SEMANTIC NORMATIVITY AND NATURALISM[*]

José L. Zalabardo
UCL

1. Kripke's normativity argument

In Wittgenstein on Rules and Private Language, Saul Kripke presented an argument, inspired by his reading of Wittgenstein, against the possibility of facts as to what someone means by a linguistic expression. The argument proceeded by considering a series of proposals as to what these facts might consist in, arguing in each case that the candidate account of meaning facts was unsatisfactory.

A proposal that figures prominently in Kripke's argument is the dispositional account of meaning. According to the dispositional account, the conditions of correct use associated with the meaning of an expression are determined by speakers' dispositions: how it would be correct to use an expression, as meant by a speaker, is determined, on this account, by how the speaker is disposed to use the expression.

Kripke raised several objections to dispositionalism. One that has received much attention in the literature arises from the normative status of facts about correct use. Here are some representative passages, focusing on the case of the meaning of the sign '+':

> The dispositionalist gives a *descriptive* account of this relation [between the meaning I ascribe to "+" and how I am disposed to answer the question "68 + 57 = ?"]: 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 rather that, if I intend to accord with my past meaning of "+", I *should* answer "125". […] the relation of meaning and intention to future action is *normative*, not *descriptive*. (Kripke 1982: 37)

> A candidate for what constitutes the state of my meaning one function, rather
> than another, by a given function sign, ought to be such that, whatever in fact I
> (am disposed to) do, there is a unique thing that I *should* do. Is not the
> dispositional view simply an equation of performance and correctness? (Kripke
> 1982: 24)

The interpretation of these and related passages is by no means straightforward, but many

have read them as advancing arguments that apply to semantics lines of reasoning originally

aimed at naturalistic accounts of ethical concepts (Boghossian 1989; Wright 1984; Blackburn

1984; McDowell 1984). In this paper I am going to consider two of these lines of reasoning. I

am not going to be concerned with the question whether they can be plausibly attributed to

Kripke.[1] My goal will be to determine whether these lines of reasoning can be construed as

providing cogent arguments against semantic naturalism.

2. <u>The naturalistic fallacy</u>

The first line of reasoning that I want to consider is based on the thought that there is an

unbridgeable logical gap between statements about how things are and statements about what

someone should do. David Hume appears to endorse this thought in the following famous

passage of the *Treatise*:

> In every system of morality, which I have hitherto met with, I have always
> remark'd, that the author proceeds for some time in the ordinary ways of
> reasoning, and establishes the being of a God, or makes observations
> concerning human affairs; when all of a sudden I am surpriz'd to find, that
> instead of the usual copulations of propositions, *is*, and *is not*, I meet with no
> proposition that is not connected with an *ought*, or an *ought not*. This change is
> imperceptible; but is however, of the last consequence. For as this *ought*, or
> *ought not*, expresses some new relation or affirmation, 'tis necessary that it
> shou'd be observ'd and explain'd; and at the same time that a reason should be
> given; for what seems altogether inconceivable, how this new relation can be a
> deduction from others, which are entirely different from it. (Hume 1978: 469)

Hume is here presenting a thought about arguments from premises about how things are to

conclusions about what someone should do. The thought appears to be that these arguments

---

[1] In (Zalabardo 1997) I argued against this attribution.

cannot be deductive, or, as I propose to construe the thought initially, that they cannot be underwritten exclusively by logic and conceptual analysis.

We can formulate the thought as concerning the character of the inferential principles that can validate these arguments. I am going to refer to the principles that concern us as *prescriptive principles*. Prescriptive principles will be statements of the form:

1. $(\forall x)(\forall S)(\phi(x, S) \supset S$ should/shouldn't do x),

where S ranges over subjects and x over actions, and ϕ is a propositional function that describes an action and possibly a subject and the relationship between subject and action.

I want to concentrate on a particular type of prescriptive principles—those in which the description from which a prescription is derived doesn't include the ascription to the subject of a desire, goal or pro-attitude. For this purpose I am going to introduce a distinction between hypothetical and categorical prescriptive principles. We shall say that a prescriptive principle is *hypothetical* just in case in order for x and S to satisfy ϕ, S needs to have some desire, goal or pro-attitude. I.e. if x and S satisfy ϕ, there is a desire, goal or pro-attitude that S has, such that in its absence x and S wouldn't satisfy ϕ. A prescriptive principle will be *categorical* if it is not hypothetical.

Thus, the obligations expressed by categorical prescriptive principles are independent of the subject's volitions. However, they are subject to be over-ridden by the demands placed on the subject's behaviour by other prescriptive principles. Let's focus on an instance of 1 specifying conditions under which you should do x. Even if it were true, and even if S and x satisfied ϕ, it might be that, all things considered, S should not do x, if for S to do x would be incompatible with what other prescriptive principles demand of S.

I want to formulate the thought that we cannot deduce a conclusion about what someone should do from premises about how things are as a claim about the status of categorical prescriptive principles. I shall refer to the claim as *Hume's thought*:

> *Hume's thought*: Categorical prescriptive principles cannot be established solely by means of logic and conceptual analysis.

Notice that Hume's thought is in principle neutral on whether or not there are true categorical prescriptive principles. On the one hand, an advocate of Hume's thought can maintain that there are true categorical prescriptive principles, so long as she doesn't take them to be derivable by logic and conceptual analysis. On the other hand, an advocate of Hume's thought can deny the existence of true categorical prescriptive principles. The rejection of categorical prescriptive principles can take two forms. One possibility is to deny the existence of facts as to what one should do over and above what follows from hypothetical prescriptive principles. The other is to assert the existence of categorical prescriptive facts, while denying that they supervene on descriptive facts. On this view, when S should do x there isn't a description of S and x such that whenever an action and a subject satisfy the description the subject should perform the action.

Hume's thought can be seen as posing an insuperable obstacle to the provision of naturalistic analyses of ethical notions. Suppose the following is a conceptual truth about moral rightness:

2.   If it would be morally right for you to perform an action, then in principle (i.e. unless the obligation is overridden by other imperatives) you should perform the action.

Now consider a naturalistic analysis of the concept of moral rightness, according to which there is a (categorical) description N of a relationship between an action and a subject such that for an action and a subject to satisfy this description is what it means for it to be morally

right for the subject to perform the action. This analysis would result in a straightforward violation of Hume's thought. $(\forall x)(\forall S)(N(x, S) \supset S$ should do $x)$ would be a true categorical prescriptive principle derivable from the conceptual analysis of moral rightness. If Hume's thought is correct, and if it is analytic that one should do what's morally right, then no naturalistic (categorical) analysis of moral rightness can be correct.

This is not as serious a blow to ethical naturalism as it might sound, since contemporary ethical naturalists don't typically see themselves as providing conceptual analyses of ethical concepts. They see themselves more commonly as trying to disclose the hidden essence of ethical concepts, not unlike the way in which science discloses the hidden essence of water or gold. On this conception of the task, a naturalistic account of an ethical concept would express a metaphysically necessary truth that could only be established a posteriori.

If a naturalistic account of moral rightness is conceived along these lines, it will no longer license violations of Hume's thought, since it won't allow us to establish $(\forall x)(\forall S)(N(x, S) \supset$ S should do $x)$ by conceptual analysis. As far as I can see, the only way to counter this move would be to strengthen Hume's thought in order to rule out the possibility of deriving categorical prescriptive principles from metaphysically necessary truths:

> *Hume's thought (second version)*: Categorical prescriptive principles cannot be
> established solely by means of logic, conceptual analysis or metaphysical necessities.

I am not going to discuss the plausibility of either version of Hume's thought. Instead I am going to assume for the sake of the argument that the second version of the principle is correct, in order to address the question, whether we can obtain from this a cogent argument against semantic naturalism.

3. The normativity of linguistic meaning

Discussion of the consequences of Hume's thought for semantics has generally focused on naturalistic accounts of the relation pairing each predicate with the property whose instances satisfy the predicate, as meant by a speaker at a time. When a predicate, as meant by a speaker, is related in this way to a property, I shall say that the predicate *refers* to the property, or that the speaker *refers* to the property by the predicate. Using this notion of predicate reference, we can now define predicate satisfaction as follows:

3.   P, as meant by S, is *satisfied* by a just in case a instantiates the property to which S refers by P.

Suppose now that the following is a conceptual truth about predicate meaning:

4.   If P, as meant by S, is satisfied by a, then S should apply P to a.

It follows that for every predicate P, subject S and object a, the following is a conceptual truth:

5. $(\forall x)(\forall S)$(x is an ascription of P to a by S & P, as meant by S, is satisfied by a $\supset$ S should do x)[2]

Now consider a naturalistic analysis of predicate reference, according to which there is a natural relation ρ between predicates, properties and speakers such that for P, Π and S to be ρ-related is what it means for S to mean Π by P. And let N be the naturalization of predicate

---

[2] In *an ascription of P to a*, as I am using the phrase, S asserts that a satisfies P.

satisfaction that results from the following definition: N(S, P, a) iff ρ(P, Π, S) & Πa.[3] Now

the analysis of predicate reference will enable us to derive the following from 5:

6.   (∀x)(∀S)(x is an ascription of P to a by S & N(S, P, a) ⊃ S should do x)

Clearly, the antecedent of 6 is a description of x and S, on the plausible assumption that being

an ascription of P to a is a description of x. Then it follows that 6 is a categorical prescriptive

principle. Furthermore, if predicate reference were correctly analysable in terms of ρ, 6 would

be a true categorical prescriptive principle derivable by conceptual analysis, and we would

have a violation of Hume's thought. Hume's thought appears to rule out naturalistic analyses

of predicate reference. The same situation will ensue if we think of the claim that predicate

reference is ρ as a necessary a posteriori truth, so long as we invoke the second version of

Hume's thought.

   One could try to resist this line of reasoning, as presented, by objecting to 4. The problem

is that it requires the subject to ascribe P to every object that satisfies it. This is, in general,

not a possibility. Hence, on the assumption that ought implies can, 4 cannot express a genuine

prescription (see (Hattiangadi 2006: 226)). There are several ways in which the argument

might be modified to overcome this objection. An appealing possibility is to replace 4 with an

injunction against ascribing P to objects that don't satisfy it:

7.   (∀x)(∀S)(x is an ascription of P to a by S & P, as meant by S, is not satisfied by a ⊃ S

         shouldn't do x)

7 doesn't fall foul of the ought-implies-can principle, since refraining from ascribing P to

objects that don't satisfy it is in principle within S's powers. And 7, together with a

---

[3] Here and elsewhere I use predicative terms as singular terms in propositions in which they
figure as objects of predication.

naturalistic analysis of predicate reference will yield a violation of Hume's thought similar to 6.

However, one could object to 7, as well as 4, on different grounds. One could argue that there simply is no categorical obligation for speakers to speak the truth or to refrain from speaking falsehoods. For speakers who have this as their goal, 7 is a legitimate maxim, but to those who have no desire to avoid saying falsehoods, 7 doesn't apply. If this is right, then to get a universally valid principle we need to replace 7 with a claim along the following lines:

8. $(\forall x)(\forall S)(x$ is an ascription of P to a by S & P, as meant by S, is not satisfied by a & S wants to avoid saying what's false $\supset$ S shouldn't do x)

I think that the status of 8 as a conceptual truth is unobjectionable. However, 8 would not bring naturalistic accounts of predicate reference into conflict with Hume's thought. It would indeed force us to treat the following as a conceptual/metaphysical truth:

9. $(\forall x)(\forall S)(x$ is an ascription of P to a by S & ~N(S, P, a) & S wants to avoid saying what's false $\supset$ S shouldn't do x)

The problem is that while 9 is a prescriptive principle, it is obviously of the hypothetical variety, and Hume's thought is not violated by treating hypothetical prescriptive principles as conceptual truths.[4]

It is important to bear in mind that this objection to 7 cannot simply rest on the possibility that the demands that it makes on the subject's behaviour might be outweighed by other prescriptive principles, as this would be compatible with 7, on our understanding of

---

[4] Notice that the same objection would apply to a version of the argument that replaced the objective obligation expressed by 7 with a subjective obligation:

$(\forall x)(\forall S)(x$ is an ascription of P to a by S & S believes that P, as meant by her, is not satisfied by a $\supset$ S shouldn't do x)

prescriptive principles. The claim is that these demands apply only to those that have a certain desire, and this is the sure sign that the categorical prescriptive principle is illegitimate.[5]

Another approach to developing the argument against naturalistic accounts of predicate reference would be to focus on prescriptive principles that express, not an obligation to speak the truth, or to refrain from speaking falsehoods, but an obligation to speak in ways that express the meanings that we want to convey.[6] This approach comes to prominence when we reflect on the contrast between factual mistakes, e.g., applying a predicate to an object that doesn't instantiate the property that the predicate refers to, and semantic mistakes, e.g., applying to an object a predicate that doesn't refer to the property that the speaker wants to attribute to the object. The claim that mistakes of the second kind are to be avoided can be expressed along the following lines:

10. $(\forall x)(\forall S)(x$ is an ascription of P to a by S & S refers to Π by P & S wants to avoid attributing Π to a in x $\supset$ S shouldn't do x)

The claim that 10 is a conceptual truth has more plausibility than the corresponding claim for 7. However, 10 doesn't bring naturalist accounts of predicate reference into conflict with Hume's thought. If predicate reference can be analysed in naturalistic terms, and if a parallel naturalistic treatment is available of the notion of which property a subject wants to attribute with a given predicate ascription, then 10 is undeniably a prescriptive principle. But just as clearly, given that its antecedent mentions S's desire to avoid attributing a certain property to the object of predication, 10 is a prescriptive principle of the hypothetical kind, and, as we know, Hume's thought is not in conflict with the conceptual status of hypothetical prescriptive principles.

---

[5] This point is made in (Hattiangadi 2006). See (Whiting 2007) for a reply.

[6] This approach is discussed in (Wikforss 2001: 209-12).

4.   The normativity of mental content

The normative character of semantic notions can also be invoked to undermine naturalistic accounts, not of linguistic meaning, but of mental content, and shifting our attention to the mind might remove some of the obstacles encountered in the linguistic version of the enterprise. This is not a fundamental change of subject matter. It would be surprising if semantic notions were to receive independent explications in the linguistic and the mental realm. Rather, we should expect that one set of notions is explicated in terms of the other. The only remaining question is whether our semantic theory treats linguistic meaning or mental content as fundamental, and while both options have advocates, starting with the mental is the favoured route among semantic naturalists.[7]

According to the representational theory of mind, the primary bearers of mental content are syntactically specifiable mental items analogous to linguistic expressions. On this account, we will have a category of mental representations that play the same role in the composition of thoughts that the predicates of a language play in the composition of sentences. The mental event of attributing a property to an object will consist, on this account, in tokening in a certain way a mental predicate that refers to this property.[8]

If we assume this picture of mental representation, we can easily develop a mental version of the argument we presented above against naturalistic accounts of predicate reference. However, this assumption is not indispensable.[9] We can also formulate a version of the

---

[7] This is the approach associated with Paul Grice. See (Grice 1957).

[8] See, for example, the Appendix to (Fodor 1987).

[9] See in this connection the debate between Colin McGinn and Paul Boghossian. (McGinn 1984: 147) argued that it was not possible to raise a problem of the normativity of mental content analogous to the problem of the normativity of linguistic meaning. (Boghossian 1989: 514) offers a reply that rests on the assumption that mental content has syntactically identifiable bearers. I am suggesting that this assumption is unnecessary.

argument for mental content that doesn't presuppose the representational theory of mind. The argument would be based on the thought that the following claim is a conceptual truth:

11. $(\forall x)(\forall S)(x$ is a mental attribution by S of Π to a & $\sim\Pi a \supset S$ shouldn't do x)

If this thought is correct, it will pose an obstacle to any naturalistic account of the notion of mental attribution of properties, whether or not the account construes the notion as involving syntactically identifiable types, along the lines of the representational theory of mind. Any such account would turn 11 into a conceptually true categorical prescriptive principle, in violation of Hume's thought.

Notice that 11 is the restriction to a class of beliefs with a particularly simple structure of the prescription not to believe falsehoods—to believe only truths—to which I'm going to refer as the *truth prescription*. Hence our question is, in effect, whether we can derive the truth prescription from logic, conceptual analysis or metaphysics.

Some authors have suggested that even if 11 were a conceptual truth, it wouldn't follow directly that mental content is normative in a sense that would pose an obstacle to a naturalist account of the notion. The thought is that the conceptual truth of 11 might in principle manifest the normative character, not of mental content, but of belief.[10]

We can see the issue more clearly if we focus on the version of 11 that we can formulate if we assume the representational theory of content:

12. $(\forall x)(\forall S)(x$ is an ascription by S of mental predicate M to a & M refers to Π & $\sim\Pi a$
     $\supset S$ shouldn't do x)

The thought is that the conceptual character of 12 could be made compatible with Hume's thought *either* by rejecting the possibility of a naturalistic account of the reference of mental

---

[10] (Boghossian 2003: 41-45) offers an interesting discussion of this issue.

predicates (the second conjunct) *or* by rejecting such an account of the notion of a mental state being an ascription of a mental predicate (the first conjunct). In the terminology used by advocates of the representational theory of mind, this is the contrast between accounts of mental content and accounts of propositional attitudes.

The idea, then, is that the semantic naturalist can accept Hume's thought and the conceptual character of 12, so long as she is prepared to forgo a naturalistic account of the notion of a mental state being an ascription of a mental predicate—of a mental predicate being 'tokened in the belief mode'.

This point is, strictly speaking, correct, but I suspect the naturalist won't find this route particularly appealing. The problem concerns the status that the resulting position would accord to facts about which mental states token a given predicate in the belief mode. These facts would have to be construed as non-descriptive, but this construal would have to be made compatible with a naturalistic account of facts about which property each mental predicate refers to. I am not going to explore the issue here, but it is hard to see how these views could be combined into an appealing package. I shall assume, accordingly, that if we could establish the conceptual status of 11, Hume's thought would be rendered incompatible with a naturalistic account of mental content.

Clearly, for the anti-naturalist it's not enough to establish the conceptual truth of a hypothetical imperative concerning mental attributions of properties, making the obligation to avoid attributing Π to objects that don't instantiate it contingent on S having the desire to avoid believing falsehoods. Another unpromising avenue would be to ground the obligation to avoid false mental property attributions on the untoward practical consequences of believing falsehoods. For we cannot expect to establish the practical undesirability of believing falsehoods exclusively on logical, conceptual or metaphysical considerations, unless a pragmatist account of truth is adopted. On any non-pragmatist account of truth, the

connection between true belief and the promotion of our goals will be a contingent matter. The anti-naturalist needs to establish on conceptual/metaphysical grounds a categorical obligation to refrain from believing what's false, independent of the subject's volitions or of the practical consequences of such beliefs.

5. Shah and Velleman on why we should believe the truth

The most prominent recent attempts to establish the truth prescription on conceptual grounds are due to Nishi Shah and David Velleman. They consist in two separate arguments, each presented in the first instance by one of these authors individually, but later developed jointly in co-authored work. My goal in this section is to provide an assessment of Shah and Velleman's arguments, summarising a more detailed discussion that I have presented elsewhere (Zalabardo 2010).

I want to start by considering Shah's argument for the truth prescription. It is based on the contention that accepting the truth prescription is a condition for possessing the concept of belief. He writes: "it is one of the conditions for possessing the concept of belief that one accept the prescription to believe that p only if p is true" (Shah 2003: 470). It follows from this that "a competent user of the concept of belief must accept the prescription to believe that p only if p is true for any activity that he conceives of as belief-formation" (Shah 2003: 470).

It is important to understand the contrast between Shah's contention and the claim that accepting the truth prescription is a condition for having beliefs. Shah is very careful to point out that his argument doesn't establish the latter claim. His argument concerns the conditions for having the concept of belief, not the conditions for forming beliefs, and, as Shah concedes, having the concept of belief is not required for forming beliefs.

This point renders the connection between Shah's contention and the truth prescription slightly problematic. If accepting the prescription were a condition for having beliefs, the link

would be straightforward: all believers would be subject to the prescription because those who didn't accept it wouldn't count as believers. But with Shah's claim things are slightly more complicated, since, according to him, someone who doesn't accept the truth prescription might still count as a believer. It would seem, then, that Shah's claim would establish the truth prescription only for subjects who have the concept of belief. But this restriction won't protect the naturalist from the consequences of Hume's thought. All it means is that Shah's claim will ground a prescription of the following form:

13. $(\forall x)(\forall S)(x$ is a mental attribution by S of Π to a & S has the concept of belief & $\sim\Pi$a

$\supset$ S shouldn't do x)

Since the naturalist can also be expected to endorse a descriptive account of who counts as having the concept of belief, her naturalistic account of mental property attribution will turn 13 into a categorical prescriptive principle, thus bringing it into conflict with Hume's thought.

Thus, for Shah, you are not conceiving of an attitude as belief unless you accept the prescription to avoid falsehood in it. Shah defends this claim with an inference-to-the-best-explanation argument. He argues that there is a phenomenon that would be adequately explained by the view that the concept of belief requires accepting the truth prescription, but would otherwise go unexplained, or receive only inferior explanations. The explanandum that plays this role in his argument is a phenomenon to which he refers as the *transparency* of *doxastic deliberation*. Doxastic deliberation is "deliberation about what to believe" (Shah 2003: 447), and the transparency of doxastic deliberation is the following feature of this activity:

> The deliberative question *whether to believe that p* inevitably gives way to the factual question *whether p*, because the answer to the latter question will determine the answer to the former. That is, the only way to answer the question *whether to believe that p* is to answer the question *whether* p. (Shah and Velleman 2005: 499)

According to Shah, treating acceptance of the truth prescription as necessary for having the concept of belief enables us to provide a satisfactory explanation of transparency, but if this view is rejected transparency will go unexplained. Here is how Shah proposes to explain transparency:

> […] a competent user of the concept of belief must accept the prescription to believe that p only if p is true for any activity that he conceives of as belief-formation. Because one accepts this prescription insofar as one is deliberating about *whether to believe that p*, determining *whether p is true* will be immediately imperative, to the exclusion of any other question, for anyone who entertains the deliberative question *whether to believe that p*. (Shah 2003: 470)

Shah's thought is that, since doxastic deliberation involves conceiving of the cognitive activity that it generates as belief formation, it requires accepting the truth prescription for this activity, and this makes the question whether to believe that p give way to the question whether p.

Shah's opponent is someone who thinks that accepting the truth prescription is not necessary for having the concept of belief—that you could in principle count as having the concept of belief even if you adhered to doxastic policies that are incompatible with the truth prescription, e.g. having only beliefs that maximize your overall utility. I am going to refer to this position as *relativism*.[11]

In order for Shah's argument to succeed, he needs to achieve two goals. First, he needs to convince the relativist that transparency is a genuine phenomenon for which an explanation is needed. Second, he needs to convince us that treating acceptance of the truth prescription as necessary for the concept of belief affords a better explanation of transparency than any available to the relativist.

Shah's first task then is to convince the relativist that transparency is a real phenomenon calling for an explanation. Notice that transparency postulates the existence of two different

---

[11] The position is independent of relativism about truth. It is a view about the goal of cognition, which is perfectly compatible with a non-relativist notion of truth.

enterprises: on the one hand, the cognitive enterprise of trying to answer the question whether p and, on the other, the deliberative enterprise of trying to answer the question whether to believe that p. Transparency is the phenomenon that obtains if the latter enterprise immediately gives way to the former.

But on the conception of the deliberative enterprise that the relativist can be expected to endorse at the outset, transparency doesn't seem to hold. For the relativist, we try to answer the question whether to believe that p by trying to determine whether believing that p would satisfy our doxastic criteria. For someone who subscribes to the truth prescription, this will consist in trying to answer the question whether p, but someone who subscribes to an alternative doxastic prescription, e.g. to believe only what maximizes her utility, would try to answer the deliberative question by trying to answer a different factual question—whether believing p would maximize her utility (call this enterprise *utilitarian reflection*). On this conception of doxastic deliberation, it is not generally a transparent enterprise.

Thus establishing that transparency is a genuine phenomenon requires invoking a different account of doxastic deliberation. Shah and Velleman make a proposal to this effect in their joint paper. To deliberate about whether to believe that p, Shah and Velleman tell us, is to engage "in reasoning that is aimed at issuing or not issuing in one's believing that *p* in accordance with the norm for believing that *p*" (Shah and Velleman 2005: 502).[12] As we have seen, Shah needs to show that utilitarian reflection doesn't count as doxastic deliberation on p. In light of this construal of the deliberative enterprise, Shah would achieve this goal if he could show that utilitarian reflection cannot be aimed at issuing (or not issuing) in your believing that p in accordance with the norm for believing that p.

In a different paper, Shah seems to put forward an argument that would yield this conclusion. He writes:

---

[12] See also (Shah 2006: 489).

> In the sense I have in mind, deliberating whether to believe that p entails intending to arrive at belief as to whether p. If my answering a question is going to count as deliberating whether to believe that p, then I must intend to arrive at belief as to whether p just by answering that question. I can arrive at the belief just by answering the question whether p; however, I cannot arrive at the belief just by answering the question whether it is in my interest to hold it. (Shah 2006: 482)

Here Shah seems to argue that the reason why I cannot aim at arriving at a belief as to whether p through utilitarian reflection is that I cannot achieve this goal just by means of utilitarian reflection.

This argument relies on a questionable principle to the effect that an activity cannot aim at an outcome if it can't bring it about. But even if we concede this principle to Shah, his reasoning runs into trouble. Let's assume for the sake of the argument that the only activity that can enable me to arrive at a belief as to whether p in the requisite way is answering the question whether p, and that it follows from this that answering the question whether p is the only activity that can be aimed at issuing or not issuing in belief as to whether p. On Shah and Velleman's construal of doxastic deliberation, this outcome would entail that the enterprise is transparent.

The problem for Shah is that the argument that he has used to establish transparency can also be used as a seemingly adequate explanation of the phenomenon that doesn't rest on the assumption that accepting the truth prescription is necessary for having the concept of belief. The alternative explanation would go as follows: (a) trying to answer the question whether p is the only activity that can enable me to arrive, in the requisite way, at a belief as to whether p; (b) it follows from this that it is the only activity that can be aimed at issuing or not issuing in a belief as to whether p, and (c) this entails, in turn, that doxastic deliberation can only take the form of trying to answer the question whether p, as transparency requires. In trying to convince his opponent that transparency is a real phenomenon, Shah has also supplied her with a seemingly adequate explanation of it. Hence Shah has failed to convince the relativist

that transparency is a genuine phenomenon that can only be explained by the assumption that accepting the truth prescription is necessary for having the concept of belief. Therefore he has failed to provide a cogent argument for the legitimacy of 11 (or 13).

Let me turn now, more briefly, to the argument originally developed by David Velleman. He has argued that "[…] an attitude doesn't qualify as a belief unless it […] has a tendency to be constrained by input in ways designed to ensure that it is true" (Velleman 2000: 255). It would follow from this that abiding by the truth prescription is necessary for having beliefs.

Velleman defends this view by contending that unless we introduce this condition in our definition of belief we won't be able to distinguish beliefs from other cognitive attitudes, such as imaginings and assumings (Velleman 2000: 247). The alternative that he considers is the motivational account of belief, which seeks to account for belief in terms of the dispositions to behaviour that beliefs motivate.

Velleman is undoubtedly right to claim that other cognitive attitudes besides belief can motivate behaviour. But this is not enough to establish his point. He needs to argue that, say, assuming or imagining that p has *the same* motivational role as believing that p, and I can't see that he provides adequate support for this point. To take one of the examples Velleman considers, imagining that you are an elephant as part of a game of make-believe may motivate you to behave in certain ways, but I don't think one can plausibly claim that it motivates you to behave in exactly the same ways as *believing* that you are an elephant would.[13] And so long as the motivational role of belief is different from the motivational roles of other propositional attitudes, it will be possible in principle to use these differences to single out belief, obviating the need to introduce regulation for truth in our account. I conclude that Velleman's argument

---

[13] Velleman appears to concede that the behavioural outcome will be different in each case (Velleman 2000: 272), but he claims that the differences are due not to a difference in the underlying behavioural disposition, but to the presence in the case of imagining of countervailing beliefs. However he doesn't offer an argument for explaining the difference in this way, rather than as arising from different behavioural dispositions.

doesn't provide adequate support for the claim that having beliefs requires abiding by the truth prescription.

6. <u>Stability</u>[14]

As I have argued in the preceding section, I think that Shah and Velleman have failed to support the claim that there is a universal prescription to believe only truths. I am not aware of any other extant proposals that succeed in this enterprise. However, there is an area which I think holds some promise for the view. In this section I want to sketch a line of reasoning that might go some way towards establishing 11 on purely conceptual grounds. I am going to argue that the doxastic policy based on the prescription expressed by 11 doesn't have many coherent alternatives.

Let me refer to the doxastic criterion according to which S should believe that p just in case p satisfies condition C (written Cp) as *the C-criterion*.[15] And let's say that S has adopted the C-criterion as her doxastic policy just in case, for every proposition p on which she wants to form an opinion, S aims to bring it about that she believes that p if and only if Cp. I am going to argue that the conditions that a criterion has to satisfy in order to generate an adequate doxastic policy might confer on the truth criterion a special status.

Consider first the effect that your doxastic policies can be expected to have on your practical policies. Suppose that your policy as to when to do φ is to do φ on occasion x just in case x satisfies condition C. What you end up doing in pursuit of this policy on a particular occasion (whether you φ or you don't φ) will depend on the belief that you form as to whether C is satisfied. Hence, if your nominal policy as to whether to φ is to do so just in case C is satisfied, your effective policy will be to do so just in case you believe that C is satisfied.

---

[14] This section can be skipped without loss of continuity.

[15] I am using propositions as their own names.

If you subscribe to the truth criterion, this point doesn't have any remarkable consequences: you aim to ϕ on occasion x just in case you believe Cx, but you aim to believe Cx just in case Cx. Hence your overall policy is still to ϕ on occasion x just in case Cx. Things are different, however, if you subscribe to a doxastic criterion other than the truth criterion. Then the policy of ϕing on occasion x just in case you believe Cx doesn't immediately collapse into the policy of ϕing on occasion x just in case Cx. Suppose that for propositions concerning whether C is satisfied your doxastic policy is the D-criterion (to believe p just in case Dp) Then your effective criterion with respect to ϕ will be to ϕ on occasion x just in case the proposition Cx satisfies condition D. In this way, which doxastic policy you adopt will affect your effective practical policies. Suppose, e.g., that your nominal dietary policy is to eat something just in case it is good for you, but your doxastic policy with respect to whether a food is good for you is to believe that it is just in case this makes you feel good about yourself. Then your effective dietary policy will be to eat something just in case believing that it is good for you makes you feel good about yourself.

Doxastic policies can have this kind of effect, not only on practical policies, but also on other doxastic policies. Suppose that your nominal doxastic policy with respect to a proposition p is to believe p just in case Cp, but your doxastic policy with respect to the proposition Cp is to believe it just in case it satisfies condition D. Then your effective doxastic policy with respect to p will be to believe it just in case DCp—or rather to believe p just in case DCp satisfies your doxastic criterion for *this* proposition…

These considerations have important consequences when we focus on global doxastic policies—policies determining the circumstances under which a subject would believe *any* proposition. Suppose that you have adopted a doxastic policy based on the C-criterion. Then you will aim to believe p just in case Cp, but you will aim to believe Cp just in case CCp, and so on ad infinitum. In this case, what will be your effective doxastic policy with respect to p?

Whether this question can receive a satisfactory answer depends on the behaviour of C. Let's say that the C-criterion is *completely stable* just in case, necessarily, for every p, Cp if and only if p. If the C-criterion is completely stable, we can plausibly argue that the infinite regress generated by a global doxastic policy doesn't pose a problem. When the C-criterion is applied to p, it dictates that you should believe that p just in case Cp, i.e., by the stability of the C-criterion, just in case p. When the scope of the criterion is extended to Cp, the verdict that it yields for p doesn't change. It dictates that you should believe that p just in case you should believe that Cp, and you should believe that Cp just in case CCp. Hence, by two applications of the stability of the C-criterion, we have that you should believe that p just in case p. By the same reasoning, if we extend the scope of the C-criterion along this hierarchy any number of times, the same verdict will ensue—you should believe that p just in case p. This suggests that when the criterion is applied along the whole infinite hierarchy, we will obtain the same outcome: you should believe that p if and only if p. Hence the infinite regress won't prevent the policy from issuing a verdict for every proposition.

In other cases, though, the regress will be problematic. Let me use $C^0(p)$ to denote p, and for every n, let $C^{n+1}(p)$ denote $CC^n(p)$. Let me say that C is *ultimately unstable* just in case, for every n there is an m greater than n such that it is not the case that necessarily $C^{m+1}(p)$ just in case $C^m(p)$. If C is ultimately unstable, then adopting a global doxastic policy based on the C-criterion will have disastrous consequences. No matter how far along the hierarchy we extend the application of the criterion, the verdict that we have reached at that point concerning whether one should believe p might be reversed at some later point as we continue to extend the scope of the criterion. At no point will we reach a definitive verdict that can't be reversed further along the hierarchy. In these circumstances, we seem forced to conclude that the criterion doesn't yield a verdict on whether one should believe that p when it is applied to

the whole infinite hierarchy. Hence an ultimately unstable criterion won't sustain a coherent global doxastic policy.

This is unquestionably an auspicious outcome for supporters of the truth prescription. The truth criterion is completely stable. Hence the infinite regress generated by global doxastic policies won't pose a problem for the policy based on the truth criterion. Furthermore, one could plausibly argue that the truth criterion is the only completely stable doxastic criterion. Let's assume this is right. Then, if completely stable criteria were the only doxastic criteria for which the regress doesn't pose a problem, we would have found a way of vindicating the idea that the truth prescription applies universally: only doxastic policies compatible with it would seem to overcome the problem.

However, this is not quite right. Completely stable criteria are not the only doxastic criteria to survive the regress, as they are not the only doxastic criteria that are not ultimately unstable. The C-criterion is *ultimately stable* (i.e. not ultimately unstable) just in case there is an n such that, for every m greater than or equal to n, necessarily $C^{m+1}(p)$ just in case $C^m(p)$. Now, completely stable criteria are ultimately stable for $n = 0$. But criteria that are ultimately stable only for higher values for n are not completely stable. Nevertheless, the reasoning that we have presented to show that completely stable criteria do not face a problem with the regress generated by global doxastic policies could be used to establish this result for every ultimately stable criterion. If the C-criterion is ultimately stable for n, then once the scope of the criterion has been extended to the nth term of the hierarchy, the verdict that it yields concerning whether one should believe that p will not be changed by further extensions. The policy will continue to dictate that one should believe that p just in case $C^n(p)$. In light of this, we can conclude that this is the verdict that the policy will sanction when the criterion is applied to the whole infinite hierarchy. Every ultimately stable doxastic criterion can avoid the difficulty posed by the infinite regress.

Are there any ultimately stable criteria that are not completely stable? I want to suggest that there might be at least one. It is certainly not the case for every p and every possible state of information that the evidence supports p just in case p. Misleading evidence is a distressingly common occurrence. It follows that the evidence criterion is not completely stable. Nevertheless one could plausibly argue that, necessarily, for every p and every possible state of information, the evidence supports p just in case the evidence supports the proposition that the evidence supports p. The right-to-left direction of this biconditional is fairly uncontroversial, as evidence for the hypothesis that there is evidence for p surely counts as evidence for p. And the left-to-right direction also has some plausibility. For one could argue that it is in the nature of evidence to be accessible. And this would seem to entail that the evidence won't support p unless there is evidence that it does. If this reasoning is correct, the evidence criterion will be ultimately stable (for n = 1).

In any case, even if the evidence criterion is ultimately stable, and even if other ultimately stable criteria exist, the line of reasoning that I have sketched has to count as a partial victory for the universality of the truth prescription. Most doxastic policies incompatible with it are ultimately unstable, and hence unsuitable as global doxastic policies. Take, for example, doxastic criteria based on the practical consequences of belief. It seems reasonable to suppose that it is not the case necessarily for any proposition p and subject S that believing that p maximizes S's utility just in case S's utility is maximized by the belief that believing that p maximizes her utility.

However, the discredit that the argument brings on ultimately unstable criteria has limits. There is nothing wrong, as far as the argument goes, with assessing beliefs with an ultimately unstable criterion, so long as we don't try to base our own doxastic policy on it. Even their use as the basis for a doxastic policy is not ruled out by the argument, so long as its scope is suitably restricted. I can use an ultimately unstable criterion for p, for the proposition that this

criterion is satisfied by believing that p, and so on. All that's needed is that sooner or later we reach a point in this sequence beyond which my doxastic policy is based on an ultimately stable criterion.

7. <u>The open-question argument</u>

Let me turn now to the second line of reasoning against semantic naturalism that I want to consider in this paper. Kripke's appeal to the normative character of meaning has sometimes been seen as putting forward a version of the open-question argument. The argument was originally presented by G.E. Moore against the possibility of defining or analysing good. Moore writes:

> The hypothesis that disagreement about the meaning of good is disagreement with regard to the correct analysis of a given whole, may be most plainly seen to be incorrect by consideration of the fact that, whatever definition be offered, it may be always asked, with significance, of the complex so defined, whether it is itself good. (Moore 1903: 15)

Moore is happy to accept that there may be necessary and sufficient conditions for something being good. What he doesn't accept is that these conditions could be taken as defining or analysing *good*. The reason that the quoted passage gives for this is that we can meaningfully ask whether these conditions are sufficient or, presumably, necessary for being good.

The argument features prominently the notion of whether it is possible to ask meaningfully whether say, maximizing utility is good. I propose to construe this notion in terms of whether or not having a certain belief involving a concept is a necessary condition for this to be the concept of *good*. In linguistic terms, the question is whether having a belief expressible with a sentence involving a predicate, say "good", is a necessary condition for this predicate to mean *good*. On this construal, the question "Is C good?" could not be meaningfully asked if someone who doesn't believe that C is a sufficient condition for the satisfaction of "good" cannot mean *good* by "good". Having this belief shouldn't be taken to require entertaining the

proposition in consciousness. The belief could be manifested instead in your linguistic dispositions by your unwillingness to contemplate the possibility of withholding "good" from something that you believe to satisfy C.

We can now provide a formulation for the open question argument, concentrating again on the concept *morally right*:[16]

14.   For every description D, it is possible for someone who means *morally right* by "morally right" to believe, for some action x, both D(x) and "~Morally right (x)".[17]

15.   If satisfying D is analytically sufficient for being morally right, then, for every x, believing D(x) and "~Morally right (x)" is incompatible with meaning *morally right* by "morally right".

Therefore:

16.   For every description D, satisfying D is not analytically sufficient for being morally right.

And a parallel argument can be formulated against the possibility of analysing predicate satisfaction:

17.   For any description D, it is possible for someone who means *satisfies* by "satisfies" to believe, for some S, P, a, both D(S, P, a) and "~satisfies (S, P, a)".

---

[16] I am not claiming that Moore would have applied to *morally right* the argument that he develops for *good.*

[17] Believing "ϕ (x)", for some x, is believing the proposition expressed by a substitution instance of "ϕ (x)".

18.   If satisfaction of D by S, P and a, is analytically sufficient for a satisfying P, as meant

by S, then for all S, P, a, believing D(S, P, a) and "~satisfies (S, P, a)" is incompatible

with meaning *satisfies* by "satisfies".

Therefore:

19.   For every description D, D(S, P, a) is not analytically sufficient for a satisfying P, as

meant by S.

Clearly, the cogency of each of these arguments depends on whether we are capable of

providing adequate support for its first premise. We shall address this point later on. But a

more immediate problem is that even if these arguments succeeded in establishing their

conclusions, the naturalist could easily dismiss them as irrelevant to her project. As we

mentioned above, contemporary naturalists don't typically see themselves as providing

analytically necessary and sufficient conditions for the satisfaction of concepts. Instead, they

tend to see the satisfaction of the concepts they are accounting for as connected by a

necessary a posteriori link to the conditions that figure in the account. The model for this

conception of their task is the semantics of natural kind terms as explicated by Saul Kripke

and Hilary Putnam (Kripke 1980; Putnam 1975). On this conception of naturalist accounts, an

account of "morally right" or "satisfies" will seek to achieve the kind of understanding of,

say, "water" provided by the discovery that water is H_2O. Clearly, on this conception of the

enterprise, the conclusion of an open-question argument has no power to undermine an

account of what determines the satisfaction conditions of "water". The point is made very

clearly by Gilbert Harman for the case of ethics:

> […] as it stands the open question argument is invalid. An analogous argument
> could be used on someone who was ignorant of the chemical composition of
> water to "prove" to him that water is not H_2O. This person will agree that it is
> not an open question whether water is water but it is an open question, at least
> for him, whether water is H_2O. Since this argument would not show that water
> is not H_2O, the open question argument in ethics cannot be used as it stands to

> show that for an act to be an act that ought to be done is not for it to have some natural characteristic C. (Harman 1977: 19)

On our construal of the notion of what counts as an open question, Harman's point is that the satisfaction conditions of "water" are determined by H_2O, even though meaning water by "water" does not require believing that being H_2O is necessary for satisfying "water".

A similar point is made by Scott Soames concerning the effect of the argument on naturalistic accounts in semantics:

> […] I am willing to grant that the skeptic might be right in maintaining that claims about what I meant are not a priori consequences of nonintentional truths. If it were clear that any necessary consequence of a set of claims P was also an a priori consequence of P, then this admission would provide the sceptic with just what he needs; for then he could force me to admit that claims about meaning might not be necessary consequences of nonintentional truths. […] However, this argumentative strategy fails. Thanks to the work of Kripke and others, it has become clear that many necessary consequences of propositions are not a priori consequences of them. Consequently, my admission that claims about meanings may not be a priori consequences of nonintentional truths need not undermine my belief that they are necessary consequences of those truths. (Soames 1998: 231)

Thus what these authors are proposing is that naturalists can sidestep the open-question argument by modelling the connection between their explicanda and their explicantia on the connection between natural-kind terms and the properties that determine their satisfaction conditions. Analysis of the concept *water* doesn't suffice for determining that it is satisfied by samples of H_2O. Similarly, for the enlightened naturalist, analysis won't suffice for determining that *morally right* is satisfied by, say, actions that maximize utility, or that *refers* is satisfied by a predicate and a property whenever, say, tokenings of the latter cause tokenings of the former under normal conditions. It follows that even if these naturalistic accounts are correct, the descriptions "x maximizes utility" or "tokenings of x cause tokenings of y under normal conditions" won't provide counterexamples to the conclusion of the open-question argument.

This point strikes me as incontestable. I want to suggest, however, that the antinaturalist might still have some room for manoeuvre. Going back to *water*, the point that the description H_2O is not analytically sufficient for the satisfaction of *water* is not open to question. But the antinaturalist might accept this point, as she should, and still contend that the behaviour of *water* as a natural-kind term renders other descriptions analytically sufficient for its satisfaction. Let me follow Putnam in using the term *operational definition* to refer to the array of criteria that a speaker (or a community) uses for applying a natural-kind term (Putnam 1975). On the standard account of the semantics of a natural-kind term, the property to which it refers is singled out as the underlying kind present in most samples in the speaker's environment that satisfy her operational definition for the term.

My suggestion is that the antinaturalist could argue that the fact that the reference of a natural-kind term is determined in this way is discoverable by analysis—independently of any fact that can only be ascertained by empirical investigation. We need empirical research to ascertain which property plays this role, but not to ascertain that the property that plays this role, if there is one, will determine the extension of the term. But if this is correct, it turns out that there is, after all, a description that is analytically sufficient for the satisfaction of *water*—namely, *x is an instance of the kind, if there is one, present in most of the samples in the speaker's environment that satisfy her operational definition of water*.[18]

Providing adequate support for the claim that this description should be treated as analytically sufficient for the satisfaction of *water* lies well beyond the scope of this paper.[19] What I am arguing is that *if* this claim could be adequately established, then the threat of the

---

[18] Notice that the analyticity of this connection is compatible with the possibility that different speakers have different operational definitions of *water*, so long as, in their respective environments, they all track the presence of H_2O.

[19] Putnam, for one, thinks that it's not possible to obtain analytically sufficient conditions for the satisfaction of a natural-kind term (Putnam 1975), but the difficulties that he raises for the candidates that he considers (Putnam 1975) do not invalidate my proposal.

open-question argument to naturalist accounts would be reinstated. For the naturalist would have to explain how, say, maximizing utility, is singled out as the property that determines the extension of *morally right*. This explanation may or may not be the same as the one given for *water*, but in either case it can be expected to invoke a second-order property M whose satisfaction by *utility maximization* confers on the latter its extension-fixing role. And if the connection between *morally right* and M can be claimed to be analytic, we would have, as in the case of *water*, an analytically sufficient description: *x is an instance of the property, if there is one, that satisfies M*.[20] The same situation could be claimed to obtain in semantics. Once again, a second-order property T would have to be invoked to explain how, say, *tokenings of x cause tokenings of y under normal conditions* is singled out for the role of fixing the extension of *refers*. And if the connection between T and *refers* is analytic, we will appear to have an analytically sufficient description: *being a predicate-property pair instantiating the relation, if there is one, that satisfies T*.

Hence, even if *morally right* and *refers* are treated as natural-kind terms, they will have analytically sufficient descriptions. Therefore, invoking this semantic model won't enable the enlightened naturalist to sidestep the conclusion of the open-question argument. Her account will still be in conflict with the conclusion of the argument.

8. Normativity?

In the preceding section I have argued that the conclusion of an open-question argument for a concept C might still be in conflict with a naturalistic account of C, even if the account is construed along the lines of the natural-kind model. Obviously this would pose a threat for the naturalist only if the argument were sound, and the soundness of the argument in each case

---

[20] An account of moral concepts along these lines has been defended by Frank Jackson. See (Jackson 1997).

cannot be plausibly asserted until support is provided for its first premise—the principle that there is no description that has to be believed to be sufficient for the satisfaction of the concept by anyone who has the concept. I want to close by considering briefly how this principle could be supported in ethics and in semantics.

The standard approach to this task in ethics is to appeal to the normative dimension of ethical concepts. For the case of *good*, Darwall, Gibbard and Railton express the thought in the following terms:

> Attributions of goodness appear to have a conceptual link with the guidance of action, a link exploited whenever we gloss the open question 'Is P really good?' as 'Is it clear that, other things being equal, we really ought to, or must, devote ourselves to bringing about P?'. (Darwall, Gibbard, and Railton 1992)

For *morally right*, the same thought is expressed by 2, above.

This link enables us to invoke the intuition that it is perfectly possible to imagine, for any action description, someone who accepts the description but doesn't think that the action ought to be performed. As Darwall, Gibbard and Railton put it:

> Our confidence that the openness of the open question does not depend upon any error or oversight may stem from our seeming ability to imagine, for any naturalistic property R, clear-headed beings who would fail to find appropriate reason or motive for action in the mere fact that R obtains (or is seen to be in the offing). (Darwall, Gibbard, and Railton 1992)

These considerations suggest an argument for 14 with the following premises:

20.   For every action description D, it is possible to believe, for some x, both Dx and that x does not have to be done.

21.   You don't mean *morally right* by "morally right" unless, for every x, if you believe that x doesn't have to be done, you believe "~Morally right (x)".

For predicate reference we might try the same approach, deriving 17 from the following premises:

22.   For any description D, it is possible to believe, for some S, P and a, both D(S, P, a) and that S doesn't have to ascribe P to a.

23.   You don't mean *satisfies* by "satisfies" unless, for all S, P, a, if you believe that S doesn't have to ascribe P to a then you believe "~satisfies (S, P, a)".

Notice, thought, that this approach would take us back to the difficulties that we encountered trying to use Hume's thought to undermine semantic naturalism. For 23 and related principles would have no plausibility unless the truth prescription could be established by conceptual analysis.

However, I want to suggest that invoking the normativity of meaning at this point is not mandatory. The alternative that I want to put forward, in closing, is that the role that normativity plays in the ethical case could be played in semantics by disquotation. What we would need to invoke is the claim that having the concept of satisfaction requires believing that a is not satisfied by P, as meant by you, whenever you believe ~Pa. This claim would supply the argument with the following premise.

24.   You don't mean *satisfies* by "satisfies" unless, for all P, a, if you believe ~Pa you believe "~satisfies (I, P, a)".[21]

Now, to complete the argument for 17 we would only need to argue that for any description D, it is possible to believe D(I, P, a) while not believing Pa. This would give us the following principle:

25.   For any description D, it is possible to believe, for some P and a, both D(I, P, a) and ~Pa.

---

[21] "I" here is the first-person pronoun.

And 25 and 24 would establish 17 in the same way in which 20 and 21 establish 14.

Following this approach we would be able to mount a version of the open-question argument

in semantics that doesn't rely on the assumption that meaning is normative.

REFERENCES

Blackburn, Simon. 1984. The Individual Strikes Back. *Synthese* 58:281-302.
Boghossian, Paul. 1989. The Rule-Following Considerations. *Mind* 98:507-49.
———. 2003. The Normativity of Content. *Philosophical Issues* 13:31-45.
Darwall, Stephen, Allan Gibbard, and Peter Railton. 1992. Toward *Fin de siècle* Ethics: Some Trends. *Philosophical Review* 101:115-89.
Fodor, Jerry A. 1987. *Psychosemantics. The Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass.: MIT Press.
Grice, H.P. 1957. Meaning. *Philosophical Review* 66: 377-88.
Harman, Gilbert. 1977. *The Nature of Morality. An Introduction to Ethics*. New York: Oxford University Press.
Hattiangadi, Anandi. 2006. Is Meaning Normative? *Mind and Language* 21:220-40.
Hume, David. 1978. *A Treatise of Human Nature*. 2nd ed. ed. Oxford: Clarendon.
Jackson, Frank. 1997. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford: Oxford University Press.
Kripke, Saul. 1980. *Naming and Necessity*. Oxford: Blackwell.
———. 1982. *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.
McDowell, John. 1984. Wittgenstein on Following a Rule. *Synthese* 58:325-63.
McGinn, Colin. 1984. *Wittgenstein on Meaning*. Oxford: Basil Blackwell.
Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
Putnam, Hilary. 1975. Is Semantics Possible? In *Mind, Language and Reality. Philosophical Papers, Volume 2*. Cambridge: Cambridge University Press.
———. 1975. The Meaning of 'Meaning'. In *Mind, Language and Reality. Philosophical Papers, Volume 2*. Cambridge: Cambridge University Press.
Shah, Nishi. 2003. How Truth Governs Belief. *Philosophical Review* 112:447-82.
———. 2006. A New Argument for Evidentialism. *Philosophical Quarterly* 56:481-98.
Shah, Nishi, and J. David Velleman. 2005. Doxastic Deliberation. *Philosophical Review* 114:497-534.
Soames, Scott. 1998. Skepticism about Meaning: Indeterminacy, Normativity, and the Rule-Following Paradox. *Canadian Journal of Philosophy Supplementary Volume* 23:211-49.
Velleman, J. David. 2000. On the Aim of Belief. In *The Possibility of Practical Reason*. Oxford: Clarendon Press.
Whiting, Daniel. 2007. The Normativity of Meaning Defended. *Analysis* 67:133-40.
Wikforss, Åsa Maria. 2001. Semantic Normativity. *Philosophical Studies* 102:203-26.
Wright, Crispin. 1984. Kripke's Account of the Argument against Private Language. *Journal of Philosophy* 81:759-77.
Zalabardo, José L. 1997. Kripke's Normativity Argument. *Canadian Journal of Philosophy* 27 (4):467-88.

———. 2010. Why Believe the Truth? Shah and Velleman on the Aim of Belief. *Philosophical Explorations* 13:1-21.