, consider the following argument:

A. For every judgment j , the semantic analysis of j must "match" the truth conditions of j (semantic theory of truth).

B. For some normative judgment of the form " r is a reason for x to ϕ " n_j , n_j is true if and only if n_j bears a favored relation to a set of normative judgments SNJ (by *Relational*).

C. Hence (by (A) and (B)), the propositional content of n_j must be that n_j bears a favored relation to a set of normative judgments SNJ .

D. Hence any semantic analysis of n_j must be specified in terms of n_j .

E. Hence any semantic analysis of n_j is circular.⁸

3. So What's the Problem?

One might reasonably wonder why commitment to a circular semantics of normative judgments is a problem. There are two reasons that strike me as decisive.

First, circular semantic analyses are themselves unacceptable. As Ralph Wedgwood writes, following Kripke: "it is precisely the task of an account of the reference of [a name] to *explain* what it is for someone to be the person referred to by [the name]. Hence, for an account of the reference of the name 'Socrates' to be acceptable, it must not use the notion of "being called 'Socrates'" or "being referred to by the name 'Socrates'", on pain of simply presupposing what is to be accounted for."⁹ Of course, constructivism is not

⁸One might wonder why I have restricted my discussion to views that accept *Unrestricted*, insofar as this principle seems to play no role in the derivation of (E). It could very well be that constructivist views that deny *Unrestricted* also face a problem here. But I restrict my discussion insofar as my solution to this problem succeeds only for views that accept it. I leave aside any potential solutions, or indeed any potential problems, for restricted versions of constructivism.

⁹Ralph Wedgwood, *The Nature of Normativity* (Oxford: Oxford University Press, 2007), 19.

giving a circular semantics of a *name*. But it is committed to an explicitly circular semantics of a class of *judgments*. And the problem is not much different. In the case of normative judgments, what we seek is a semantic analysis sufficient to determine their propositional content, what they refer to, that to which the judgment is committed, etc. But a circular semantic analysis—one that implies the sort of regress I explored in the previous section—cannot accomplish this task. Without a prior understanding of the propositional content of “*r* is a reason for *x* to ϕ ”, we cannot understand the propositional content of “*r* is a reason for *x* to ϕ ”. Call the first problem the “problem of circularity”.

A second problem has to do with the structure of constructivism itself. Constructivism about reasons requires us to be able to determine when a given normative judgment withstands scrutiny from the standpoint of other normative judgments. But it is difficult to see, without a substantive analysis of what these judgments *mean*, whether and when any individual normative judgment will survive such scrutiny.¹⁰ Hence without providing a substantive semantic analysis, or at least a semantic analysis that says more than the one constructivism is stuck with, constructivism about practical reasons, morality, or any other domain offers a set of truth conditions that is, at best, entirely indeterminate.

The best way to see this problem is to consider an explicit example. Imagine that my set of normative judgments contains only two, specified as follows:

2. That it will kill everyone is a decisive reason for me not to start a nuclear conflict between the United States and Russia.
3. That it will kill everyone is a decisive reason for me to start a nuclear conflict between the United States and Russia.

Does (2) withstand the scrutiny of (3)? One would expect that the answer should be “no”. As Street notes:

As soon as one takes anything whatsoever to be a reason, one thereby ‘legislates’ standards according to which, by one’s own lights as a valuing agent, one is making a mistake, whether one knows it or not, if one endorses certain other normative judgments. . . For one normative judgment to *withstand scrutiny* from the standpoint of other normative judgments, then, is for that judgment not to be mistaken as determined by the standards of

¹⁰I thank Jimmy Lenman and an anonymous reader for inspiring further attention to this problem.

correctness that are constitutively set by those other normative judgments in combination with the non-normative facts.¹¹

Stated in this way, it certainly seems as though (2) and (3) do not withstand the scrutiny of each other. From (2), we are to take the proposition “starting a nuclear conflict between the United States and Russia will kill everyone” to be a decisive reason not to start such a conflict. (3) violates that standard of correctness by insisting that the very same proposition is decisive reason *to* start such a conflict. This is a violation of the standard of correctness set up by (2): in insisting that a particular proposition is a decisive reason to ϕ , one is also insisting that it is *not* a decisive reason to $\neg\phi$.

But there are two problems here for constructivism. First, insofar as constructivism is unable to offer any informative semantic analysis of normative judgments, it is difficult to see how one could support the claim that (2) and (3) do not withstand scrutiny of each other. After all, to know whether they do, in fact, survive such scrutiny, one has to understand their semantic or propositional content. Otherwise we have no idea whether (3) “makes a mistake” by (2)’s lights. Hence without an informative account of the propositional content of (2) and (3), whether I actually have a reason to start a nuclear conflict seems indeterminate.

Second, if the semantic analyses of (2) and (3) must match their truth conditions, (2) and (3) must be analyzed, respectively, like this:

4. “That it will kill everyone is a decisive reason for me not to start a nuclear conflict between the United States and Russia” withstands scrutiny from the standpoint of “That it will kill everyone is a decisive reason for me to start a nuclear conflict between the United States and Russia”.

5. “That it will kill everyone is a decisive reason for me to start a nuclear conflict between the United States and Russia” withstands scrutiny from the standpoint of “That it will kill everyone is a decisive reason for me not to start a nuclear conflict between the United States and Russia”.

But if (2) and (3) are properly analyzed by (4) and (5), to know whether (2) and (3) withstand the scrutiny of each other, we must ask: does (4) withstand the scrutiny of (5)? Answer: *absolutely!* (4) merely says that a particular judgment, viz., “That it will kill everyone is a reason for me not to start a nuclear conflict between the United States and Russia” withstands

¹¹Street, 230.

the scrutiny of another judgment, viz., “That it will kill everyone is a decisive reason for me to start a nuclear conflict between the United States and Russia”. (5) is merely the vice versa. But these judgments clearly withstand scrutiny from the standpoint of each other: there is no “standard of correctness” set up by (4) that could possibly be violated by (5). Indeed, one might be tempted to make an even stronger claim: (4) and (5), far from failing to withstand each others’ scrutiny, positively *support* each other. So if (4) and (5) withstand scrutiny, and the proper semantic analysis of (2) and (3) is (4) and (5), respectively, we cannot—or so it would appear—accept the claim that (2) and (3) do not withstand the scrutiny of each other.

One might complain that (4) and (5) are *false*, and hence (2) and (3) are also false. Surely, so it may be claimed, this is precisely the verdict that constructivism intends. But one can plausibly say that (4) and (5) are false only if one offers a semantic analysis of (2) and (3), sufficient to render (2) and (3) non-scrutiny-withstanding. But if constructivism requires that (2) and (3) are analyzed as (4) and (5), then (4) and (5) are false only if *(4) and (5)* do not withstand scrutiny. But they do, and hence (4) and (5) are, or so it would seem, true.

The problem on display here traces directly to the semantic analysis of normative judgments forced on constructivism by a semantic theory of truth. To see this, consider the reading of (2) and (3) described by David Brink as follows: “As many have observed, moral discourse is typically declarative or assertive in form. We say things like ‘The government’s tax plan is unfair’, ‘Waldo is just’, ‘It would be wrong to work for that cause’, and ‘My obligation to Maurice is greater than my obligation to Malcolm’. This language is putatively fact-stating (because it is declarative in form) and certainly seems to ascribe moral properties to persons, actions, policies, and so forth.”¹² Call this a “surface” semantic analysis of normative judgments. This surface semantics seems committed to metaphysically significant moral (or normative) properties, such as reasons, obligations, and so forth. On a surface reading of (2) and (3), they certainly do not withstand the scrutiny of each other: one cannot ascribe the properties ascribed in (2) and (3) as they are ascribed by both (2) and (3). The property of being a decisive reason to ϕ simply rules out the possibility of being a decisive reason to $\neg\phi$. Hence, on this surface analysis, (2) and (3) do not withstand each other’s scrutiny.

But there are two important differences between the semantic analysis forced on constructivism by a semantic theory of truth and the surface

¹²Brink, 25.

semantics described by Brink. First, though normative judgments ascribe properties, they ascribe only the property of scrutiny-withstandingness. Second, these judgments ascribe no property to “persons, actions, policies, and so forth”, but rather state facts about the relation between individual normative judgments. But if this is correct, constructivism cannot vindicate the plausible thought that (2) and (3) do not withstand scrutiny because they do not ascribe the sort of property that would entail a failure of (2) to withstand the scrutiny of (3). They ascribe only scrutiny-withstandingness to themselves, and in so doing (as we’ve seen with (4) and (5)), they do survive such scrutiny. At best, constructivism seems, given its inability to provide a substantive semantic analysis, unable to provide any determinate answer to the question of what reasons one has. At worst, constructivism—given that normative judgments assert only that they withstand scrutiny—seems unable to declare *any* normative judgments false. For the sake of charity, call the second problem the “problem of indeterminacy”.

4. *Constructivism: Metaphysical not Alethic*

A response to the above argument should be considered here. Recall Street’s account of metaethical constructivism: “the fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A ’s other judgments about reasons.”¹³ Here Street says very little about what makes a normative judgment *true*, or what a normative judgment *means*. Rather, Street offers what might be called a *metaphysical account of reasons*: what must be in place for r to be a reason for x to ϕ .¹⁴

But if this is right, constructivism can avoid the problem of circularity. On this view, normative judgments refer, simply, to normative entities: reasons, full stop. In so doing, it shares a truth condition with realism: normative judgments of the form “ r is a reason for x to ϕ are made true by the fact that r is a reason for x to ϕ . But this doesn’t render the resulting view any less constructivist: constructivism can be re-invented as a distinctive *metaphysical* position. As Street urges, the fact that r is a reason for x to ϕ is constituted by the scrutiny-withstandingness of the relevant normative judgments. If we accept *this* version of constructivism, constructivism needn’t be committed to a circular semantic analysis of normative judgments.

¹³Street, 223.

¹⁴See Street, 241-2.

There are two problems with this response. First, even if constructivism is construed as a metaphysical thesis, to avoid this problem it must also reject *Relational*. But though this is a position in logical space, many constructivist views do, in fact, accept *Relational*. If so, though this needn't be a problem for every constructivist view, it remains a problem for a wide range of them.¹⁵ Indeed, Street herself suggests explicit commitment to *Relational* at various points:

According to constructivism...for a normative judgment...to be correct is for it to stand up to the specified sort of reflective scrutiny; the normative judgment's correctness is *constituted* by the fact that it withstands this scrutiny. To speak more metaphorically...the standards that determine the correctness or incorrectness of the normative judgment in question are thought to be "given from within," or "legislated by," some further practical standpoint: to be correct is to withstand scrutiny from that standpoint.¹⁶

Second, even we allow constructivism to reject *Relational*, and hence to avoid semantic circularity, a constructivist metaphysics of reasons still has problems at the semantic level: given that (2) and (3) ascribe a property that is metaphysically constituted by the scrutiny-withstandingness of (2) and (3), the problem of indeterminacy arises again. The problem is this: any standard of correctness established by the ascription of a normative fact, as in (2), must be determined by that which *constitutes* this normative fact. But the metaphysical substratum, as it were, of the normative fact ascribed by (2) is the scrutiny-withstandingness of (2) itself. Hence any "standard of correctness" set up by (2) must be set up by (4), i.e., the claim that (2) withstands the scrutiny of (3). But, as we have seen, the claim that (2) withstands the scrutiny of (3) sets up no standard of correctness whatsoever that conflicts with the property ascribed in (3), i.e., that (3) withstands the scrutiny of (2). If so, it is again difficult to see how (2) and (3) might fail to withstand the scrutiny of each other: they do not ascribe properties that are mutually inconsistent or even in tension.

¹⁵Indeed, it seems to me that specifying constructivism as a theory of what makes normative judgments true rather than as a metaphysical thesis of that which makes up reasons is independently motivated: the former view can claim an additional advantage of ontological parsimony over normative realism.

¹⁶Street, 209. Of course, Street here uses the term "correctness" rather than "truth", but this point is a red-herring. Even if correctness is distinct from truth, it seems to me no less platitudinous to claim that the correctness of a given judgment is a function of a proper relation between the meaning of that judgment and extra-linguistic reality.

The root of this problem is semantic. To vindicate the claim that (2) does not withstand the scrutiny of (3), one must offer a semantic analysis of (2) and (3) that renders (2) and (3) in some sort of tension. On the current proposal, (2) and (3) refer to “reasons”. But if reasons are metaphysically constituted by the simple scrutiny withstandingness of (2) and (3), problems arise again. In referring to a property constituted by (4), (2) can establish no standard of correctness that is incompatible with (3). After all, to establish such a standard of correctness, the metaphysical constituent of the property ascribed in (2) must be in tension with the metaphysical constituent of the property ascribed by (3), i.e., (5). But (4) and (5) are most certainly not in tension. If so, (2) fully withstands the scrutiny of (3) and *vice versa*. To solve this, one must adopt a different semantic analysis of (2) and (3), and hold that they refer to a more robust property than “reasons” (on the current analysis), e.g., some judgment-independent property possessed by the fact of mass deaths, sufficient to rule out (3) on the assumption of (2). But this semantic proposal, though it is sufficient to deliver the plausible claim that (3) does not withstand the scrutiny of (2), is obviously incompatible with constructivism. First, it would be hard to see how constructivism’s metaphysical analysis of reasons would avoid obsolescence: normative judgments do not refer to properties that are made up by the scrutiny-withstandingness of judgments. Second, and perhaps more importantly, if normative judgments like (2) and (3) refer to robust, judgment-independent properties, and if we accept the semantic theory of truth, (2) and (3) will be true *if and only if* these judgment-independent properties hold. Their scrutiny-withstandingness would be simply beside the point.

To sum up: rejecting *Relational* does not solve constructivism’s problems. For the remainder, then, I return to consideration of variants that accept *Relational*.

5. *Good News and Bad News*

So long as the meaning of a judgment must bear “the right relation” to its truth condition, a constructivist, who defines the truth conditions in terms of a normative judgment’s bearing a favored relation to other normative judgments, is forced to offer a semantic analysis of normative judgments that, even independent of its circularity, is problematic indeed. For constructivism to survive, we must break the link between the meaning of normative judgments and their truth conditions. And, as far as I can tell, the only way to do this is to reject a semantic theory of truth.

This is good news and bad news. The good news is that one can reject a semantic theory of truth without absurdity, and can thus maintain a constructivist metaethic in the wake of the problems I note here. I argue for this latter claim in the next two sections. The bad news, however, is that a semantic truth predicate has been thought to be simply platitudinous, a feature of the mere *concept* of truth. If so, many will claim that despite the non-absurdity of the alternative I propose, any revision to our understanding of truth is just too radical a move to be justified simply to save constructivism. For my purposes, however, I am satisfied with a conditional conclusion: *if* we are to accept constructivism, we must accept my alternative approach to truth in the normative domain, or an alternative equally incompatible with a semantic theory. In addition, I hope to have shown that for constructivism that there is *a* way forward, whether or not there are others, and whether or not the going is too rough.

6. *Constructivism, Coherence, and Truth*

Constructivists must reject a semantic theory of truth. What, then, should they accept? To answer this question, it is helpful to say a bit more about the nature of constructivism as a general view of practical reasons. Though this categorization is imprecise, constructivism is best thought of as a form of *idealism* about normativity. For the constructivist, practical reasons are simply a product of the various relationships between our own ideas or judgments about normative facts. Constructivism's general idealism helps to guide the choice of truth predicates, however. Idealists have generally been attracted to a *coherence* theory of truth. I suggest we explore this traditional pairing. On this view, (1), for instance, is true if and only if (1) is coherent with the other normative judgments to which I am committed.

Before I investigate the advantages of a coherence truth predicate for the normative domain, a few words must first be said about the coherence theory itself. First, such a view would define truth for normative judgments in terms of their coherence. However, a problem arises. What is coherence? An adequate explanation of coherence has been a classic sticking point for coherence theories in all domains. However, this problem is less daunting given that we seek to offer a coherence theory in a way that *conforms to the general requirements of constructivism*. If so, what it means for a set of normative judgments to be coherent will be defined by whatever favored relation is referred to by *Relational*. For Street, for instance, coherence will

be defined in terms of “scrutiny-withstandingness”.¹⁷

I should address an immediate form of skepticism about a coherence theory of truth. As Richard Kirkham writes: “If we believe that there is a world independent of our thoughts, then no proposition that purports to describe that world can be considered true if it is inconsistent with that world, no matter how well it coheres with other propositions. And if it does express that world accurately, then it cannot be false, no matter how much it fails to cohere.”¹⁸ Given, however, that we certainly believe that constructivism as a general view is not true of, say, the domain of *physics*, or any other non-normative domain, wouldn’t it seem a wildly radical maneuver to adopt a coherence theory on the basis of constructivism about normative judgments?

In response, the constructivist will say that *Relational* is only applicable to the normative domain, and hence the coherence theory of truth itself is acceptable only for the normative domain. This will commit constructivism to a position that has come to be known as “alethic pluralism”, the view that alternative domains of discourse can permit of different truth-constitutive properties. The problem on the table arises only if we assume that the property of coherence constitutes truth for judgments in *all* domains of discourse. But we should not do this. A constructivist about normativity is not committed to being a constructivist about all domains, including, say, the domain of physics. In domains for which a form of ontological realism is appropriate, the coherence theory of truth is inappropriate.

7. Questions

My proposal, then, is to reject a semantic theory of truth—not on the whole, but simply for the normative domain—and replace it with a coherence theory of truth. I hasten to add that this view is not novel. Similar views have been explored by Quine,¹⁹ Crispin Wright,²⁰ and Michael Lynch, who writes:

¹⁷For a more thorough description of how to make one’s actual set of normative judgments into a coherent set, see Dale Dorsey, “A Coherence Theory of Truth in Ethics” in *Philosophical Studies* 117 (2006). See also Dale Dorsey “Truth and Error in Morality” in *New Waves in Truth*, ed. Wright and Pedersen (New York, NY: Palgrave Macmillan, 2010).

¹⁸Kirkham, 111.

¹⁹W. V. Quine, “On the Nature of Moral Values” in *Theories and Things* (Cambridge, MA: Harvard University Press), 63.

²⁰Wright, 227

If, as we are assuming, moral judgment takes place within the space of reasons—we provide reason and evidence for our moral beliefs—our moral judgments are subject to significant rational norms. Consequently, a natural suggestion is that moral judgments are made true by a property that is constructed out of those epistemic norms. . . Given coherentism’s plausibility as an account of the structure of warrant in morality it seems natural for anyone attracted to an epistemic account of the property that makes moral judgments true to appeal to it.²¹

Lynch offers one reason to accept a coherence theory in the normative domain: that normative judgments are made true by the conditions by which they are warranted, and they are warranted by their coherence. I am strictly neutral on this motivation, as I explore another: a method by which to salvage constructivism about practical reasons. But whatever the motivation, there surely remains a number of unanswered questions about this project. I address the most important ones here.

Question 1: It’s all well and good to reject a semantic truth predicate. But don’t constructivists still have to accept a circular semantics?

No. Recall that constructivism was required to accept a circular semantics *as a result* of a semantic truth predicate, which holds that a judgment is true if and only if its meaning bears the “right relation” to extra-linguistic reality. But if we reject this claim, there is no pressure to offer a semantic analysis for any normative judgment that “matches” that normative judgment’s truth conditions. If we know that a normative judgment is true, and we know that this normative judgment is true if and only if it coheres with other normative judgments, we are unlicensed to make any claims about its semantics. We are licensed only to say that it coheres with other normative judgments, whatever their semantic analysis may be.

Question 2: But then what sort of semantics should constructivists accept?

Given the problems stated in §3, we must be given a semantics: what normative judgments *mean* that doesn’t simply reduce (2) and (3) to (4) and (5). And though I hasten to add that a constructivist cannot be totally neutral with regard to the semantic choices it makes, it has wide latitude.

²¹Michael Lynch, *Truth as One and Many* (Oxford: Oxford University Press, 2009), 164, 168.

Constructivists are licensed to select any semantic programme they wish that can also deliver the proper verdicts concerning which normative judgments withstand scrutiny, or bear the favored relation, to other normative judgments.

In particular, and most plausibly, constructivism is licensed to accept the “surface” semantics as described by Brink, i.e., that normative judgments “ascribe [normative] properties to persons, actions, policies, and so forth,” and ascribe properties that are more robust than simple scrutiny-withstandingness. Constructivists, for instance, could accept the “normonic” semantics discussed in §2: one could hold that “ r is a reason for x to ϕ ” means that a certain property—a “normon”—holds of r . The “normon” could be understood naturally or non-naturally, perhaps as a form of *sui generis* normative property. The ascriptions of “normons” needn’t be committed to any metaphysics that refers back to normative judgments themselves, and hence these ascriptions can set up “standards of correctness” just as any standard judgment-independent property ascription would.²²

Question 3: Wouldn’t this semantic analysis yield an error theory for normative judgments?

No. Judgments in a particular domain d permit of an error theoretic analysis if and only if judgments in domain d are truth-apt, *and* all judgments in domain d are false. Indeed, as I have argued elsewhere,²³ far from succumbing to an error theory, accepting a coherence theory of truth for normative judgments is a novel method by which to avoid an error theory. Arguments for error theories generally start by offering a semantics of judgments of the domain in question. In this case, normative judgments refer to “normons”. Next, error theories offer a metaphysical claim: “normons” do not exist. (Famously, Mackie’s arguments for the metaphysical claim include the argument from relativity and the argument from queerness.²⁴) But the inference from the claim that no such metaphysically queer properites exist to an error theory—that is, that all normative judgments are false—goes through *only if we accept that normative judgments are judged true or false on the basis of their bearing the “right relation” to such nonexistent meta-*

²²One note: this semantic analysis seems to render a constructivist’s metaphysics of reasons, once again, obsolete: no normative judgments actually refer to any entities made up of the scrutiny-withstandingness of judgments. But no matter. Constructivism remains a distinctive metaethical position in virtue of its acceptance of *Relational*.

²³See “Truth and Error in Morality”, op. cit.

²⁴Cf. Mackie, 36-42.

physically queer properties. But normative judgments, on this view, are not judged as true or false in this way. Rather, they are judged true or false on the basis of their coherence. Hence even if such norms don't exist, judgments that ascribe them needn't be false if we judge the truth or falsity of the judgments in question on the basis of their coherence.

Question 4: Isn't this a realist semantics? And if so, doesn't this mean that a coherence theory of truth—by your own admission—is inappropriate?

A “normonic” semantics of normative judgments is paradigmatically realist—a semantics that refers to judgment-independent normative properties. If so, it would seem that accepting a coherence theory of truth for normative judgments—and their realist propositional content—is explicitly accepting something I have so far deemed inappropriate: the combination of realism with a coherence theory of truth.

This analysis, however, is mistaken. A coherence theory of truth is inappropriate for a domain d if and only if we believe that facts about domain d are facts about a “world independent of our thoughts”, to borrow Kirkham's phrase. To put this in another way, realism and the coherence theory are an inappropriate mix given realism's *metaphysics*, not realism's *semantics*. Just because we accept a realist *semantics* of normative judgments doesn't mean we have to commit to evaluating normative judgments according to a semantic theory of truth. Rather, because constructivists do not believe that facts about normativity are facts about a “world independent of our thoughts”, they are licensed to choose a coherence theory, and evaluate normative judgments differently. The semantic analysis of normative judgments, then, says nothing about whether the *truth* of such judgments is *itself* semantic. That is settled by further, metaphysical, questions about the normative domain. Insofar as constructivism retains a realist semantics while rejecting a normative “world independent of our thoughts,” the constructivist is licensed to accept a coherence theory of truth for that domain.

Question 5: Isn't the definition of coherence circular?

Constructivism needn't be committed to a circular semantics. But a serious problem looms. Consider the nature of *coherence*. Surely any plausible account of coherence is going to be put in terms of the entailment relations between judgments, i.e., whether judgments are inconsistent with others, whether they can be inferred from a set of other normative judgments, etc. But entailment relations are understood in terms of truth: for p and q to

be inconsistent means that p and q cannot be true together. That p entails q means that p cannot be true and q false, and so forth. Hence it would appear that coherence is defined in terms of truth. But because truth, at least for normative judgments, is defined in terms of coherence, any account of coherence will be viciously circular.²⁵

This is a serious problem. After all, constructivism relies on a sensible understanding of what it means for two judgments to “cohere” or “withstand scrutiny” or bear whatever favored relation to each other. But if truth is defined in terms of this favored relation, and if this favored relation is defined in terms of truth, we seem to be led back to the problem caused by a circular semantics: we seem unable to say with any certainty whether two judgments will bear this favored relation, because we don’t have any informative account of what this favored relation *is*.

Accepting alethic pluralism, however, allows us avoid a circular analysis of coherence. Whether the members of any given set of judgments bear any particular relation to each other is *not* a question that is properly evaluated in the normative domain: the coherence of a belief set is a purely non-normative question. However, because non-normative domains make use of a non-coherence truth predicate, the coherence of any particular belief set is evaluated *given a non-coherence account of truth*. Let me put this point in a slightly different way. Assume that non-normative truth-bearers are susceptible to a semantic truth predicate. Because a judgment of the form “system of belief S is coherent” is not a normative judgment, this judgment is evaluated by means of a semantic truth predicate; this judgment is true just in case its meaning bears the right relation extra-linguistic reality. But if that is correct, the coherence of S is determined by the various entailment relations between its constitutive judgments, i.e., whether they can be true together, whether they set up “standards of correctness” that others survive, etc. But because the domain in question is non-normative, the applicable truth predicate is semantic. Of course, if we accept constructivism, all normative judgments, as evaluated by a semantic truth predicate, will come out false (assuming a “surface” or “metaphysically queer” semantics). But this does not mean they cannot be evaluated for their various entailment relations given a non-coherence truth predicate. Hence the definition of “coherence” is not viciously circular. Truth, for normative judgments, is understood in terms of coherence, which is thereby understood in terms of the ability of particular judgments to be true together assuming a *non-*

²⁵Kirkham, 107.

coherence, or *semantic*, truth predicate.²⁶

By way of a conclusion to this section of the paper, let me sum up my positive proposal and its rationale. The problems of circularity and indeterminacy are foisted upon constructivism *so long as constructivism accepts a traditional, semantic, theory of truth*. Hence to avoid them, constructivist views must reject a traditional, semantic truth predicate in favor of an alternative. My proposal: accept a metaphysically robust “surface” semantics of normative judgments, together with a coherence theory of truth for normative judgments. This account captures the heart of a constructivist view. After all, a coherence theory of truth for normative judgments guarantees both *Truth-Apt* and *Relational*. For the coherence theory, normative judgments are true if and only if they are coherent with an agent’s other normative judgments, where “coherence” is defined according to the favored relation as specified by a particular constructivist view. Thus a coherence theory of truth can capture the essential element of constructivism: that the truth of a normative judgment is set by its bearing of a favored relation (coherence) to other normative judgments.

8. *The Problems of Alethic Pluralism*

Recall that constructivism need not claim that judgments in all domains permit of a coherence truth predicate. Rather, the constructivist must claim that non-normative judgments permit of a different truth predicate than normative ones.

This result is significant in itself. As it turns out, to avoid the problem of semantic circularity, the constructivist must embrace alethic pluralism, and so must address the various problems that confront it. This might be thought an unattractive predicament. Alethic pluralism is accompanied by familiar problems. Three are worth noting here. First, it is unclear that pluralism about truth can plausibly accommodate inferences involving judgments permitting of different truth predicates.²⁷ Second, pluralism cannot seem to provide a coherent analysis of logical connectives including, e.g., conjunctions, disjunctions, etc. Third, pluralism would seem to leave open the possibility that semantically identical judgments can be evaluated according to different truth predicates. But if so, how are we to determine when a particular judgment is to be evaluated by, e.g., a coherence or se-

²⁶See Dorsey, “A Coherence Theory of Truth in Ethics”, 502-4.

²⁷Cf. Christine Tappolet, “Mixed Inferences: A Problem for Pluralism about Truth Predicates”, *Analysis* 57 (1997).

mantic truth predicate?

The point of this section is to introduce these problems as problems for a constructivist view that require solutions: constructivists must view these as *their problems*. I attempt to sketch solutions to the problems, but I stress that these solutions are tentative, and are certainly not the last word on the subject. Indeed, my solutions require substantial retooling of some important bits of our philosophical conceptual scheme. As I note in §9, whether such retooling is ultimately acceptable in light of the plausibility of constructivism is as yet up for grabs.

8.1. *Mixed Inferences and Mixed Connectives*

Mixed inferences are a problem for truth pluralism because truth pluralism appears unable to accommodate the sense in which valid inference is “truth preserving”. Tappolet writes:

[T]here is a simple and equally powerful objection to the claim that there is a plurality of truth predicates. Consider the following inference:

- (1) Wet cats are funny.
 - (2) This cat is wet.
- Ergo*, this cat is funny.

The validity of an inference requires that the truth of the premisses necessitates the truth of the conclusion. But how can this inference be valid if we are to suppose with Crispin Wright that two different kinds of truth predicates are involved in these premisses? For the conclusion to hold, some unique truth predicate must apply to all three sentences. But what truth predicate is that?²⁸

This problem applies to a constructivist’s truth pluralism. Take the following inference:

- A. That a particular action causes pain is a reason for anyone not to do it.
- B. Torturing babies causes pain.
- C. Hence, anyone has a reason not to torture babies.

This inference is certainly valid. But notice that, given the proposal I suggest here, its premisses do not permit of the same truth predicate. But if an

²⁸Tappolet, 209-210.

inference starts from normative and non-normative premises, and results in a normative conclusion, what is it preserving? Certainly not semantic truth, because the conclusion is evaluated according to a coherence truth predicate. Certainly not coherence truth, either. One of the premises is semantic.

Take now the problem of mixed connectives. Consider:

6. “That cookies are delicious is a reason for me to eat them, and snow is white.”

(6) is certainly true. But under which predicate? To say that it is susceptible to the semantic truth predicate would seem to render the left-hand conjunct false (at least on a constructivist view). To say that it is susceptible to a coherence truth predicate would seem to render us idealists about the whiteness of snow. This presents a problem. Which truth predicate evaluates mixed connectives? It seems the answer is: none at all.

One possible solution to this problem—and one I fully license—is to accept a version of alethic pluralism Michael Lynch calls “truth functionalism”. According to Lynch, to be true is to bear a property that plays the “truth-role”, where the “truth-role” can be distinct in different domains of discourse.²⁹ Here truth is a *functional* property that is preserved in valid inference. It just so happens that what plays the “truth-role” may be distinct for different premises of valid mixed inferences. (Lynch’s solution to the problem of mixed connectives requires holding that a different property plays the truth-role for connectives: the property of its truth-value being “grounded” in its atomic propositions.³⁰)

This solution is OK with me. But I want to try out a somewhat different suggestion. This proposal begins by introducing a new truth-like predicate. Call this predicate “truth*³¹”. A truth-bearer is true* if and only if that truth-bearer is true *as evaluated by the truth predicate appropriate for its domain*. A non-normative truth-bearer is true* if and only if extra-linguistic reality is as the truth-bearer says it is. A normative judgment is true* if and only if that judgment coheres properly with the rest of the speaker’s normative judgments.

Now take the problem of mixed connectives. The introduction of truth* can solve this problem with a mere tweak of the semantics of connectives. Instead of claiming that a conjunction *means* that both conjuncts are *true*,

²⁹Lynch, ch. 4.

³⁰Lynch, 90-91.

³¹See Dorsey, “A Coherence Theory of Truth in Ethics”, 504-6; for what it’s worth, I take my solution to the problem of mixed inferences expressed there to be unnecessarily complex, though the thrust of my solution stands.

we say that a conjunction *means* that both conjuncts are *true**. Hence (6) is true if and only if each conjunct is *true**, i.e., “that cookies are delicious is a reason for me to eat them” coheres with my other normative judgments, and whether “snow is white” bears the right relation to extra-linguistic reality. Given *this* analysis of connectives, any connective can properly be said to be susceptible to a standard, semantic truth predicate. Whether any particular judgment, or other truth-bearer, is *true** is a non-normative question, settled simply by facts about the coherence of my normative judgments. Hence it would appear that a slight alteration in the semantic analysis of connectives allows us to plausibly assign a single truth predicate, i.e., a *semantic* truth predicate, to all connectives (including mixed connectives). Though, on a semantic theory of truth, “I have reason to eat the cookies” might very well be false (depending on its semantic analysis), it is, nonetheless, *true**, i.e., true as assessed by the predicate appropriate to its domain.

If we accept this account of mixed connectives, the proper account of mixed inferences is trivial. We can perfectly preserve the claim that valid inferences are truth preserving if we understand that which inferences preserve to be *truth** rather than truth. However, given the definition of *truth**, this guarantees that, for valid mixed inferences, whenever the premises are true, the conclusion is true as well.³² Though I believe either solution succeeds, I am neutral regarding the relative merits of my proposal versus Lynch’s truth functionalism. Suffice it to say, however, there is at least *a* solution.

8.2. Normative and Non-Normative Normative Judgments

The third problem runs as follows. Consider:

7. That this suit is well-made is reason for me to purchase it.

If we accept a “normonic” semantics for judgments like (7), (7) would appear to mean:

8. Normonic property “reason for me to purchase this suit” holds of the fact that this suit is well-made.

But (8), quite obviously, can be assessed *non-normatively*. And if we wish to deny that normons exist, (8) is false. Hence (7) is false. But, if we assume

³²It is important to note here that this way of understanding mixed inferences puts constraints on a proper understanding of coherence. In particular, it must be the case that coherent normative judgments must also be consistent with true non-normative claims. But I take this to be an uncontroversial requirement.

that (7) withstands the scrutiny of my other judgments about reasons, (7) is true. But if this is correct, (7) is both true and false. How could this be?

Of course, (7) *is* both true and false. But this isn't as strange as it sounds. Such a claim is absurd only if the truth predicate used to evaluate a proposition as true and false is identical, i.e., if the domain of discourse is the same. But, quite obviously, (7) is evaluated as true in the normative domain, false in a non-normative domain. In this way, (7)'s truth and falsity might be compared to that of another sentence:

9. Sherlock Holmes lives at 221b Baker Street.

In discussing *fiction*, we are inclined to say that (9) is true. But when in a non-fictional domain of discourse, we are tempted to say that (9) is false. But this is not mysterious. Neither is it in the case of (7). Of course, one might wonder what distinguishes the normative use of (7) from a non-normative use of (7). Though I don't have much of a theory to offer here, it seems to me that the same thing said about (9) should be said about (7), viz., that it depends not on the semantics of a given judgment, but rather its pragmatic elements, as well as the context of utterance.³³ Though I don't have the space to develop this suggestion any further, it seems to me that something like this is roughly correct.

9. *Conclusion: Isn't this All Just Wildly Ad Hoc?*

Before I conclude, a nagging objection should be dealt with. I have offered a solution to the problem of semantic circularity for constructivism that rearranges our understanding of the nature of truth in the normative domain. Not only this, to accommodate the truth pluralism this requires, I suggested a number of further semantic maneuvers, including an alternative semantic analysis of connectives. Isn't all this just *ad hoc*? Is there any independent reason to believe that any of this stuff is actually true?

This objection should be broken down into two different questions. First: is there any reason, independent of the prior acceptance of constructivism, to accept a coherence theory of truth in the normative domain, and its theoretic accoutrements? Second: if the answer to the first question is no, is this an objectionable feature of the solution I propose?

Take the first question first. I think there is, indeed, independent reason to believe that a coherence theory of truth in the normative domain,

³³See, for instance, Jody Azzouni, *Talking About Nothing* (Oxford: Oxford University Press, 2010), 251-2, for a helpful analysis of the fictional case.

along with its alethic pluralism, succeeds. As I note above, I have argued elsewhere that such a move plausibly eliminates the spectre of an error theory.³⁴ Furthermore, I have argued that a coherence theory for normative judgments can be motivated by certain Quinean reflections about the relationship between normativity and observation.³⁵ Indeed, there are a number of independent motivations for a more general alethic pluralism, beyond constructivism.³⁶ Assume for the purposes of argument, however, that these prior motivations fail. Let's say that the only genuine motivation one might have for accepting the view I propose here is a prior commitment to constructivism. This raises the second question: is this problematic for the solutions I propose?

Well, yes and no. It may very well be that there is pressure not to alter, e.g., the semantic truth predicate, not to alter the semantics of connectives in the way I propose, etc. But it is important not to overstate this pressure. The conceptual, linguistic, or logical apparatus proposed for any domain is only as acceptable as its ability to allow us to say things we wish to say about the domain in question. In other words, if we find constructivism about the nature of practical reasons plausible, this should be all the evidence we require to reject a semantic theory of truth. If so, the acceptability of my positive proposals will depend on the strength of the *metaethical* argument for constructivism. But this is as it should be.

Allow me to put this point in slightly different terms. The problem for constructivism is serious: under a semantic theory of truth, constructivism must accept a semantic analysis of normative judgments that is seriously problematic for the constructivist project as a whole. Hence, to salvage constructivism, we should accept an alternative theory of truth for normative judgments. Any such alternative (including my own favored proposal) raises a number of complications which, though I believe can be adequately dealt with, require some reorganization of our philosophical conceptual scheme. At least for the sake of argument, I concede that none of these maneuvers is motivated independently of constructivism. Hence at this point we have a choice: we can either accept these technical suggestions as the price of accepting constructivism, or we can reject these technical suggestions and with them reject constructivism. Which alternative we should choose depends fully on the strength of the *argument for constructivism*. If this argument is good enough, there should be no hesitation to reject a semantic

³⁴See "Truth and Error in Morality", op. cit.

³⁵See "A Coherence Theory of Truth in Ethics", op. cit.

³⁶Lynch, ch. 2.

truth predicate in the normative domain. Such a move is not *ad hoc*, but is simply a necessary corollary of a metaethical view we have reason to accept. Nevertheless, as I do not mean to assess the strength of the argument for constructivism here, I leave its all-things-considered fate an open question.