17

THE NORMATIVITY OF MEANING REVISITED[1]

Paul Boghossian

Introduction

In this essay, I revisit a topic that Allan Gibbard and I have been debating off and on for about three decades and which was ignited by Saul Kripke's (1982, 37) remark:

> The relation of meaning and intention to future action is *normative*, not *descriptive*.

Ever since Kripke made this remark, philosophers have been arguing over what this 'meaning is normative' slogan means and whether it is true on any of its legitimate interpretations.

Kripke attributed the insight to Wittgenstein (1953), although he makes it clear that, at least as far as this aspect of Wittgenstein's thought is concerned, he is sympathetic to it.

Kripke saw in the slogan a powerful weapon with which to combat proposed naturalistic reductions of meaning, in particular dispositional analyses. If the relation between the meaning of a word and its use is normative, so Kripke seems to have thought, then meaning can't consist in dispositions to use the word in a certain way, because the relation between a disposition and its exercise is descriptive not normative.

Gibbard joins Kripke in finding the slogan important; and he has been exploring it sympathetically over the past few decades. Recently, he has published a magisterial volume,

---

1 I've had the pleasure of discussing this and other topics with Allan Gibbard ever since I had the good fortune of becoming his junior colleague at the University of Michigan at Ann Arbor in 1984. Allan was an inspiring interlocutor, an insightful critic, and a generous mentor. This contribution to a volume in his honor is a totally inadequate expression of thanks.

Meaning and Normativity, which pulls together a lifetime's worth of profound reflections on the nature of normativity, the nature of meaning, and the relations between them.

Prima facie, it is somewhat surprising that Gibbard should be such a fan of the 'normativity of meaning' slogan. His other views, you might think, would disincline him from embracing it.

For one thing, Gibbard is a committed naturalist, who thinks that all facts are ultimately grounded in natural facts and properties. But doesn't the normativity of meaning make trouble for naturalism, as Kripke thinks?

For another, one of the ways in which Gibbard tries to secure a naturalistic outlook is by being an expressivist about normative discourse. Seemingly, though, expressivist meaning requires a contrast with nonexpressivist factual meaning. However, if the very notion of meaning is itself normative, then meaning talk must itself be given an expressivist characterization. And the result will be that the very distinction between expressivist and nonexpressivist meaning will itself have to be a distinction with merely expressivist meaning. Won't this be a problem? Wouldn't an expressivist want the contrast between expressivism and factualism to be itself a factual contrast?

Despite all this, Gibbard insists that the slogan is importantly true.

Normativity of Meaning as Correctness

In order to set some context for the debate, I will say a little about my own relation to this slogan.

In the literature, I have seen it said by some that I am a strong *proponent* of the normativity of meaning thesis, and by others that I am a strong *opponent*. This is not usually a very good sign about the degree of clarity that one has brought to the topic. I will try to clarify what might be going on here.

In my earliest writing on Kripke's book, "The Rule-Following Considerations," of 1989, I put forward a particular view of what the slogan meant, an interpretation that in the literature is apparently referred to as the 'orthodox interpretation.' To quote a bit from that paper:

> Suppose the expression 'green' means *green*. It follows immediately that the expression 'green' applies *correctly* only to *these* things (the green ones) and not to *those* (the non-greens) . . . The normativity of meaning turns out to be, in other words, simply a new name for the familiar fact that, regardless of whether one thinks of meaning in truth-theoretic or assertion-theoretic terms, meaningful expressions possess conditions of *correct use* . . . Kripke's insight was to realize that this observation may be converted into a condition of adequacy on theories of the determination of meaning: any proposed candidate for the property in virtue of which an expression has meaning, must be such as to ground the 'normativity' of meaning—it ought to be possible to read off from any alleged meaning constituting property of a word, what is the correct use of that word. (513)

On this view, clearly, the label 'normativity of meaning' is being used fairly lightly. Most of the philosophers who get excited about Kripke's slogan think that he is pointing to something that (even if not entirely without precedent) would be considered highly controversial, and which would be denied by most naturalistic reductionists about meaning.

By contrast, I say that it is just a new label for a very familiar phenomenon that the meaning of a linguistic expression—say a predicate—is inextricably linked to some notion of correct application: the main candidates being either true application or warranted application. And this latter claim no one denies.

For all that this use of 'normativity' is light, it is not entirely without interest or import. As Gideon Rosen (1997) has rightly emphasized, the notion of *correctness* here is a quite *general* one that applies even to performances of scores of music, or of dances, or of ways of folding a foldable bike, cases where there is no truth or warranted application at issue.

This is importantly right and one of the main reasons why *all* of these topics fall under the general heading of 'rule-following': they all have the following structure in common:

> A certain *standard* for doing something—playing a piece of music, dancing a dance, playing a game, using a word—which in and of itself may not be *rationally mandated*, is *adopted*. Once it is adopted, it serves as a standard for subsequent performance, pronouncing that performance as correct or incorrect.

Is Correctness Real Normativity?

A propos of this way of thinking about the slogan 'meaning is normative', many have complained that it construes 'normativity' too lightly (see, e.g., Wikforss 2001, Glüer and Wikforss, 2015).

The complaint is: 'correct' as it is being used here is not a genuinely normative term; hence the label 'normativity of meaning' is inappropriate and misleading. This is not real normativity and hence does not pose even a prima facie problem for naturalism's ability to account for meaning.

There are three questions here that I would like to look at briefly (too briefly, given the complexities involved):

1. What is it for a notion to be genuinely normative?
2. Does 'correct' count as genuinely normative?
3. How interesting is the previous question (2) for the purposes of evaluating the naturalistic program in the theory of meaning?

The critics tend to assume that for a notion to be genuinely normative, claims involving it should analytically entail subjective ought-claims. Furthermore, they point out, semantic

correctness claims do not imply subjective oughts. Hence, they conclude, semantic correctness is not genuinely normative.

I think that the critics are right that semantic correctness claims are not subjective ought-entailing (see my 2005). If I mean addition by '+' it would be *correct* for me to say '68 + 57 = 125'. But it doesn't immediately follow, without further normative assumptions, that if I am asked about that sum, I *ought* to say '125.' If I mean addition, and I want to tell the truth, I ought to say '125.' But if I want to mislead my audience, I ought to say something else, something incorrect. So, correctness is not necessarily ought-entailing or 'action-guiding.' (Notice that I am assuming here that if linguistic meaning were the source of norms, it would have to be the source of norms on assertion. This is something that, as we will see later on, Gibbard wants to deny.)

At any rate, my view, along with that of, among others, Rosen and Peter Railton, is that correctness claims count as normative even though they are not subjective ought-entailing. The notion of correctness is *evaluative*—it appeals to a standard—even if it is not prescriptive.[2]

The distinction between these two kinds of normativity shows up clearly when we say that *a false belief is incorrect*. That's not just a way of repeating the claim that the belief is false; it implies that something has gone wrong with that belief, that it is defective. And, yet, it could simultaneously be true that the thinker subjectively *ought* to have come to that belief given the (misleading) evidence available to her.

So, meaning claims are correctness-entailing. And semantic correctness is a kind of normativity, although it is not of the 'action-guiding' variety.

Some of this helps explain why I have been regarded by some as pro normativity of meaning, and by others as anti. It is partly about where you stand on the 'correctness as normative' issue.

Another potential confound is that while I do regard correctness as normative, I am also inclined to think that the question whether semantic correctness is genuine normativity is an issue that is not all that interesting in connection with the naturalistic program in the theory of meaning. I will explain why I think that later on in this essay.

Before doing that, however, I want to take a look at Gibbard's recent attempt to revive the claim that meaning is normative in the full-throated sense of being analytically subjective ought-entailing.

Gibbard on the Normativity of Meaning

Gibbard's overall view is, as you might expect, complicated, subtle, original, and impressive. I'm not entirely confident I understand all of its different parts, or how they fit together.

In this section, I am just going to look at one central argument that Gibbard provides for thinking that the concept of meaning is normative. Gibbard intends the argument to be preliminary. He takes the real test of his theory to be whether, when taken as a whole,

2 For a related distinction, see McPherson and Plunkett (2017).

it illuminates the vexed topic of meaning in a way that other theories have failed to do. So, what I will be assessing here is not his overall view but the preliminary support for it that he offers. A full assessment of his theory would be a daunting task.

Gibbard starts by distinguishing between a weak Normativity of Meaning thesis and a strong one. According to the *Weak Thesis* (22):

> Claims about meaning, all by themselves, without the help of further normative assumptions, analytically imply *ought* claims.

The *Strong Thesis* (22), by contrast, asserts not only the Weak Thesis, but also that:

> Meaning can be fully defined through some combination of normative and naturalistic concepts.

Gibbard accepts both theses. As a result, according to Gibbard:

> Every means implies an ought.

And

> For every means, there is an ought that implies it.

Notice that Gibbard is making these claims about the *concept* of meaning, not the *property* of meaning. It's one of the distinctive features of Gibbard's view that he works with a very strong concept/property distinction. While he insists that the concept MEANING is normative, he equally insists that the property of meaning, the property in virtue of which things fall under the concept MEANING, is as naturalistic as you please. Indeed, he believes that the idea of a normative *property* makes no sense. Only concepts can be normative or not. This explains why he doesn't fear that his normativity of meaning theses will have anti-naturalist consequences. By limiting the normativity theses to the *concept* of meaning he leaves it open that the property of meaning is constituted purely naturalistically.

I will raise a question about whether we have any good reason to believe even the Weak Thesis.

Gibbard's Argument

Gibbard offers the following argument in its support:

> A chief reason to believe the weak normativity thesis . . . is that a certain basic kind of ought follows from a means *invariably*. (16; emphasis added)

Gibbard explains that the 'ought' he is talking about here is Ewing's (1939) ought: an ideal, primitive ought that ignores costs and limitations on our powers of reasoning.

What's an example of the sort of entailment at issue? Gibbard says:

> With Ewing's primitive ought, we can say, an ought does follow from the meaning of 'nothing', and follows invariably. One ought not to believe both that snow is white and that nothing is white. This follows in a way not dependent on any normative principles it would make sense to doubt. This is explained if oughts are built into characterizing the very meaning of a word like 'nothing,' oughts that comprise the logic of the word. It is explained if the logic of the word 'nothing' is a matter of certain oughts that govern the beliefs couched with the word. (15)

Gibbard's point in emphasizing that the entailment is *invariable* is to rule out *prudential* explanations for why one ought not to believe contradictions. Prudentially, one might well have reason to believe both that snow is white and that nothing is white (suppose you know that an evil demon will blow up the world unless you believe both those things).

But Gibbard thinks that, even if this were so, it would still be true, using Ewing's ought, that one ought not to believe those two things together.

Let us grant, then, that it is an analytic truth that

> (Belief Nothing, BN) One ought not to believe both that snow is white and that nothing is white.

Still, even granting this, there is a prima facie problem with this way of arguing for the normative nature of the MEANING concept.

Let's agree that (BN) is true. How does anything follow about the meaning of *linguistic expressions*? (BN) doesn't involve the meaning of anything linguistic at all, it would seem. It just says that one ought not both *believe* that snow is white and that nothing is white. Perhaps when we talk about beliefs we are talking about *propositions* or *thoughts* (although even this is controversial); still nothing about language has entered into the picture, yet. So, how did we get from "one ought not . . . believe both that snow is white and that nothing is white" to "presumptively . . . a *means* seems to qualify as normative" (16)?

Let me come back to the issue of how we are going to get *language* to enter the picture.

Perhaps what we have here is, at least in the first instance, an argument that *mental content* is normative, rather than that linguistic meaning is. Should we at least agree that it follows analytically from

> p is the proposition that snow is white and nothing is white

that

> One ought not to believe p?

While this may seem more plausible, it is also plausible that whatever norm is in the vicinity here, it stems from the nature of *belief* rather than from the nature of mental content.[3]

We have already seen reason to maintain that belief is a normative concept, from the fact that for someone to properly understand the concept of belief they need to understand that

(Norm on Belief-1) A false belief is defective.

Another plausible norm on belief, stated let's say in terms of Ewing's ideal primitive ought, might be this:

(Norm on Belief-2) If X has undefeated evidence that p, then X ought to believe that p.[4]

Since, in a suitably broad sense of 'evidence,' I have undefeated evidence that it can't be the case both that snow is white and that nothing is white, I have an explanation for why (BN) holds that appeals only to norms on belief and not any norms on meaning.

Not only is there no *positive* reason to attribute the normativity here to mental content, rather than to belief, there is reason *against* doing so. After all, mental content features in a host of *other* propositional attitudes besides belief, desiring and hoping, for example, to which *different* norms apply, if any. Wouldn't it make more sense, then, to think of these different norms as deriving from the distinct attitudes that they govern, rather than from the mental contents that they have in common?

In any case, how are we going to get from these observations about *belief* to the normativity of *linguistic meaning* (here I am returning to the question I had earlier postponed). Referring to our earlier exchange, Gibbard says:

> In arguing against the normativity of content and in favor of the normativity of belief, Boghossian dismisses the normativity of linguistic meaning as something that would have to be put in terms of norms of assertion. My own attempt in this book goes by a different route, which Boghossian does not consider. (17, n32)

We saw above that there are no analytic entailments from a meaning claim to an ought claim for *assertion*. From

I mean plus by '+'

it follows that

3 I am drawing here on my 2005.

4 I borrow this formulation from Hill 2013.

It is *correct* for me to say '68 + 57 + 125'

but *not* that

I *ought* to *say* '68 + 57 = 125'

even if this were Ewing's ideal ought. At best we could claim that

I ought to *believe* that 68 + 57 = 125.

Now, Gibbard's idea is that we can convert this into a defense of the normativity of linguistic meaning by claiming that, in many of these cases,

> Talk of believing thoughts . . . amounts to talk of accepting sentences in the thinker's own language, along with what those sentences mean. (27)

Let's formulate this as the following view:

> (Belief as Sentence Acceptance, BSA): To believe that p amounts to accepting a sentence of one's own language that means that p.

If (BSA) were true for all beliefs, then we would have an answer to the objection I was pressing that the norms in this area stem from belief and not meaning or content. If belief just amounts to acceptance of a sentence of one's own language, along with what it means, then the norm that

If X has undefeated evidence that p, then X ought to believe that p.

just amounts to:

If X has undefeated evidence that p, and X's sentence S means that p, then X ought to accept S.

And this does seem to be an example where a meaning claim implies an ought claim.

The trouble, of course, is that (BSA) is not true for beliefs in general. There are prelinguistic beliefs (as in infants) and nonlinguistic beliefs (as in many animals).

Indeed, even in linguistic creatures like us, I doubt very much that most belief is linguistic in this sense. As I visually survey the scene in front of me, I am forming many new beliefs. But as I now survey this scene, I am certainly not busy accepting sentences of English. If I were asked, I could express most of those beliefs in English, but that is to give voice to them, not to explain what they consist in.

Gibbard says he takes this talk of belief as sentence acceptance over from Paul Horwich. But, for Horwich, belief is constituted by the acceptance of sentences in a 'language of thought,' not by the acceptance of sentences of a public language. However, even if one found talk of a 'language of thought' acceptable, it would do very little to make plausible a claim like (BSA). To say that belief takes place in a language of thought is just a metaphorical way of saying that the states that constitute belief have something analogous to a syntactic structure, something with parts that can be combined and recombined in certain ways. That is very far removed from supposing that belief consists in some sort of psychological relation to the independently identifiable sentences of one's public language.

In any event, as long as we have to concede that belief talk in general is not reducible to talk of the acceptance of sentences, we will have to recognize norms for belief that are not reducible to norms on linguistic meaning.

But then, once we have done that, we may as well see all the norms in this area as emanating from the norms on belief, just as we have previously said.

This is all by way of questioning the chief initial motivation that Gibbard gives for taking seriously the 'meaning is normative' slogan. For all I've shown, it might be that once we appreciate the virtues of his overall view, we will be willing to pay whatever price these remarks have identified. However, I'm not there yet.

Back to Correctness

Let's go back to meaning as setting a standard of correctness for the application of words. As I've said, I am inclined to regard this as a genuine type of normativity; but I don't attach the same importance to this claim that Kripke did.

For Kripke, the question whether the notion of meaning is normative seems to have been of paramount importance because he thought that, if it could be shown to be true, it would have an immediate and negative impact on the *naturalistic reducibility* of meaning.

Thus, after giving many specific criticisms of the dispositional theory, Kripke says that it should have been obvious from the start that a dispositional theory could not be right because the relation between meaning and use is normative whereas the relation between a disposition and its exercise is descriptive (37).

As this remark shows, Kripke is thinking of the normativity of meaning as something that would *immediately* show that no dispositional analysis of a meaning claim could be right, because it so obviously misconstrues the relation between meaning and use. There are two respects in which I disagree with Kripke here.

The first is that the most one might extract from the point is something about the a priori reducibility of meaning to dispositions, via some sort of conceptual analysis. And that would fall very short of showing that meaning is not ultimately grounded in naturalistic facts, in the way that Gibbard favors.

The second is that even if we allow that correctness is normative, I don't believe that that would immediately show that a dispositional analysis is ruled out because, as I explained in previous writings (Boghossian 1989), any dispositional theory would have to include some appeal to a set of conditions that are 'ideal' conditions under which our cognitive mechanisms are incapable of making mistakes. Our ordinary dispositions with respect to a given term are bound to include dispositions to apply that term to things to which it does not apply. Thus, reading off the meaning from our ordinary dispositions is bound to lead to incorrect verdicts about what our terms mean. Hence, any dispositional theory will have to appeal to the dispositions that we have under ideal conditions. However, if meaning is constituted by dispositions under ideal conditions, then it seems possible to capture whatever normativity there exists here in the difference between how a person would ideally respond as compared to how she actually responds.

Meaning's Normative Role

Let's go back to the crucial passage in which Kripke enunciates his 'meaning is normative' slogan and let me quote it more fully.

> Suppose I do mean addition by '+' What is the relation of this supposition to the question how I will respond to the problem '68 + 57'? The dispositionalist gives a *descriptive* account of this relation: if '+' meant addition, then I will answer '125'. But this is not the proper account of the relation, which is *normative*, not descriptive. The point is *not* that, if I meant addition by '+' I *will* answer '125', but that . . . I *should* answer '125' . . . The relation of meaning and intention to future action is *normative*, not *descriptive*.
>
> In the beginning of our discussion of the dispositional analysis, we suggested that it had a certain air of irrelevance with respect to a significant aspect of the sceptical problem—that the fact that the sceptic can maintain the hypothesis that I meant quus shows that I had no *justification* for answering '125' rather than '5,. How does the dispositional analysis even appear to touch this problem? Our conclusion in the previous paragraph shows that in some sense . . . we have returned full circle to our original intuition. Precisely the fact that our answer to the question of which function I meant is *justificatory* of my present response is ignored in the dispositional account and leads to all its difficulties. (37)

Notice that Kripke clearly regards talk of the normative nature of meaning as *equivalent* to saying that our grasp of the meaning of a word must be able to *justify* our use of that word.

There are two possible problems with this claimed equivalence.

First, and in general, to say of X that it is constitutively normative and to say that it has a justificatory role are two different things. Something can have a *normative role* without itself *being* constitutively normative. It can have that role as a matter of necessity, and yet not itself *be* normative.

To see a clear example of this, consider perception. Perception has a normative role: it can justify belief. Indeed, if perception has a normative role, it necessarily has that role.

But there is no intuitive sense in which perception is a normative concept or state. You don't need to understand anything normative (in the prescriptive sense) to understand the idea of something being perceptually presented to you as being the case.

Second, in the particular case of meaning, meaning is constitutively normative, because of its necessary connection to correctness conditions; but its *justificatory* role is a *further* aspect of its nature. It follows, therefore, that we might be left with a problem about meaning's justificatory role even *after* we have said all that needs saying about its relation to correctness conditions.

To explain. Suppose I mean *green* by 'green,' and, on seeing a green thing under good lighting conditions, I utter the sentence 'Here is a green thing.' In this case, let's suppose, I correctly and justifiably believe that the thing is green.

Now, when I voice my belief by uttering the sentence 'Here is a green thing,' what I say is not only correct, given what I mean by 'green,' but is also *justified* by my grasp of the meaning of 'green.'

The point is that language use is a *rational activity*.[5] When I use the word 'green' to express my belief that something is green, I have a *reason* for using that word as opposed to another, a reason that is provided by what I mean by the word.

So, our grasp of the meaning of an expression has to be able to justify our use of that expression (just like perception has to be able to justify our beliefs). Call this the *Justificatory Thesis*.

Notice how all these claims about meaning are entirely parallel to the intuitive way in which we would characterize an instance of rule-following:

> If S is following a rule R on a given occasion by doing A, then (1) S has *accepted* R, (2) S's acceptance of R determines whether what S did was *correct*, (3) S's acceptance of R *explains* why S does A, and (4) S's acceptance of R rationalizes or justifies her doing A.

This is why, intuitively, meaning something by a word of a public language looks like a special case of following a rule with respect to that word. They have identical structures.[6]

In both cases, it looks as though there is a time at which a meaning is grasped (a rule internalized or accepted); and that grasp then stands in the following three relations to subsequent behavior: it sets a *standard* of correctness for that behavior; it explains it; and it *justifies* it.

I think these three aspects of both rules and meaning deeply puzzled Wittgenstein. And I agree that they are indeed deeply puzzling.

---

5 A point I think of as emphasized by Christopher Peacocke among others.

6 For reasons explained in my 1989, meaning things by words of mentalese, as opposed to a public language, could not have this structure.

They are undoubtedly puzzling on a dispositional view. In particular, the dispositionalist will clearly have a problem with the justificatory requirement. For how does my being disposed to apply 'green' to green things *justify* me in applying it to this green thing? Just because I am disposed to do something doesn't imply that I am justified in doing it.

One might think that ideal conditions could be invoked here to help out the dispositionalist. Couldn't we say that the justification for what I am disposed to say here is grounded in the fact that this is what I would say if conditions were ideal?

The problem with this reply, though, is that the justification that we intuitively have for the rational use of language is some sort of internalist justification: it must be something that I am aware of and whose relevance to my action I can somehow see. That is why talk of being *guided* by one's understanding of the meaning of a word is so intuitively compelling.

But this would not be true for a justification grounded in some counterfactual about what I would do if conditions were ideal. I don't normally know what those conditions are and don't know how to compute what I would do if I were in them.

In any event, the problem here is not unique to the dispositionalist but may be thought to present a quite general problem for any account of meaning. The general problem is this:

There is one case of the understanding supplying a justification for our using a word in a certain way that we sort of understand reasonably well. This is when my grasp of the meaning of a sentence, S, consists in my grasp of some sort of explicit definition for S; and my basis for assenting to S consists in my inferring S from its definition (what I have elsewhere called 'Frege analyticity'). We could give such a story for why I assent to "All squares are four-sided figures."

But this is clearly a very special case—special both in that grasp rarely consists in grasping an *explicit definition*, and in that basis rarely means *inferential basis*.

But if the relation between grasp and assent is not like that, what else could it be like?

The cases of one thing justifying another noninferentially that we understand well enough are the cases where a perceptual state that p justifies our belief that p. A perceptual state can serve as an epistemic basis for assent because it is a *presentation* of the world as being a certain way, a *seeming*—that feature is essential to its ability to justify belief.

But how could the relation between grasp of the meaning of 'green' and the application of 'green' be analogous to the relation that obtains in the perceptual case?

It is hard to see what could count as an internalistically acceptable answer to this question, if not a story that appeals, in the way that philosophers and linguists have found quite natural to do, to the notion of an intuition or an intellectual seeming. Our grasp of the meaning of 'green' supplies us with an intuition to the effect that it correctly applies to the presented object.

Of course, to say this is to supply the merest hint of where an answer to our question should be sought. We are still some distance from knowing exactly how to fill this story out.

Conclusion

My current view, then, is that the really difficult problem in the theory of meaning concerns how my grasp of the meaning of a word could *justify* future applications of that word.

This has often been conflated with the issue whether meaning is constitutively normative.

But the Justificatory Thesis and the Normativity of Meaning thesis are distinct claims, each interesting in its own right.[7]

References

Boghossian, P. (1989). "The Rule-Following Considerations." *Mind* 98, no. 392: 507–49.

——— (2005). "Is Meaning Normative?" In C. Nimtz & A. Beckermann (Eds.), *Philosophie und/als Wissenschaft—Hauptvorträge und Kolloquiumsbeiträge zu GAP.5*. Paderborn: Mentis.

Ewing, A. C. (1939). "A Suggested Non-Naturalistic Analysis of Good." *Mind* 48, no. 189: 1–22.

Gibbard, A. (2012). *Meaning and Normativity*. Oxford: Oxford University Press.

Glüer, K., and Å. Wikforss (2015). "Meaning Normativism: Against the Simple Argument." *Organon* 22: 63–73.

Hill, C. S. (2013). "Review of Allan Gibbard, Meaning and Normativity." *Notre Dame Philosophical Reviews*.

Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.

McPherson, T., and D. Plunkett (2017). "The Nature and Explanatory Ambitions of Metaethics." In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett, 1–28. London: Routledge.

Rosen, G. (1997). "Who Makes the Rules Around Here?" *Philosophy and Phenomenological Research* 57, no. 1: 163–71.

Wikforss, Åsa (2001). "Semantic Normativity." *Philosophical Studies* 102, no. 2: 203–26.

Wittgenstien, Ludwig (1953). *Philosophical Investigations*, G. E. M. Anscombe and R. Rhees (Eds.), G. E. M. Anscombe (trans.), Oxford: Blackwell.

7 I am grateful to the audience at the GibbardFest in Ann Arbor in spring 2016 and to Billy Dunaway, Paul Horwich, Josh Peterson, and David Plunkett for helpful comments.