

Noncognitivism and Epistemic Evaluations

Bob Beddor

Forthcoming in *Philosophers' Imprint*

Abstract

This paper develops a new challenge for moral noncognitivism. In brief, the challenge is this: beliefs—both moral and non-moral—are epistemically evaluable, whereas desires are not. It is tempting to explain this difference in terms of differences in the functional roles of beliefs and desires. However, this explanation stands in tension with noncognitivism, which maintains that moral beliefs have a desire-like functional role. After critically reviewing some initial responses to the challenge, I suggest a solution, which involves rethinking the functional relationship between desire and belief.

1 Noncognitivism as Boxology

In simpler times, it was easy to distinguish moral cognitivists from noncognitivists. Cognitivists believed in moral beliefs; noncognitivists did not.¹ But these days most noncognitivists want to ‘save the appearances’ of cognitivism: they want to allow that belief reports such as ‘Jane believes that stealing is wrong’ can be true. Semantically descending, most contemporary noncognitivists are willing to say that there are moral beliefs.²

How, then, are we to tell the two views apart? While contemporary noncognitivists allow for moral beliefs, they typically hold that moral beliefs are much more similar to desires than they are to prosaic beliefs (*grass is green* and the like).³ In particular, noncognitivists usually maintain that all it is to believe that stealing is wrong is to have some conative attitude towards stealing, for example, to desire that no one steals, or to disapprove of stealing, etc.

Given this, it’s natural to interpret noncognitivism as a thesis about the functional role of moral beliefs. Let the ‘B-Role’ refer to the functional role of prosaic beliefs. Let the ‘D-Role’ refer to the functional role of desires. Noncognitivism amounts to the view that the functional role of moral beliefs resembles the D-Role much more closely than the B-Role. (See Fig. 1.) On this view, belief isn’t a natural psychological kind: just as

¹For influential statements of ‘old-school’ noncognitivism, see Ayer 1936: 108; Russell 1935: 231–232.

²See, for example, Blackburn 1993, 1998; Gibbard 2003.

³I borrow the term ‘prosaic’ from Gibbard (2003), who uses it to refer to non-normative beliefs.




Figure 1: Dueling Boxologies

‘jade’ picks out both jadeite and nephrite, so ‘belief’ picks out both prosaic belief and some desire-like state.⁴

While seldom formulated in exactly these terms, it’s common to find noncognitivism characterized along roughly these lines. For example, a number of authors characterize noncognitivism as the view that moral beliefs have a desire-like direction of fit.⁵ According to this characterization, prosaic beliefs have a mind-to-world direction of fit: they aim to accurately represent the world. Moral beliefs are like desires in that they have world-to-mind direction of fit: they aim not to represent the world, but to change it. If we understand the difference between two states’ direction of fit as a difference between their functional roles (as many of these authors do), then this way of understanding noncognitivism amounts to a version of my boxological construal.⁶

⁴For the analogy with jade, see Ridge 2009, 2014.

⁵E.g., Ridge 2006a; Schroeder 2010: chp.5; Zangwill 2011; Streumer 2013; Björnsson and McPherson 2014.

⁶Why focus on functional roles, rather than direction of fit? Mainly to forestall confusion. Some authors prefer to understand direction of fit in explicitly normative terms (Anscombe 1957; Platts 1979; Shafer-Landau 2003; Gregory 2012). On such views, a state’s direction of fit is not a matter of its functional role; instead, it’s a matter of whether that state ought to conform to the world, or whether the world ought to conform to it. Furthermore, even those who understand direction of fit in non-normative terms often construe this notion in very different ways. (Compare, for example, the account provided by Smith (1987, 1994) with the ‘higher-order’ approach in Humberstone 1992.) Thus to prevent confusion, I’ll formulate my discussion entirely in terms of functional roles, leaving others to translate it into ‘direction of fit’ talk as they see fit.

A boxological characterization of noncognitivism also seems natural in light of the arguments that noncognitivists employ. A standard argument for noncognitivism starts by appealing to motivational internalism: the thesis that there is a necessary connection between having a moral belief and being motivated to act on it. This is then conjoined with a Humean theory of motivation, according to which motivation requires the presence of a desire. This ‘Argument from Motivation’, whatever its merits, naturally suggests a conception of noncognitivism along the lines I’ve sketched—as the view that the functional role of moral beliefs closely resembles that of desires.⁷

In this paper, I present a new challenge for noncognitivism, thus construed: the Epistemic Evaluability Challenge. The challenge starts with the observation that both moral and prosaic beliefs are epistemically evaluable, whereas prototypical desires are not. It’s tempting to explain this difference in epistemic evaluability in terms of some difference in the functional roles of beliefs and desires. But this explanation stands in tension with the noncognitivist idea that the functional role of moral beliefs resembles the D-Role far more closely than the B-Role.

This paper proceeds as follows. §2 develops the Epistemic Evaluability Challenge in more detail. §§3-6 consider initial responses to the challenge and find them wanting. §7 develops what I take to be the most promising response. The response I advocate is to advance a new version of noncognitivism—what I call ‘grounding noncognitivism.’ On this view, moral and prosaic beliefs have the same functional role. What distinguishes moral from prosaic beliefs is that the former are fully grounded in desire-like states. I suggest that this view may offer just what is needed to resolve the challenge.

2 The Epistemic Evaluability Challenge

2.1 The Challenge Expounded

Prosaic beliefs are epistemically evaluable. In everyday life, we talk about whether prosaic beliefs amount to knowledge—an epistemic evaluation *par excellence*. We also epistemically evaluate prosaic beliefs in other ways—as rational or irrational, as justified or unjustified.

By contrast, prototypical desires do not seem to be epistemically evaluable. Suppose I desire to slake my thirst. Perhaps this desire can be evaluated as rational or irrational, as justified or unjustified. But if so, it seems to be *practical* rationality and justification that’s at issue. Moreover, describing a desire as an item of knowledge seems like a category mistake: I can know *that* I want a drink, but my want isn’t itself an item of knowledge.

This raises what I’ll call the ‘Epistemic Evaluability Question’ (EEQ):

EEQ Why are beliefs epistemically evaluable, whereas desires are not?

⁷For uses of the Argument from Motivation, see e.g., Stevenson 1937: 16; Blackburn 1998: 61; Gibbard 2003: chp.7. Of course, it remains controversial whether motivational internalism is true, as well as whether its truth would support noncognitivism. For relevant discussion, see Darwall et al. 1992; Svavarsdóttir 1999; Björnsson et al. 2015.

It's natural to try to answer EEQ in terms of functional roles. More precisely:

Functional Role Hypothesis Beliefs are epistemically evaluable in virtue of some features of their functional role. Desires lack these features, and so are not apt for epistemic evaluation.

Two points of clarification are in order. First, the Functional Role Hypothesis does not claim that beliefs are the *only* epistemically evaluable states. At the very least, we should allow that *degrees of belief* are epistemically evaluable. My .5 credence that it will rain can be epistemically rational; on some views, it can even constitute knowledge (Moss 2013). One might also maintain that a psychological state can be epistemically evaluated—at least in certain respects—if it entails having some belief, or some degree of belief. For example, *fearing that a burglar will break in* seems to entail having a non-negligible credence that a burglar will break in. Perhaps if it's epistemically irrational to have a non-negligible credence that a burglar will break in, this makes the corresponding fear epistemically irrational. None of this is ruled out by the Functional Role Hypothesis.

Second, the Functional Role Hypothesis does not tell us which features of the B-Role suffice for epistemic evaluability. To see how one might flesh out the Functional Role Hypothesis, it is worth taking a brief excursus through standard functionalist accounts of belief. On standard accounts, the functional role of belief includes input and output conditions. The input conditions are often thought to involve some connection between beliefs and appearances, e.g.:

Ceteris paribus, if it appears to someone that p , they'll be at least somewhat disposed to believe p .⁸

The output conditions are usually thought to include some link between belief and action, e.g.:

Ceteris paribus, if someone believes p , they'll be disposed to treat p as true for the purposes of practical reasoning.⁹

This bare-bones story could be complicated in various ways. One could add further input and output conditions. One could also add internal role conditions specifying connections between beliefs and other mental states.¹⁰

Let C be some input, output, or internal role condition, or some combination of such conditions. There will be a possible view that maintains C is what makes beliefs epistemically evaluable. Of course, some candidate conditions will be more plausible than

⁸See Loar 1981 and Smith 1987. A closely related suggestion is that the input conditions include some connection between beliefs and evidence. I discuss this proposal in more detail in §5.3.

⁹This statement of the connection between belief and action is rather vague, and there are various options for formulating it more precisely (an issue that I take up in §7). For discussion, see Armstrong 1973; Loar 1981; Stalnaker 1984; Railton 2014, among many others.

¹⁰See, for example, Schwitzgebel 2002.

others. My own suspicion is that any plausible candidate will include some reference to the connection between belief and action.¹¹ However, there's no need to prejudge this issue. All that the Functional Role Hypothesis requires is that the epistemic evaluability of a state (or lack thereof) is explained by some features of that state's functional role; it does not take a stand on which features bear the explanatory load.

Our clarifications finished, we are now in a position to see why the Functional Role Hypothesis stands in tension with noncognitivism. According to our boxological construal, noncognitivism maintains that the functional role of moral belief resembles the D-Role much more closely than the B-Role. Given the Functional Role Hypothesis, noncognitivists seem to be committed to the conclusion that moral beliefs are not epistemically evaluable.

But this is surely wrong. Our everyday conversation is rife with talk of moral knowledge. We say things like, 'You knew you shouldn't have done that'; we talk about teaching children to know the difference between right and wrong. Our practice of ascribing moral knowledge is even enshrined in the law. Take, for example, the M'Naghten rule, which makes criminal liability dependent on whether the accused knew that their action was wrong.

Arguably, moral beliefs can be epistemically evaluated in other ways as well. It's natural to describe moral beliefs as rational or irrational, justified or unjustified: 'It's irrational to believe that euthanasia is wrong'; 'He was justified in thinking it was the right course of action.' Admittedly, it's less clear whether these descriptions are distinctly *epistemic* evaluations. Perhaps, some may suggest, when we describe moral beliefs in these ways, we're really talking about *practical* justification/rationality. However, I think there are grounds for push-back. In the epistemology literature, it's widely held that in order for a belief to amount to knowledge, it must be *epistemically* justified/rational; practical justification/rationality is not enough.¹² Indeed, some have even proposed defining epistemic rationality as the species of rationality that is required for knowledge.¹³ If these views are correct, then any moral belief that qualifies as knowledge will also qualify as epistemically justified/rational.

Common sense, then, takes moral beliefs to be epistemically evaluable. They qualify as knowledge; arguably, they also qualify as epistemically rational and epistemically justified. In these regards, they resemble prosaic beliefs and differ from prototypical desires. Noncognitivists who seek to accommodate the realist trappings of moral thought must either provide an alternative answer to EEQ—an answer that doesn't make epis-

¹¹While desires play a role in guiding action, their role is fundamentally different. Desiring *p* does not dispose one to treat *p* as true; rather, it disposes one to *make p* true. More on this in §7.

¹²We can motivate this by considering cases where the practical rationality of a belief pulls apart from its epistemic rationality. Suppose Jones believes that he will win the race, even though his evidence strongly indicates he will lose. Suppose, moreover, that this belief gives him a burst of confidence, causing him to win. At the start of the race, Jones' belief is practically rational, in virtue of its beneficial consequences. Intuitively, however, his belief does not qualify as knowledge. And the reason for this seems to be that it is epistemically irrational.

¹³See e.g., Schroeder 2015.

temic evaluability dependent on functional role—or explain how, despite appearances, the Functional Role Hypothesis can be reconciled with noncognitivism. Call this the ‘Epistemic Evaluability Challenge.’

The Epistemic Evaluability Challenge is first and foremost a problem for noncognitivism. Is it also a problem for *expressivism*? This depends on how one conceives the relation between expressivism and noncognitivism. Typically, expressivism is characterized as the thesis that moral discourse does not purport to represent the world. While expressivism and noncognitivism often march hand-in-hand, at least some writers have argued they can be separated. For example, Horgan and Timmons (2006) advocate a ‘cognitivist expressivism’, which combines an expressivist semantics with a cognitivist boxology. I will not take a stand on whether this is a stable combination. However, insofar as it is, the view evades the Epistemic Evaluability Challenge.¹⁴

In the rest of this section, I further clarify the Epistemic Evaluability Challenge by considering two natural concerns. The first is that the challenge has already been solved. In a number of places, Blackburn and Gibbard have sketched noncognitivist accounts of moral knowledge, which some might think can be used to answer the Epistemic Evaluability Challenge. The second concern is that the Epistemic Evaluability Challenge is not a *new* challenge; rather, it is simply a new way of dressing up one of noncognitivism’s more familiar headaches. Addressing these worries will put us in a better position to both appreciate the novelty of the Epistemic Evaluability Challenge and assess what would constitute a satisfactory response.

2.2 Noncognitivist Accounts of Moral Knowledge

Whereas old-school noncognitivists would be happy to chalk our everyday attributions of moral knowledge up to speaker error, contemporary noncognitivists tend to be less dismissive. Indeed, contemporary noncognitivists have gone some lengths towards explaining how desire-like states can satisfy the various conditions on knowledge.

In order to explain the truth condition on knowledge, noncognitivists typically maintain that ‘*p* is true’ is equivalent to ‘*p*’, and so someone who asserts the former expresses whatever mental state would be expressed by asserting the latter.¹⁵ And so when I say, ‘Jane knows stealing is wrong’, I not only report that Jane desires (or prefers, etc.) that people avoid stealing; I also express that I share her desire. In a similar vein, both Blackburn (1996, 1998) and Gibbard (2003) develop noncognitivist glosses on a ‘No Defeaters’ condition. For example, Blackburn proposes that when I say that Jane’s belief is immune

¹⁴Recently, Chrisman (2008, 2015) has argued that traditional versions of expressivism should be replaced with *inferentialism*—the view that the meanings of sentences should be explained in terms of their inferential roles. On this view, the difference between prosaic and moral sentences is that the former typically license various prosaic conclusions, whereas the latter typically license various actions. There is an interesting question as to how the inferential role of a *sentence* bears on the functional role of the underlying belief. I will not tackle this question here. The important point is that insofar as inferentialists reject the idea that moral beliefs have a desire-like functional role, they can also sidestep the Epistemic Evaluability Challenge.

¹⁵See Blackburn 1984, 1993, 1998; Gibbard 2003.

to defeat, I am saying that no improvement in her epistemic position would undermine her belief. Here the notions of ‘improvement’ and ‘undermining’ are themselves normative, and given an expressivist treatment.

While there are difficult questions about points of detail, I am happy to grant that a strategy along these lines can succeed: noncognitivists may well be able to explain how a moral belief can satisfy the various conditions on knowledge. But will doing so provide a solution to the Epistemic Evaluability Challenge? I think not. After all, we still won’t have explained why prototypical desires—e.g., my desire to slake my thirst—cannot similarly satisfy these conditions. If some desire-like states can satisfy the conditions on knowledge, why can’t they all do so?

To clarify: at this point, I am not claiming that this question is unanswerable. Indeed, the rest of the paper will be devoted to assessing whether noncognitivists can provide an answer. My point is that the answer does not simply fall out of the noncognitivist accounts of moral knowledge that have been proposed to date. And so even if we are happy to embrace these accounts, the Epistemic Evaluability Challenge remains.

2.3 Why the Epistemic Evaluability Challenge is Distinct

Noncognitivism already faces a long list of challenges; the Epistemic Evaluability Challenge is just the latest addition. But is it even an addition at all? Will a solution to one of the more familiar challenges carry over to solve the Epistemic Evaluability Challenge? In order to see why the Epistemic Evaluability Challenge is importantly distinct, let us briefly compare it to some of the more familiar worries facing noncognitivists.

So far the challenge that’s received the most press is the Frege-Geach problem: the problem of explaining the embedding behavior and semantic properties of moral terms.¹⁶ While discussions of the Frege-Geach problem have tended to focus on negation and conditionals, a complete solution to the Frege-Geach problem should provide an analysis of *any* sentence that embeds moral vocabulary. This would include moral belief reports, as well as sentences that epistemically evaluate moral beliefs, for example:

- (1) Jane $\left\{ \begin{array}{c} \text{knows} \\ \text{rationally believes} \end{array} \right\}$ that stealing is wrong.

It might be thought that an analysis of such sentences will *ipso facto* solve the Epistemic Evaluability Challenge.

However, I think this is too quick. True, a complete solution to the Frege-Geach problem had better provide an analysis of sentences like (1). And in order to be recognizably noncognitivist, this analysis had better entail that a speaker who utters (1) is epistemically evaluating one of Jane’s desire-like states. But even if such an analysis succeeds,

¹⁶For a classic statement of the problem, see Geach 1964; for an overview, see Schroeder 2008b. Strictly speaking, the Frege-Geach Problem is a problem for expressivism rather than noncognitivism. However, given the close connection between the two, I will ignore this difference.

it won't necessarily tell us why we can epistemically evaluate this particular desire-like state, but not her desire to quench her thirst.

Some may think that this reply overlooks an important facet of the Frege-Geach problem. A full solution to the Frege-Geach problem requires explaining the semantic relations between moral sentences, for example, why 'Stealing is wrong' is inconsistent with 'Stealing is not wrong.' The standard expressivist strategy is to appeal to a notion of 'disagreement in attitude': a disagreement that arises between conative attitudes rather than prosaic beliefs.¹⁷ However, it is not clear that *all* clashes in conative attitudes amount to a form of disagreement. Perhaps, then, if we can explain why some (but not all) conative attitudes generate disagreement, we can leverage this into an explanation of why some (but not all) conative attitudes are epistemically evaluable.

A full response to this suggestion will be left to §5, where I consider some of the most promising candidates for desire-like states that generate disagreement in attitude—in particular, preferences and plans.¹⁸ There I argue that these states also fail to be epistemically evaluable in the same ways as moral and prosaic beliefs. If this is right, then just because a state generates disagreements does not mean that it is epistemically evaluable. And so even if we can explain inconsistency in terms of disagreement in attitude, we still will not have solved the Epistemic Evaluability Challenge.¹⁹

The Epistemic Evaluability Challenge is also importantly different from Dorr's (2002) objection to noncognitivism. According to Dorr, noncognitivists have trouble explaining why it's rational to form beliefs on the basis of certain arguments that contain both moral and prosaic premises. Imagine someone who goes through the following reasoning:

- (2) If lying is wrong, then the souls of liars will be punished in the afterlife.
- (3) Lying is wrong.
- (4) The souls of liars will be punished in the afterlife.

This person could be perfectly rational in so reasoning. But, Dorr contends, noncognitivists have trouble accounting for this. After all, if noncognitivism is correct, this person

¹⁷For relevant discussion, see Stevenson 1944; Blackburn 1984; Gibbard 2003.

¹⁸Of course, it is controversial whether any satisfactory notion of disagreement in attitude is forthcoming. (See e.g., Schroeder 2008a: chp 3; Beddor forthcoming for worries on this front.) For present purposes, I will set such concerns aside.

¹⁹Two further remarks on the relation between the Frege-Geach problem and the Epistemic Evaluability Challenge are in order. First, I am conceiving of the Frege-Geach problem as a relatively circumscribed problem in semantics: it's the problem of giving a compositional expressivist semantics for moral terms that explains their semantic properties and embedding behavior. Others may conceive of the Frege-Geach problem as a more general philosophical challenge to explain all of the properties of moral thought and talk. Given this more liberal construal, I am happy to view the Epistemic Evaluability Challenge as a special case of the Frege-Geach problem (though, for reasons that will become clear, I take it to be a particularly difficult case). Second, I make no claim that the Epistemic Evaluability Challenge is entirely independent of the Frege-Geach problem, even on its more narrow semantic construal. Indeed, the positive proposal I develop in §7 will build on aspects of Gibbard's approach to Frege-Geach worries. The point is simply that a solution to the latter problem will not automatically translate into a solution to the former.

is forming a belief about the world on the basis of a desire-like attitude, which amounts to a form of wishful thinking.

Dorr's objection is similar to mine in that we both think that noncognitivists have trouble explaining the epistemic statuses of moral beliefs. But Dorr's focus is on whether noncognitivists can vindicate our intuitive epistemic verdicts on specific forms of moral reasoning. By contrast, the Epistemic Evaluability Challenge focuses on whether noncognitivists can explain why moral beliefs are apt for epistemic evaluation in the first place. These are very different questions. This is not to say that there is *no* relation whatsoever between them. Perhaps thinking carefully about the conditions under which a desire-like attitude is epistemically rational will help us understand why some desire-like attitudes are apt for epistemic appraisal, whereas others are not. However, there is no reason to assume in advance that a solution to Dorr's challenge will generalize smoothly into a solution to the Epistemic Evaluability Challenge.²⁰

Perhaps the Epistemic Evaluability Challenge's closest cousin in the literature is a challenge raised in passing by Schroeder (2010): the 'Multiple Kinds Problem' (96-97). Schroeder observes that moral beliefs and prosaic beliefs have an awful lot in common. He points out that if noncognitivism is correct, this is a massive coincidence, since, according to noncognitivists, moral and prosaic beliefs are very different attitudes. The Epistemic Evaluability Challenge could be viewed as a special instance of the Multiple Kinds problem. However, it is a particularly thorny instance. Consider some of the other properties shared by moral and prosaic beliefs: they can stand in relations of agreement and disagreement; they can be coherent or incoherent; they are truth apt. It seems a state can possess these properties without being epistemically evaluable. As noted above, arguably preferences and plans stand in disagreement and coherence relations, but they are not epistemically evaluable in the same ways that moral beliefs are. (Or so I'll argue.) Furthermore, suppositions and imaginings are truth apt, but do not seem to be epistemically evaluable. For example, if Jane supposes it's raining in Florence, this supposition is either true or false. But this supposition doesn't seem like it could amount to knowledge. If this is correct, then an explanation of *these* commonalities between prosaic and moral belief will not necessarily yield an explanation of their mutual epistemic evaluability.

Having clarified the Epistemic Evaluability Challenge, let us now consider how noncognitivists might try to answer it.

²⁰To bolster this point, note that many of the proposed solutions to Dorr's challenge do not generalize to solve the Epistemic Evaluability Challenge. For example, Enoch (2003) argues that anyone who has reason to believe (2) also has reason to believe the conditional: *If lying is wrong, I disapprove of lying*, and anyone who has reason to believe (3) also has reason to believe: *I disapprove of lying*. According to Enoch, these alternative descriptive premises are what justifies the agent in believing (4). Note that this proposal makes no headway on the Epistemic Evaluability Challenge; it doesn't explain why the state of believing (3) is itself epistemically evaluable, whereas other desires are not.

3 Are Some Desires Epistemically Evaluable?

One line of response is to insist that some desires are epistemically evaluable. There are a few different ways one might try to motivate this position. One way is to point out that some desires are based on beliefs, which are in turn epistemically evaluable. Perhaps, some may suggest, a desire can inherit the epistemic status of a belief on which it is based. For example, suppose I desire to drink pond sludge, but only because I irrationally believe that drinking pond sludge will make me a better philosopher. Perhaps my desire for a pond potation is epistemically—and not just practically—irrational.

However, there are two difficulties with this response. First, even if desires can inherit *some* of the epistemic statuses of the beliefs on which they are based, it doesn't seem that they can inherit all of these statuses. Suppose I know that I will enjoy a particular play. As a result, I form a desire to see it. Even if we grant that this desire is epistemically rational, it seems odd to describe the desire as knowledge. And so a restricted version of EEQ remains:

Why can beliefs (both moral and prosaic) qualify as knowledge, whereas desires cannot?

As with EEQ, it is tempting to answer this question by appealing to some difference in the functional role of beliefs and desires. But, as before, any such answer stands in tension with noncognitivism.

Even if we set aside the issue of knowledge, a further difficulty remains. According to the response under consideration, desires are at best *derivatively epistemically evaluable*: whenever a desire has some epistemic status, it derives this status from a belief on which it is based. However, beliefs are importantly different. Take any foundational prosaic belief, for example, *It appears to me that such-and-such*. This belief can be epistemically rational, even though it does not inherit this status from any other beliefs. The same goes for foundational moral beliefs. My belief that *pain is bad* can be epistemically rational, even though it does not seem to derive this status from other beliefs.

This difference gives rise to another variant of EEQ:

Why can beliefs (both moral and prosaic) be *non-derivatively* epistemically rational, whereas desires cannot?

Once again, noncognitivists will have to give an answer that does not depend on functional role.²¹

²¹Some might insist that moral beliefs are only derivatively epistemically rational, and that even my belief that *pain is bad* derives its epistemic status from some prosaic belief of the form, *Pain has properties $P_1 \dots P_n$* . But this just pushes us back a step: what prosaic belief serves as the basis for my belief, *Anything with properties $P_1 \dots P_n$ is bad*? Alternatively, some might concede that the most basic moral beliefs have their epistemic statuses by default: my belief that pain is bad is epistemically rational simply in virtue of being foundational. But then why wouldn't the same hold for basic desires—say, the desire to avoid pain?

Another way of motivating the idea that some desires are epistemically evaluable is to point out that noncognitivists already face the ‘Moral Attitude Problem.’ This problem stems from the fact that not every desire constitutes a moral belief, even by the noncognitivist’s lights. After all, I might desire to slake my thirst without thinking that thirst-slaking is morally good. Thus noncognitivists owe us some story about what distinguishes the ‘moral attitude’ from ordinary, non-moral desires.²² One might hope that a solution to this problem will also yield an answer to the Epistemic Evaluability Challenge: perhaps once we know what distinguishes the moral attitude, we’ll also know why this attitude is epistemically evaluable.

However, this response also faces two difficulties.²³ The first concerns its plausibility. On the face of it, the question, *What makes an attitude qualify as moral?* seems very different from the question, *What makes an attitude epistemically evaluable?* After all, morality is one thing, epistemology is another. Thus noncognitivists who opt for this response need to provide some reason for thinking the two questions are connected.

If we turn to look at some of the leading solutions to the Moral Attitude Problem, we find scant grounds for optimism. Consider for example, the ‘emotional ascent’ solution advocated by Blackburn (1998). Simplifying somewhat, Blackburn proposes that a desire counts as moral as long as it’s accompanied by a higher-order desire for other people to share it. For example, my desire that no one steals constitutes a moral desire because I also desire that others share this desire. This proposal has the advantage of explaining why we care whether others share the moral attitude. And perhaps it thereby explains why we make claims about the conditions under which the moral attitude is warranted or unwarranted. But what explains why we make claims about whether the moral attitude is *epistemically* warranted, rather than practically warranted? Why would meeting this higher-order desire condition suffice to make a desire *epistemically* rational or irrational, let alone a candidate for knowledge?

A second difficulty comes from parsimony considerations. Even if being a moral attitude is sufficient for epistemically evaluability, it is clearly not necessary. After all, prosaic beliefs are epistemically evaluable. What explains this? If noncognitivists allow that prosaic beliefs are epistemically evaluable in virtue of some features of their functional role, they will be saddled with a disjunctive answer to EEQ. Specifically, they will be committed to saying that prosaic beliefs are epistemically evaluable in virtue of some features of their functional role, but moral beliefs are epistemically evaluable in virtue of meeting some other condition—namely, whatever condition makes a desire qualify as distinctively moral. By contrast, cognitivists have a more parsimonious answer to EEQ:

²²The label, ‘Moral Attitude Problem’ is due to Miller (2003), who raises difficulties for a number of proposed solutions. For discussion of how noncognitivists might solve the problem, see Kauppinen 2010; Köhler 2013; Björnsson and McPherson 2014.

²³A caveat: I am here focusing on noncognitivists who think the moral attitude is a specific type of desire. My criticisms do not apply to those who want to solve the Moral Attitude Problem by taking the moral attitude to have a more complicated functional role—one that differs from the D-Role in certain respects. I discuss such proposals in §5.

they'll say that all beliefs—moral and prosaic—are epistemically evaluable in virtue of some features of their functional role.

A third and final way of motivating the claim that some desires are epistemically evaluable is worth discussing. Perhaps, some may suggest, certain desires are epistemically evaluable, but not *qua* desires. Rather, we need to conceive of them under the 'guise of belief' in order for them to be epistemically evaluable. According to this proposal, a state's epistemic evaluability (or lack thereof) is relative to our mode of describing the state. Call it a desire, and one cannot properly ascribe it any epistemic status. Call it a belief, and suddenly one can.²⁴

This is an intriguing option; in my view, it is the most promising way of motivating the idea that desires are sometimes epistemically evaluable. That said, I think it faces a significant obstacle. Proponents of this 'guise of belief' response will need to maintain that epistemic evaluations create opaque contexts, and it is questionable whether this claim can be independently motivated. To see why they need to make this claim, consider the following argument:

- (5) Jane's belief (that stealing is wrong) is epistemically justified.
- (6) Jane's desire (that no one steals) is not epistemically justified.
- (7) Jane's belief \neq Jane's desire.

In order to resist the conclusion, proponents of the 'guise of belief' response will need to say that one of the constituent expressions—'Jane's belief', 'Jane's desire', 'epistemically justified'—shifts its extension in the course of the argument. But it is one thing to say this, and another thing to motivate it. Is there independent reason to posit this sort of semantic shift?

To appreciate the difficulty, consider an analogous argument that often crops up in debates over whether constitution is identity:

- (8) The statue is necessarily a statue.
- (9) The clay is not necessarily a statue.
- (10) The statue \neq the clay.

While many find this argument persuasive, those who hold that constitution is identity cry foul; they insist that one of the premises creates an opaque context, and hence the conclusion does not follow.²⁵ Regardless of whether we find this response persuasive, we should at least admit that its semantic premise is on comparatively firm ground. After all, it is independently plausible that modal terms create opaque contexts. Take an astronomically updated version of Quine's (1961) famous example: eight is necessarily greater than seven; the number of planets is not necessarily greater than seven; but from this we cannot conclude that there are not eight planets. This is precisely the sort of

²⁴Thanks to an anonymous reviewer for suggesting this response.

²⁵See e.g., Gibbard 1975. For criticisms of this response, see e.g., Baker 1997.

independent motivation that proponents of the ‘guise of belief’ response need, but seem to lack.

None of this is intended to be a decisive refutation of the ‘guise of belief’ response. My point is simply that proponents of this response incur substantive and controversial semantic commitments. If—as I’ll be arguing in §7—there is a way of solving the Epistemic Evaluability Challenge that avoids these commitments, then this provides a reason to favor such an alternative.

4 Noncognitivism Extended

Another response to the Epistemic Evaluability Challenge is to switch tack and ask, ‘What’s involved in epistemically evaluating a psychological state in the first place?’

Epistemic evaluations resemble moral evaluations in some important respects. In particular, both are normative: calling a belief ‘irrational’ is a form of criticism; calling a belief ‘rational’ is a form of praise. This has led some moral noncognitivists to embrace *epistemic noncognitivism*. According to epistemic noncognitivism, epistemically evaluating a psychological state *s* involves adopting a particular desire-like attitude—an ‘epistemic attitude’—towards *s*.²⁶ To illustrate, suppose Susie believes that global warming has human causes, and Fred deems this belief epistemically rational. On a simple version of epistemic noncognitivism, for Fred to believe that Susie’s belief is rational is for Fred to approve—perhaps in some distinctly epistemic way—of Susie’s belief. Given the assumption that epistemic rationality is a necessary condition on knowledge, this entails that knowledge attributions also have a noncognitive element: part of what it is for Fred to believe that Susie *knows* that global warming has human causes is for Fred to adopt some epistemic attitude towards her belief.

Epistemic noncognitivism is controversial, but let us suppose for the sake of argument that some version of it proves viable. For epistemic noncognitivists, answering EEQ will amount to answering what I’ll call the ‘Epistemic Attitude Question’ (EAQ):

EAQ Why do we adopt epistemic attitudes towards beliefs, but not towards desires?

However, it’s doubtful that EAQ is any more tractable than EEQ. In order to answer EAQ, noncognitivists will still need to identify some features that moral and prosaic beliefs possess, and that prototypical desires lack—features that explain why we adopt epistemic attitudes towards the former but not the latter. But what are these features, if not aspects of the functional role of moral and prosaic beliefs?

Perhaps, some may suggest, this just shows that we should take our noncognitivism a step further. In addition to being noncognitivists about epistemic evaluations, perhaps we should also be noncognitivists about belief ascriptions: perhaps regarding some

²⁶For sympathetic discussions of epistemic noncognitivism, see Gibbard 2003; Ridge 2007a; Field 2009, 2018; Kappel 2010; Chrisman 2007; Carter and Chrisman 2012; Grajner 2015; Beddor 2016: chp.2, among others.

psychological state s as a ‘belief’ involves adopting a conative attitude towards s . For example, one might hold that part of what it is to regard s as a ‘belief’ is to be willing to adopt epistemic attitudes towards s . At first blush, such a view appears to offer an easy way out of the challenge. If regarding a state as a ‘belief’ requires being willing to epistemically evaluate it, then there is no great mystery as to why beliefs—both moral and prosaic—are epistemically evaluable.

But does this move really help? We still don’t have an answer to EAQ: we still don’t have a story about why we’re willing to adopt epistemic attitudes towards some states and not others. Presumably, this isn’t a matter of whim or convention. (If it were, we would expect to find communities that adopt the epistemic attitudes towards all sorts of psychological states: desires for thirst-quenching, headaches, you name it.) And so noncognitivists about belief attributions still need to identify some further features of the psychological states we call ‘beliefs’, features that explain why we’re willing to adopt the epistemic attitudes towards them in the first place. Again, it is unclear what those features could be, if not some aspects of their functional role.²⁷

5 Enriching the Boxology

Another response to the Epistemic Evaluability Challenge is to insist on a richer boxology. Our cognitive architecture involves more than two boxes: mental states such as plans, intentions, and preferences have functional roles that resemble the D-Role in some respects, but differ from the D-Role in others. And some noncognitivists have been careful to identify moral beliefs with these *desire-like* states. For example, Gibbard (2003, 2013) takes moral belief to be a type of plan; Dreier (2006, 2009) and Silk (2015) take moral belief to be a type of preference. Perhaps noncognitivists who go this route can retain the Functional Role Hypothesis: perhaps they can use the functional differences between moral beliefs and desires to explain why the former are epistemically evaluable in ways that the latter are not.

5.1 Preferences and Plans

But can we find any *desire-like* states that are epistemically evaluable in the same ways as beliefs? Suppose that Jane prefers attending a party to staying home. Suppose that Jane also plans to attend the party. Would we be willing to describe either her preference or her plan as *epistemically* rational? Perhaps we would if both are based on a belief

²⁷This is not to say that a state’s functional role completely determines whether we classify it as a belief. For example, one might adopt the view defended in Shah and Velleman 2005, according to which there are a variety of attitudes of acceptance, all of which have similar functional roles. Shah and Velleman maintain that treating one of these attitudes as a ‘belief’ involves adopting a conative attitude towards it. Note that on this view the functional role of a state does not fully determine whether it is classified as a belief. However, it does impose significant constraints, precluding certain states (in particular, conative attitudes) from being the kind of things that we could properly regard as beliefs.

that possesses this epistemic status—for example, a belief that her friends will attend the party. But this at best gets us derivative epistemic rationality. Moreover, it seems odd to describe either the preference or the plan as knowledge, even derivatively.

Thus neither preferences nor plans are epistemically evaluable in the same ways as moral beliefs. This is instructive, since one motivation for identifying moral beliefs with such states is that they display some of the hallmarks of prosaic beliefs. For example, Dreier’s reason for thinking that noncognitivists should identify moral beliefs with preferences is that preferences are subject to coherence constraints: in this regard they resemble prosaic beliefs and differ from desires. According to Dreier, if I prefer to take a trip and also prefer to stay home, my preferences are incoherent; if I desire to take a trip and also desire to stay home, my desires are only conflicted (2009: 105-106). Similarly, one of Gibbard’s reasons for identifying moral beliefs with plans is that plans, like prosaic beliefs, stand in agreement and disagreement relations. To illustrate with one of Gibbard’s examples, suppose Caesar plans to go to the Senate. Suppose Brutus has a conflicting ‘contingency plan’: his plan, conditional on being in Caesar’s situation, is to stay home. According to Gibbard, Caesar and Brutus disagree in plan (2003: 68-69).

It is not entirely clear to me that desires differ from preferences and plans in these respects. But let us grant, at least for the sake of argument, that they do: let us grant that preferences can be evaluated for coherence/incoherence, and that plans can be evaluated for agreement/disagreement. What my discussion shows is that whatever functional features render a state evaluable in *these* ways do not thereby render it *epistemically* evaluable.

5.2 Beyond Folk Psychology

Suppose noncognitivists concede that it was a mistake to identify moral beliefs with preferences or plans. ‘Still,’ they may insist, ‘this doesn’t show that there couldn’t be a desire-like state with a slightly different functional role—call it the “D*-Role”—that renders this state epistemically evaluable in the same ways as prosaic beliefs. Admittedly, it is difficult to find any clear examples of such states in our folk psychological inventory. But so what? There are more states than are dreamt of in folk psychology.’

One worry about this response is that we risk losing our grip on noncognitivism. Long before taking our first metaethics class, we had a pre-theoretical understanding of various conative attitudes (desire, preference, and the like), an understanding that came with our mastery of folk psychology. Then we started doing metaethics and met the noncognitivist, who identified moral beliefs with these conative attitudes. Regardless of whether we agreed with this proposal, we at least understood it. But now the noncognitivist has changed her tune. Now moral belief is claimed to be an unfamiliar attitude, unknown to folk psychology. This raises the worry that we lack any clear understanding of the attitude in question.

In order to address this concern, noncognitivists who look beyond folk psychology owe us an account of the states they posit. They need to give us some sense of how the

D*-Role differs from the D-Role, and why this difference makes an epistemic difference. Without such an account, noncognitivists have traded boxology for black box-ology.²⁸

Providing such an account is easier said than done. The natural temptation is to look to the B-Role for guidance: to find some feature of the B-Role that plausibly underpins epistemic evaluability, and insist that part of what it is for a mental state to play the D*-Role is for it to possess this feature. But this strategy comes with a risk. The more that the functional role of moral belief resembles the functional role of prosaic belief, the harder it is to maintain that the functional role of moral belief resembles the D-Role much more closely than it does the B-Role.

Thus noncognitivists who pursue this strategy must walk a fine line. On the one hand, they must make the functional role of moral belief close enough to the B-Role to secure epistemic evaluability. On the other hand, they must make the functional role of moral belief distant enough from the B-Role to preserve their noncognitivist *bona fides*. In the rest of this section, I consider what strikes me as the most promising way of treading this tightrope.

5.3 Evidence Responsiveness

Prototypical beliefs are responsive to evidence. Suppose I initially believe my car is in the shop, but I later gain compelling evidence that it isn't there. I'll be disposed to revise my belief. Some have suggested that this evidence responsiveness is part of the functional role of belief:

Ceteris paribus, gaining sufficiently strong evidence for (against) p disposes one to believe (cease believing) p .²⁹

Arguably, evidence responsiveness is a feature of moral beliefs as well as prosaic beliefs. Suppose Jane initially believes that the US should impose sanctions on a particular country, but later receives evidence that the US shouldn't do so. If Alex accepts this evidence, we would expect her to be at least somewhat disposed to revise her belief.

By contrast, prototypical desires and preferences don't seem to be evidence responsive (Smith 1987). Of course, they are *reasons responsive*: gaining a sufficiently strong practical reason not to attend a performance (e.g., reading a bad review) will lead me

²⁸Recently, Gibbard has acknowledged that his notion of a plan differs in certain ways from our ordinary concept of a plan (2013: chp.8). However, none of the differences that Gibbard mentions seem to make for an epistemic difference. For example, Gibbardian plans, unlike ordinary plans, apply to hypothetical situations in which we will never find ourselves (e.g., being in Caesar's sandals, contemplating whether to go to the Senate). This difference in the scope of the plans doesn't seem epistemically relevant. From the epistemic point of view, a Gibbardian plan to go to the Senate, if in Caesar's position, seems on a par with an ordinary plan to stop at the store: both are at best derivatively epistemically evaluable.

²⁹For sympathetic discussion of the idea that evidence responsiveness is partially constitutive of belief, see Adler 2002; Velleman 2000; Shah and Velleman 2005. For dissent, see Bayne and Pacherie 2005.

to abandon my desire to attend. But gaining evidence that I will not attend the performance will not have the same effect. Indeed, learning that tickets are extremely scarce (and hence that I am unlikely to secure one) may increase my desire to attend.

Are intentions and plans evidence responsive? This is less clear. According to what's sometimes called the 'Strong Belief Thesis', intending to ϕ entails believing one will ϕ . If this is right, then gaining strong evidence that I won't see the play will dispose me to cease intending to see it. However, the Strong Belief Thesis is controversial.³⁰ And even if it's correct, intentions will be at most derivatively evidence responsive: they'll only be evidence responsive because they entail some further evidence responsive state (belief).

Perhaps, then, evidence responsiveness (or non-derivative evidence responsiveness) gives us what we want. Arguably, it's part of the functional role of both moral and prosaic belief. Moreover, it's the sort of feature that might be thought to suffice for epistemic evaluability. At the same time, evidence responsiveness doesn't seem to be sufficient for belief. And so noncognitivists could maintain that while moral beliefs are evidence responsive, they play a desire-like role in practical reasoning. Hence they differ from prosaic beliefs in their output conditions.

While this all seems promising, a serious hurdle remains. According to the present proposal, evidence responsiveness is part of the functional role of moral beliefs. But a state is evidence responsive just in case it is responsive to evidence for its *content*. In the case of moral belief, this will be a moral content. This raises a worry: if noncognitivists stop here, they will be helping themselves to the notion of a moral content, which runs contrary to the explanatory ambitions of noncognitivism. After all, one of the main goals of noncognitivism is to *explain* the contents of moral thought in purely naturalistic terms—to locate 'ought's within the world of 'is's'.³¹

To overcome this hurdle, advocates of the evidence responsiveness strategy will need to develop some independent account of what it is to have evidence for or against a moral content—an account that, on pain of circularity, does not itself rely on the notion of moral belief. Can this be done? Perhaps. But we should not underestimate the difficulties. To begin with, note that while noncognitivists have developed accounts of moral knowledge (§2.2) and epistemic evaluations more generally (§4), none have—at least to my knowledge—provided what's needed here, which is an account of how evidence can bear on some moral content. Moreover, when we consider how noncognitivists might fill in the details, we find that some initially attractive routes lead to dead ends.

For example, one natural strategy would be to proceed as follows. Start with your preferred account of evidence. Next, analyze claims about 'evidential support' as claims about what credences an agent should adopt, given some body of evidence. For example, 'Evidence e counts in favor of p ' could be paraphrased as 'Someone who gains e should raise their degree of belief in p .' We could go on to give this 'should' an expressivist gloss, analyzing such claims as expressions of an epistemic pro-attitude towards someone rais-

³⁰For discussion, see Bratman 1987; Ross 2009.

³¹See Blackburn 1998 and Gibbard 1990, 2003 for clear statements of the naturalistic agenda.

ing their degree of belief in p upon gaining e . While this analysis is promising, it won't give proponents of the evidence responsiveness strategy what they need. After all, this analysis relies on the notion of *degrees of belief*. When p is moral, this will be a degree of belief in a moral content. And so the account relies on the very sort of psychological relation towards a moral content that we sought to explain.³²

Let's take stock. While it's natural to try to resolve the Epistemic Evaluability Challenge by insisting on a richer functional role for moral beliefs, this response runs into a number of hurdles. It won't help to model moral beliefs on preferences or plans, since these states are not epistemically evaluable in the same ways as moral beliefs. And if noncognitivists choose to identify moral beliefs with some other state, they owe us an account of this state's functional role—an account that explains why this state is epistemically evaluable. Providing such an account is not easy. One natural strategy—appealing to evidence responsiveness—led the noncognitivist to incur further explanatory commitments. It remains to be seen whether these commitments can be discharged and, if not, whether some alternative strategy fares better.

6 Hybrid Approaches

Thus far, we have focused on 'pure' noncognitivist views, according to which having a moral belief is just a matter of having some conative attitude. Recently, a number of authors have proposed hybrid views that integrate cognitivist and noncognitivist elements. According to hybrid theorists, having a moral belief involves having both a conative attitude and a prosaic belief.³³ For example, a simple hybrid view might hold that believing torture is wrong involves being in two states: (i) a prosaic belief that torture has some natural property F (say, causing suffering), (ii) a desire to avoid actions that have F . If noncognitivists are willing to go hybrid, can they escape the Epistemic Evaluability Challenge?

At first blush, the answer appears to be 'Yes.' After all, hybrid theorists can adopt a

³²A different strategy would be to suggest that a moral belief can be evidence responsive without being responsive to evidence for/against its content. All that's required is for it to be responsive to evidence for/against some prosaic belief on which the moral belief is partially based. For example: Jane believes that eating meat at a certain restaurant is morally permissible, but only because she believes the restaurant uses humanely treated animals. Were she to gain evidence against the prosaic belief, she would revise her moral belief. But *this* sort of evidence responsiveness is also exhibited by run-of-the-mill plans and preferences. For example, I plan to go to the store, but only because I believe it will be open. Were I to gain evidence the store is closed, I would revise my plan. And so this indirect evidence responsiveness won't give us what we need, which is a special sort of evidence responsiveness that distinguishes moral beliefs from ordinary plans/preferences, in virtue of which the former but not the latter are epistemically evaluable.

³³For discussion of hybrid views, see Ridge 2006b, 2007b, 2014; Boisvert 2008; Schroeder 2009; Fletcher and Ridge 2014; Laskowski forthcoming. As Schroeder (2013) observes, a hybrid view can be thought of as a special case of a 'relational' view, according to which one has a moral belief if and only if a particular relation between prosaic beliefs and desire-like states obtains. (For an example of a non-hybrid relational view, see Toppinen 2013.) While I focus on hybrid views, everything I say applies to relational views in general.

version of the Functional Role Hypothesis: they can insist that a state is epistemically evaluable as long as it plays the B-Role, or is *partially composed of a state that plays the B-Role*. Moral beliefs are epistemically evaluable, since they contain a prosaic component. Desires are not epistemically evaluable, since they contain no such component.

Alas, things are not so simple. Consider: how exactly does the prosaic component render the moral belief epistemically evaluable? The simplest answer would be that the moral belief inherits the epistemic status of its prosaic component. But if we consider the details of standard hybrid views, we see that this can't be right. Suppose that vicious Vic values suffering. According to the simple hybrid view, Vic's belief that torture is good consists in (i) a prosaic belief that torture causes suffering, (ii) a desire to perform actions that cause suffering. Now, suppose Vic's prosaic belief (that torture causes suffering) is in perfectly good epistemic standing: it is rational and justified; it qualifies as knowledge. Still, we presumably do not want to accord his belief that torture is good the same epistemic standing. Most obviously, we would be loath to call it knowledge (since knowledge is factive); we would also be reluctant to deem it rational or justified.³⁴

What this shows is that the epistemic status of a moral belief is not simply inherited from its prosaic component. Instead, it is somehow affected by the conative component. Hybrid theorists owe us an account of how this works: how exactly does Vic's desire preclude the moral belief from qualifying as knowledge, or as epistemic rational? More generally, how does a desire which is not itself epistemically evaluable affect the epistemic status of a composite belief? These questions are far from trivial, and it is by no means obvious how to provide satisfactory answers. Until such answers are forthcoming, hybrid theorists cannot claim to have fully resolved our challenge.³⁵

7 Grounding Noncognitivism

We have explored a number of initially attractive options for resolving the Epistemic Evaluability Challenge, only to run into further obstacles. Perhaps, then, it is time to go back to the beginning and reconsider our initial construal of 'pure' noncognitivism. In this section, I suggest construing noncognitivism as a thesis about the *grounding* re-

³⁴This problem is not an artifact of the simple hybrid view I've chosen for the sake of illustration; it also affects other hybrid views. Take, for example, the view in Ridge 2006b. On this view, believing that x is good involves both (i) a prosaic belief that some advisor would approve of x , (ii) approval of the advisor mentioned in (i). Suppose Vic's belief that some advisor would approve of torture has impeccable epistemic credentials. Still, it seems that Vic's belief that *torture is good* is epistemically defective.

³⁵To be clear on the difficulty facing hybrid theorists: I am not claiming that there is anything wrong with holding that states which are not epistemically evaluable influence the epistemic status of other states. After all, experiences seem to work this way; my perceptual experience of a sunset is not epistemically evaluable, but it can affect the epistemic status of my belief that I'm seeing a sunset. The difficulty is rather in providing a general explanation of how desires wield their epistemic influence. In the case of experiences, there is arguably a simple story to be told. Any experience with content p makes it *prima facie* rational to believe p (Pryor 2002). Clearly no story along these lines will work in the case of desire. Desiring p certainly does not make it rational to believe p —otherwise wishful thinking would be rational!

lations between desires and beliefs. While this version of noncognitivism has not been discussed in the literature, I argue that it has significant advantages. In particular, it is fully consistent Functional Role Hypothesis. As a result, it offers a promising way to resolve our challenge.

7.1 Introducing Grounding Noncognitivism

We began by construing pure noncognitivism as a disjunctive view of belief: to have a belief is to either be in a state that plays the B-Role or to be in a state that plays the D-Role. Thus understood, noncognitivism amounts to a ‘Two Box’ conception of belief. It says there are two kinds of belief, prosaic and moral, each with its own functional role. While this is a natural construal, it is not the only option. Here is an alternative. Suppose we agree with the cognitivist that there is only one belief box: prosaic and moral beliefs have the same functional role. Where we part ways with the cognitivist is in our conception of the relation between the B-Role and the D-Role. In particular, we advance the following thesis: whenever an agent has a moral belief, this is grounded in the fact that they have certain desire-like attitudes. For example, whenever an agent believes that stealing is wrong, this fact is grounded in the further fact that they have some distinctly moral species of disapproval towards stealing. Call this view, ‘Grounding Noncognitivism.’ (See Fig. 2.)

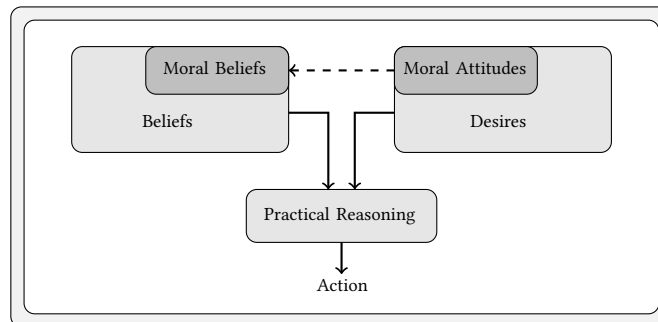


Figure 2: Grounding Noncognitivism. Solid lines represent causal dependence; dashed lines represent metaphysical dependence.

Grounding noncognitivism is an unfamiliar view; a few words of clarification may help bring it in clearer perspective. First, what notion of ‘grounding’ is at issue? Here I am drawing on the metaphysics literature, which takes grounding to be a non-causal dependence relation, of the sort that is often conveyed by the expressions, ‘because’ and ‘in virtue of.’ By way of illustration, consider the following claims: (i) the vase is fragile in virtue of its categorical properties; (ii) Jones is in pain because his nociceptors are firing, (iii) the set {Socrates} exists in virtue of the fact that Socrates exists. Each of these claims postulates a non-causal dependence relation between one fact and another.

‘Grounding’ is a term of art used to denote this relation.³⁶

By now, much ink has been spilled on the grounding relation. I will set subtle questions about the nature of grounding aside, since they are largely orthogonal to my purposes. What I mainly need is a fairly minimal assumption: the grounding relation is distinct from the identity relation. This seems plausible in light of the preceding examples: dispositional properties are not identical to their categorical bases; Socrates’ singleton set is not identical to Socrates.³⁷ Given this assumption, we are in a position to see how grounding noncognitivism differs from the more familiar view, embraced by ‘Two Box’ noncognitivism, that moral beliefs are identical to conative states. Both views agree that there is an intimate modal connection between moral belief and desire, specifically:

Modal Connection Necessarily, S has a moral belief if and only if S is in a certain conative state (specifically, the moral attitude).³⁸

But they differ on the exact relation between the moral belief and the desire. For the grounding noncognitivist, the relation between the belief and the desire is the same as the relation between {Socrates} and Socrates—one of dependence rather than identity.

How does grounding noncognitivism help with the Epistemic Evaluability Challenge? The answer is that grounding noncognitivists can embrace the Functional Role Hypothesis without revisions or reservations. After all, grounding noncognitivists maintain that prosaic and moral beliefs have the same functional role. And so grounding noncognitivists can happily allow that this role suffices to make a state epistemically evaluable. Similarly, grounding noncognitivists will maintain that desires have an importantly different functional role—a role that deprives them of epistemic evaluability. To illustrate, suppose that Jane has a distinctively moral desire for people to give to charity. According to grounding noncognitivists, this desire is not epistemically evaluable, since it lacks the epistemic evaluability-conferring features of the B-Role. However, the fact that Jane has this desire grounds the fact that she has a certain moral belief—specifically, the belief that giving to charity is good. And this belief is epistemically evaluable, in virtue of its functional role.

Some might find this rather mysterious. If the desire-like attitude grounds the moral belief, how can the belief possess a property—e.g., the property of being knowledge—that the desire lacks? But the literature on grounding affords plenty of cases where the existence of x grounds the existence of y , yet y possesses properties that x lacks. Take

³⁶For seminal work on grounding, see Fine (2001, 2012); Schaffer (2009); Rosen (2010). As these authors note, while the term ‘grounding’ is recent, the notion is old. We find it at work as far back as the *Euthyphro*, when Socrates asks whether something is pious *because* it is loved by the gods.

³⁷Indeed, many hold that grounding is an asymmetric, hence irreflexive, relation. The main motivation for this is that grounding is an *explanatory* relation, and explanation seems to be asymmetric (Cameron 2008; Schaffer 2009; Rosen 2010).

³⁸At least, the grounding noncognitivist agrees on this point given the plausible assumption that grounding is a species of metaphysical entailment: if A fully grounds B at a world w , then in any world v where A obtains, B obtains. For defense of this assumption, see Fine 2012; Trogon 2013.

our earlier example: the fact that Jones is in a certain brain state grounds the fact that Jones is in pain, but the pain has the property of being mental, whereas the brain state does not. Or take the example of Socrates and his singleton set. Socrates' existence grounds the existence of his singleton set, but his singleton set has the property of being a set, whereas Socrates does not.

Grounding noncognitivism thus offers a promising way of escaping the Epistemic Evaluability Challenge. That said, the view gives rise to difficult questions. In what follows, I canvass some of the most pressing. §7.2 tackles an explanatory question: *how* do desire-like states manage to ground moral beliefs? §7.3 explores whether grounding noncognitivism merits the noncognitivist title: how does grounding noncognitivism differ from sophisticated versions of subjectivism, such as Dreier's (1990) view? Finally, §7.4 considers whether grounding noncognitivism can explain why only some desire-like states ground beliefs.

7.2 How Do Desire-Like States Ground Moral Beliefs?

As it stands, grounding noncognitivism may appear to have all the advantages of explanatory theft over toil. How exactly does possession of a conative state manage to ground possession of a moral belief?

Some might resist this demand for explanation. It is not clear whether all grounding relations can be explained; perhaps some grounding facts are brute. However, this response does little to satisfy the craving for a more genuinely explanatory account. Moreover, it is a risky strategy dialectically. After all, noncognitivists often appeal to the Humean theory of motivation. And people are usually attracted to the Humean theory because they want to avoid positing brute necessary connections between beliefs and desires. But surely positing brute grounding connections is no better, especially since the grounding connections entail necessary connections!

It is thus incumbent on grounding noncognitivists to provide some further explanation here. In what follows, I'll sketch one way of delivering the goods. My explanation proceeds in two parts. The first part is semantic: I sketch an expressivist account of moral contents. The second part is psychological: I develop a general account of the functional roles of belief and desire, from which a noncognitivist treatment of the functional role of moral belief emerges as a special case. Combining these two parts delivers an explanation of why moral beliefs necessarily depend on conative states—an explanation that preserves a unitary functional role for all belief, moral and prosaic. Note that this two-part explanation is only intended as a 'proof of concept': I aim to sketch one possible way of developing grounding noncognitivism into a genuinely explanatory theory.

7.2.1 The Semantic Stage: Contents for Expressivists

Developing the first stage of my explanation requires shifting from moral psychology to moral semantics. The natural semantic ally for a noncognitivist psychology is expres-

sivism. According to expressivism, the function of a moral utterance is not to represent the world, but rather to express the speaker's conative attitudes.

By now, a variety of sophisticated implementations of this central expressivist idea have been developed. For present purposes, I will focus on one implementation, due to Gibbard (1990, 2003). I choose Gibbard's strategy because it is one of the most prominent and promising expressivist semantics to date. However, I should stress at the outset that nothing crucial hinges on this choice of implementation.

In order to introduce Gibbard's framework, it is useful to start with a possible worlds semantics. Begin by assuming that prosaic claims represent the world as being a certain way, and hence can be modeled with a set of possible worlds. For example, 'Grass is green' represents the world as being one where grass is green; hence it can be modeled with the set of worlds in which grass is green. Since moral utterances do not represent the world as being a certain way, their contents cannot be modeled by sets of worlds. Nonetheless, we can still model their contents using a conservative extension of a possible worlds semantics. In particular, we can take the content of a moral utterance to be a set of ordered pairs whose first member is a world and whose second member is some formal entity that represents the moral attitude. Different versions of this approach are possible, depending on what one takes the moral attitude to be. Gibbard (1990) uses a norm (representing the content of a state of norm acceptance); Gibbard (2003) and Yalcin (2012) use a hyperplan (representing the content of a planning state). In order to remain as neutral as possible concerning the nature of the moral attitude, I will take the entity to be a *moral perspective*, where a moral perspective is a representation of any actual or hypothetical moral attitude.³⁹ Nothing substantive will be assumed about the moral attitude, only that it is a conative state, and that it can be used to rank various actions and outcomes.

To illustrate this approach, take a moral utterance such as:

- (11) Charitable giving is good.

The content of this utterance will be a set of world, moral perspective pairs—specifically, the set of w, m pairs where m assigns high marks to charitable giving at w . For ease of reference, let us call this set 'GOOD':

$$\text{GOOD} = \{ \langle w, m \rangle \mid m \text{ highly ranks charitable giving at } w \}.$$

Two advantages of Gibbard's framework are worth highlighting. First, it enables expressivists to mimic the standard possible worlds semantics for negation, disjunction, conjunction, and so on. As in possible worlds semantics, these operations are analyzed in terms of set theoretic operations—complementation, union, and intersection. The only difference is that these operations are now defined over sets of world, moral perspective pairs. Second, the Gibbardian framework allows us to assign the same sort of content

³⁹ Cf. the notion of a 'normative perspective' in Ridge 2014: 113-121.

to both moral and prosaic sentences. To see this, note that for any set of worlds Γ we can define a ‘boring’ set of world, moral perspective pairs that is equivalent to Γ .⁴⁰ Thus rather than taking the semantic content of ‘Grass is green’ to be the set of worlds in which grass is green, we can take it to be the set of world, moral perspective pairs in which grass is green. As a result, Gibbard’s framework extends naturally to sentences that ‘mix’ moral and non-moral content—e.g., ‘Grass is green and charitable giving is good.’ Taken together, these features make the Gibbardian framework one of the most promising strategies for handling the Frege-Geach problem.

Equipped with an expressivist account of content, we can extract expressivist truth conditions. Of course, old-school noncognitivism shunned all talk of moral truth. But just as most contemporary noncognitivists want to preserve our everyday attributions of moral belief, most want to preserve our everyday attributions of moral truth—assuming these attributions can be analyzed in a way that is compatible with antirealism.⁴¹

Here’s a natural way of getting the desired analysis out of our Gibbardian contents. On a standard intensional semantics in the tradition of Kaplan (1989), contents have truth-values relative to a *circumstance of evaluation*. (Think of a circumstance of evaluation as a tuple of whatever features need to be settled in order for a content to sensibly be described as true or false.) Expressivists can help themselves themselves to this framework provided they take circumstances of evaluation to include not just a world, but also a moral perspective. They can then adopt a standard definition of truth at a circumstance of evaluation: a content p is true at a circumstance of evaluation $\langle w, m \rangle$ if and only if $\langle w, m \rangle \in p$. To illustrate, this approach yields the following truth conditions for (11):

Expressivist Truth Conditions ‘Charitable giving is good’ is true at $\langle w, m \rangle$ if and only if $\langle w, m \rangle \in \text{GOOD}$, which obtains if and only if m highly ranks charitable giving at w .

Some may worry that even if expressivists are willing to talk about moral truth, the notion of moral *truth conditions* is anathema to expressivism. However, I think that there is no need for expressivists to renounce truth conditional semantics altogether. Rather they should distinguish between *representational truth conditions* and *non-representational truth conditions*. Representational truth conditions place conditions on the world and the world alone: if a sentence has representational truth conditions, then in order to settle its truth-value all one needs to do is settle what the world is like.⁴² Our expressivist truth conditions are non-representational. To find out whether the content of (11) is true, it’s not enough to settle on a world. One also needs to settle on a moral perspective. By distinguishing between representational truth conditions and

⁴⁰More precisely, say a set of world, moral perspective pairs p is boring just in case for any moral perspectives m_1, m_2 and any world w , $\langle w, m_1 \rangle \in p$ if and only if $\langle w, m_2 \rangle \in p$. (Cf. the notion of a boring set of centered worlds in Egan 2006: 107.)

⁴¹See the discussion of the factivity condition on knowledge in §2.2 and the references therein.

⁴²More precisely, a sentence has representational truth conditions if and only if its content can be modeled by a set of worlds, or by a boring set of world, moral perspective pairs.

non-representational truth conditions, we can reap the advantages of truth conditional semantics while preserving the core expressivist idea that moral assertions do not aim to represent the world.⁴³

7.2.2 The Psychological Stage: The Functional Role of Belief

So much for the semantics. The next step is to develop an account of the functional roles of belief and desire. According to a popular line of thought, belief and desire are interrelated dispositional states. Stalnaker famously suggested that, as a first pass, we could develop this thought as follows:

To desire that p is to be disposed to act in ways that would tend to bring it about that p in a world in which one's beliefs, whatever they are, were true. To believe that p is to be disposed to act in ways that would tend to satisfy one's desires, whatever they are, in a world in which p (together with one's other beliefs) were true. (1984: 15)

For our purposes, it will be helpful to employ a modal variant of this idea:

Desire (D) S desires (at w) p if and only if in all accessible worlds v where S 's beliefs (at w) are true, S acts so as to fulfill p (together with S 's other desires at w).

Belief (B) S believes (at w) p if and only if in all accessible worlds v where p (together with S 's other beliefs at w) is true, S acts so as to fulfill the desires S holds at w .

If there are accessible worlds where S 's beliefs and desires differ from S 's beliefs and desires at the world of evaluation, this account runs into obvious counterexamples. Thus we should impose the following constraint on the accessibility relation:

Psychological Sameness Constraint For any worlds w and v : v is accessible from w only if S 's beliefs and desires are the same at v as at w .

Given this constraint, D and B offer a promising first pass account of the functional relations between desire and prosaic belief. But there is an obvious problem applying

⁴³See Yalcin (2011) for a similar defense of the use of truth conditions in an expressivist semantics. Readers may observe a resemblance between the expressivist truth conditions defended here and the assessor relativist truth conditions defended by e.g., Kölbel (2003); Lasersohn (2005); Stephenson (2007); Egan (2010, 2012); and MacFarlane (2014) for various classes of expressions. For present purposes, I will set aside the question of whether there are important differences between the two approaches. However, it is worth noting that even if there is no important difference in the truth conditions, there may still be important differences in the accompanying conceptions of belief. Assessor relativism about some subject matter is typically paired with a cognitivist view of beliefs involving that subject matter, whereas I will be using the expressivist truth conditions to implement a version of noncognitivism.

this account to moral beliefs: B presupposes that moral beliefs are true at worlds. According to our non-representational truth conditions, moral contents are not true or false at worlds, but only at world, moral perspective pairs.

Luckily, there is a straightforward way of extending D and B. Here's the big-picture, intuitive thought: for an agent to believe p is for her to be disposed to treat p as true *relative to her moral perspective*.

To make this more precise, let m_S^w denote S's moral perspective at w . (Think of this as a specification of the contents of S's moral attitudes at w .) We now give our functional roles a noncognitivist twist:

Noncognitivist Desire (ND) S desires (at w) p if and only if in all accessible worlds v such that S's beliefs (at w) are true at $\langle v, m_S^v \rangle$, A acts so as to fulfill p (together with S's other desires at w).

Noncognitivist Belief (NB) S believes (at w) p if and only if in all accessible worlds v such that p (together with S's other beliefs at w) is true at $\langle v, m_S^v \rangle$, S acts so as to fulfill the desires S holds at w .

These noncognitivist functional roles are conservative in the sense that they are equivalent to our original modal conditions (D and B) when it comes to prosaic beliefs. After all, the truth conditions for prosaic contents are insensitive to the moral perspective in the circumstance of evaluation.

So the difference between the two characterizations of the functional roles of belief and desire only emerges when we turn to moral beliefs. And it is here where we finally get the noncognitivist explanation we sought. In particular, by combining ND and NB with our expressivist semantics, we can derive the result that a moral belief—say, a belief that charitable giving is good—is necessarily grounded in some conative attitude.

To see this, let's look at how this approach validates the Modal Connection, according to which certain conative attitudes—specifically, the moral attitudes—are both necessary and sufficient for moral beliefs. Start with the necessity direction:

Necessity If S believes charitable giving is good, then S has the moral attitude towards charitable giving.

To prove this, it will be helpful to use a consequence of our expressivist truth conditions. Our truth conditions tell us that GOOD is true at some $\langle w, m \rangle$ if and only if m highly ranks charitable giving at w . Now consider the special instance of these truth conditions where the moral perspective is S's moral perspective at w (m_S^w). We get the result that GOOD is true at $\langle w, m_S^w \rangle$ if and only if m_S^w highly ranks charitable giving. Given that moral perspectives are representations of moral attitudes, this in turn holds if and only if S has the moral attitude, at w , towards charitable giving at w . So we have:

Truth-Desire Link GOOD is true at $\langle w, m_S^w \rangle$ if and only if S has the moral attitude, at w , towards charitable giving.

Equipped with the Truth-Desire Link, we can now prove Necessity:

Proof. Assume S believes GOOD at some world w . Let v be any accessible world that meets the following constraint: S's beliefs (at w) are true at $\langle v, m_S^v \rangle$. So GOOD is true at $\langle v, m_S^v \rangle$. By Truth-Desire Link, it follows that at v S has the moral attitude towards charitable giving. By the Psychological Sameness Constraint, it follows that at w S also has the moral attitude towards charitable giving. \square

Two remarks about this proof are in order. First, this proof assumes that there is at least one accessible world v such that S's beliefs (at w) are true at $\langle v, m_S^v \rangle$. Is this assumption warranted? The intuitive thought here is that all of your beliefs *could have* been true—i.e., there's at least one accessible circumstance of evaluation where all of your beliefs come out true. Note that the prosaic version of this assumption—viz., that there's at least one accessible world where all of your prosaic beliefs come out true—is already presupposed by our original functional roles for prosaic belief and desire (B and D). This assumption is arguably an idealization; for example, it cannot capture agents with inconsistent beliefs. But insofar as we are willing to make this idealization when it comes to prosaic beliefs, we should be willing to make an analogous idealization when it comes to moral beliefs.

Second, this proof did not actually use ND or NB. Rather, it used our expressivist truth conditions, together with a particular presupposition of ND and NB (the assumption remarked on in the preceding paragraph). However, we do need ND and NB themselves—not just their presupposition—in order to prove the other direction of the Modal Connection, to which we now turn:

Sufficiency If S has the moral attitude towards charitable giving, S believes charitable giving is good.

Proof. Assume that, at w , S has the moral attitude towards charitable giving. As before, let v be some accessible world such that S's beliefs (at w) are true at $\langle v, m_S^v \rangle$. Since S has the moral attitude towards charitable giving at w , it follows, by the Psychological Sameness Constraint, that S has the moral attitude towards charitable giving at v . By the Truth-Desire Link, GOOD is true at $\langle v, m_S^v \rangle$. Moreover, since S's beliefs (at w) are all true at $\langle v, m_S^v \rangle$, we can infer that at v S acts so as to fulfill S's desires at w (by ND (\Rightarrow)). Since v was arbitrarily chosen, it follows that for every accessible world u such that GOOD, together with S's other beliefs at w , is true at $\langle u, m_S^u \rangle$, S acts so as to fulfill the desires S holds at w . By NB (\Leftarrow), S believes GOOD at w . \square

So by combining our expressivist semantics with ND and NB, we validate the Modal Connection. Of course, grounding noncognitivism ultimately want more than this; they want a grounding connection. The foregoing account promises to deliver this as well. According to the account just sketched, to hold some moral belief is to be in a particular dispositional state—a state that *depends* on having a certain moral perspective. Clearly,

the dependence in question is metaphysical rather than causal. And so the conjunction of our expressivist semantics with NDB delivers what we sought: a principled explanation of *why* moral beliefs are necessarily grounded in conative states.

Taking stock: here I've sketched one way of meeting the explanatory demand facing grounding noncognitivism. The approach just sketched may well require refinement. In particular, relied on a modal variant of a simple dispositional theory of belief and desire. Both the simple dispositional theory and the modal variant thereof may well require refinement.⁴⁴ In future work, I hope to explore in more detail the various options for formulating a unitary functional role for moral beliefs and desires that delivers grounding noncognitivism as a consequence.

7.3 How Is This Noncognitivism?

Let us now turn to consider another pressing question for grounding noncognitivism: is the view noncognitivist in name only? Sure, the view forges a close connection between conative attitudes and beliefs, but so do many versions of subjectivism. How, then, does grounding noncognitivism merit its title?

By way of answer, it will be helpful to compare grounding noncognitivism in some detail to one of the most sophisticated versions of subjectivism: Dreier's speaker relativism (1990). On Dreier's view, in any given context of utterance the content of a moral expression such as 'good' will be some natural property selected by the contextually determined moral perspective ('moral system', in Dreier's parlance). To illustrate, suppose that in the simplest case the contextually determined moral perspective is the speaker's. And suppose that Jane's moral perspective gives top marks to actions that promote happiness. Then the content of her utterance of (11) ('Charitable giving is good') will be the proposition: *Charitable giving promotes happiness*.

There are at least two important differences between grounding noncognitivism and speaker relativism. The first concerns their pictures of moral assertion. To use the terminology introduced in §7.2.1, Dreier takes moral assertions to have representational truth conditions. Thus for Dreier, the truth-value of Jane's utterance of (11) depends on the world and the world alone; it is true if the world is one where charitable giving promotes happiness, false otherwise.

As a result, speaker relativism runs into familiar difficulties when it comes to moral disagreement. Recall vicious Vic, who values suffering. When Vic scoffs:

(12) Charitable giving is not good!

he seems to disagree with Jane. Speaker relativists have trouble explaining this. After all, the content of his utterance is a proposition such as *Charitable giving does not promote suffering*, which is perfectly consistent with the content of Jane's utterance of (11).

⁴⁴For concerns about a simple dispositional theory, see Stampe 1987; Velleman 2000: chp.11.

Speaker relativism likewise struggles to explain why Vic might respond to Jane's utterance by exclaiming, 'That's false!', as well as why Jane might respond to Vic's utterance in kind.⁴⁵

By contrast, while grounding noncognitivism can be implemented *via* different semantic frameworks, we have seen that one natural implementation is the Gibbardian approach developed in §7.2.1. On this semantics, the content of Jane's utterance of (11) is a particular set of world, moral perspective pairs: GOOD. And the content of Vic's utterance of (12) is the set: $\{\langle w, m \rangle \mid m \text{ does not highly rank charitable giving at } w\}$. These two contents are inconsistent: there is no $\langle w, m \rangle$ that belongs to both. Hence this framework has no trouble accounting for the intuition that Jane and Vic disagree. Our Gibbardian machinery likewise explains why Jane and Vic will be inclined to deem each others' utterances *false*. The content of Jane's utterance is false relative to Vic's moral perspective, and the content of his utterance is false relative to her moral perspective.⁴⁶

Thus one important difference between speaker relativism and grounding noncognitivism (as developed here) concerns their rival pictures of moral assertion. A second—and more crucial—difference concerns their pictures of moral *belief*. What, according to speaker relativism, is required for having a moral belief? As Dreier (2007) notes, when Jane says:

(13) Vic believes that torture is good.

she does not seem to convey that Vic believes that torture instantiates the property ranked highest by *her* moral perspective (promoting general happiness). More plausibly, she conveys that Vic believes that torture instantiates the property ranked highest by *his* moral perspective. Generalizing: 'S believes *x* is good' is true if and only if S believes that *x* exhibits some property F_S , where F_S is whatever natural property is ranked highest by S's moral perspective.⁴⁷

Thus developed, speaker relativism leads to a hybrid account of moral belief. After all, the view holds that believing *x* is good requires both (i) possessing a prosaic belief

⁴⁵Of course, speaker relativists are not without replies. See Dreier 2009 for detailed discussion of the problem of disagreement; see Plunkett and Sundell 2013 for a general argument that disagreement does not require inconsistency in content. Unfortunately, discussion of these responses would take me too far afield. For my purposes, it is enough to highlight that moral disagreement poses at least a *prima facie* challenge for speaker relativism.

⁴⁶Here I assume that people will be inclined to judge an utterance false if it is false relative to *their* circumstance of evaluation. (In this regard, my approach more closely resembles assessor relativism than speaker relativism—see fn. 43 and the references therein.) However, this inclination need not be indefeasible. One might hold that its strength is affected by a variety of conversational factors, such as the question under discussion. (See Beddor and Egan 2018 for an assessor relativist view of epistemic modals that incorporates this sort of variability.)

⁴⁷This account of moral beliefs is also necessary if the speaker relativist wants to underwrite motivational internalism about *belief*, as opposed to motivational internalism about assertion. (If Jane's utterance of (13) conveyed that Vic believes that torture exhibits whatever property *her* moral perspective ranks most highly, we'd have no reason to expect Vic to be motivated to act by the belief in question, except insofar as we have reason to think that Vic's moral perspective resembles Jane's.)

that x is F , (ii) possessing certain desire-like states that rank F highly. Consequently the view faces the same difficulties that arose for hybrid noncognitivism more generally (§6). In particular, it faces the problem of explaining why Vic's belief that torture is good does not qualify as knowledge, or as epistemically rational, even though the prosaic component of his belief has these statuses.

Grounding noncognitivism avoids this problem. After all, grounding noncognitivists deny that moral belief is some distinct state composed of prosaic and conative elements. Rather, they propose that moral belief has the very same functional role as prosaic belief. Consequently, a moral belief gets its epistemic status in much the same fashion that any prosaic belief does.

Some may think that this is too quick. After all, grounding noncognitivism doesn't actually tell us the conditions under which a belief—moral or prosaic—acquires some particular epistemic status. While this is correct, grounding noncognitivism is easier to integrate with standard epistemological answers to this question. For example, grounding noncognitivists can help themselves to a standard-issue evidentialist view of epistemic rationality and justification, according to which S 's belief is epistemically rational/justified if and only if S has sufficiently strong evidence in favor of its content. Of course, many tricky details remain to be filled in—for example, explaining what it means to have evidence in favor of some moral content—but, as we have seen, there are at least some *prima facie* plausible ways of answering these questions.⁴⁸ By contrast, Dreier cannot adopt this simple evidentialist picture. After all, Vic's belief that murder is good is not epistemically rational, even though its content (that torture causes suffering) is well-supported by his evidence.

Grounding noncognitivism is thus importantly different from sophisticated subjectivist views. Moreover, even if some readers remain reluctant to bestow the noncognitivist title upon it, they should hopefully concede that the view captures many central noncognitivist ideas. It captures the idea that moral beliefs do not aim to accurately represent the world. It likewise captures the Modal Connection, according to which one has a moral belief if and only if one has a certain desire. As a result, it is compatible with the Argument from Motivation. For my purposes, this is enough. Ultimately, the interesting issue is not one of nomenclature—what's in a name?—but rather what the view delivers. If I'm right, grounding noncognitivism delivers many core tenets of noncognitivism without succumbing to the Epistemic Evaluability Challenge.

⁴⁸For example, in §5.3 I suggested understanding talk of evidential support as talk about the credences one ought to adopt, where the 'ought' here is an expression of an epistemic attitude. (There I argued that noncognitivists cannot appeal to this account insofar as they are trying to use evidence responsiveness to explain the functional role of moral belief. This criticism does not apply to grounding noncognitivists who adopt NDB, since NDB makes no reference to evidence responsiveness.)

7.4 Why Aren't Ordinary Desires Epistemically Evaluable?

A final concern for grounding noncognitivism is that it overgenerates beliefs. If some desire-like attitudes can ground beliefs, why can't a non-moral desire, such as a desire to quench my thirst, also ground a belief? But if it can, the Functional Role Hypothesis would seem to entail that the resulting belief is epistemically evaluable.

While this is a serious concern, grounding noncognitivists have two possible replies. The first is to embrace the supposedly objectionable consequence, but deny that the consequence is, on examination, objectionable. According to this response, every desire-like attitude grounds a normative belief, though only the moral attitude grounds moral beliefs. For example, my desire to quench my thirst grounds a belief that it would be good to quench my thirst. It's just that the relevant sense of 'good' is prudential rather than moral.⁴⁹ Moreover, the response runs, this prudential belief is epistemically evaluable: we can describe it as epistemically rational or irrational; under the right conditions, it qualifies as knowledge. And here too the epistemic evaluability of the belief is explained by its functional role.

Doesn't this lead to the conclusion that the desires in question are epistemically evaluable, in contradiction to our starting observation (§2)? No. Recall that for grounding noncognitivists, it is crucial to distinguish between the *grounds* and the *grounded*—between the moral attitude and the moral belief that this attitude grounds. While the moral belief is epistemically evaluable, the moral attitude is not (§7.1). Similarly, while the prudential belief (that thirst-quenching is good) is epistemically evaluable, the underlying desire to quench my thirst is not. And the explanation for this is once again given by the Functional Role Hypothesis: the desire to quench my thirst is not epistemically evaluable because it does not play the B-role. And so this response is perfectly consistent with our starting observation that it seems odd to describe desires as epistemically rational, or as items of knowledge.

An alternative response is to deny that all desires ground beliefs. According to this response, there is something special about the moral attitude in virtue of which it—and it alone—gives rise to normative beliefs. What is this special something? Different answers are possible, and a full assessment of the various alternatives will need to be left to future work. However, let me briefly sketch one potentially promising option.

Suppose we follow Blackburn in holding that part of what distinguishes the moral attitude from ordinary desires is that it is important to us that others share our moral attitudes. Given this importance, it would be useful to have a linguistic mechanism for coordinating our moral attitudes. In particular, it would be useful to have a speech act whose conventional function is to get one's interlocutors to share one's moral perspective. A speech act with the sort of Gibbardian content sketched in §7.2.1 is well-suited to serve this role. Why is this? In contemporary philosophy of language and linguistics, it's common to think of an assertion as a proposal to get one's audience to believe its content

⁴⁹A response along these lines is in line with Gibbard's (2003) framework. On Gibbard's view, all 'ought' judgments are planning states, and only some of these judgments are specifically moral.

(Stalnaker 1978). By combining this picture of the pragmatic force of assertion with our Gibbardian semantics, we get the result that when Jane asserts (11), she is proposing that her interlocutors believe GOOD. Given our account of what it is to believe a set of world, moral perspective pairs (§7.2.2), anyone who believes this content will thereby come to have a moral perspective that assigns high marks to charitable giving. Thus Gibbardian contents are apt vehicles for coordinating our moral attitudes: by making an assertion with this sort of content, you not only express your moral perspective, you also try to get your audience to share it.⁵⁰

Of course, we *could* enrich our language with terms whose function is to coordinate our non-moral desires. Were we to invent such terms, the semantics values of sentences containing these terms would presumably be some sort of Gibbardian contents.⁵¹ However, according to the response under consideration, there is a reason why our language lacks such terms. There is no need to coordinate our ordinary, non-moral desires, whereas there is a need to coordinate our moral attitudes.

There are, then, two ways for grounding noncognitivists to respond to the overgeneration worry. First, they can let a thousand attitudes ground: all desires ground normative beliefs, and each of these is epistemically evaluable, though the underlying desire is not. Second, they can find some feature that distinguishes the moral attitude from other desires, and then tell a story about why an attitude with this feature is uniquely suited to ground moral beliefs.

8 Conclusion

Moral and prosaic beliefs are epistemically evaluable; ordinary desires are not. I've argued that this creates a problem for standard versions of noncognitivism, which take moral beliefs to have a desire-like functional role. After raising difficulties for some initial responses, I outlined a possible solution. The solution is to adopt a new form of noncognitivism, according to which moral beliefs are grounded in desires. Unlike traditional versions of noncognitivism, this view preserves a unified functional role for all beliefs, both moral and prosaic. As a result, it yields a simple solution to our problem: moral and prosaic beliefs are both epistemically evaluable in virtue of their functional role.⁵²

⁵⁰This emphasis on the coordinative aspects of moral assertion fits naturally with recent proposals by Yalcin (2012) and Carballo and Santorio (2016).

⁵¹In this case, their contents would not be sets of world, *moral perspective* pairs, but rather sets of ordered pairs of worlds and some entity that represents the content of any old desire-like state, be it moral or non-moral. (Cf. Gibbard's (2003) notion of a hyperplan.)

⁵²Earlier versions of this paper were presented at Oxford University, the Central APA, and the University of Colorado-Boulder; I am grateful for the feedback I received on each of these occasions. Many thanks are due to David Black, Annette Bryson, Roger Crisp, Andy Egan, Alvin Goldman, Hsueh Qu, Nico Kirk-Giannini, Abelard Podgorski, Ernie Sosa, Weng Hong Tang, and four anonymous referees at *Philosophers' Imprint* for very helpful comments on earlier drafts. I am particularly grateful to Simon Goldstein, who read multiple drafts and gave tremendously useful comments on each.

References

- Jonathan Adler. *Belief's Own Ethics*. MIT Press, Cambridge, MA, 2002.
- G.E.M. Anscombe. *Intention*. Harvard University Press, Cambridge, MA, 1957.
- D.M. Armstrong. *Belief, Truth, and Knowledge*. Cambridge University Press, London, 1973.
- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- Lynne Rudder Baker. Why constitution is not identity. *The Journal of Philosophy*, 94:599–621, 1997.
- Tim Bayne and Elisabeth Pacherie. In defence of the doxastic conception of delusions. *Mind and Language*, 20(2):163–188, 2005.
- Bob Beddor. *Reduction in Epistemology*. PhD thesis, Rutgers University, 2016.
- Bob Beddor. Subjective disagreement. *Noûs*, forthcoming.
- Bob Beddor and Andy Egan. Might do better: Flexible relativism and the QUD. *Semantics and Pragmatics*, 11, 2018.
- Gunnar Björnsson and Tristram McPherson. Moral attitudes for non-cognitivists: Solving the specification problem. *Mind*, 123(408):1–38, 2014.
- Gunnar Björnsson, Caj Strandberg, Ragnar Francén Olinder, John Eriksson, and Frederik Björklund, editors. *Motivational Internalism*. Oxford University Press, New York, 2015.
- Simon Blackburn. *Spreading the Word*. Clarendon Press, Oxford, 1984.
- Simon Blackburn. *Essays in Quasi-Realism*. Oxford University Press, Oxford, 1993.
- Simon Blackburn. Securing the notes: moral epistemology for the quasi-realist. In Timmons, editor, *Moral Knowledge? New Readings in Moral Epistemology*. Oxford University Press, New York, 1996.
- Simon Blackburn. *Ruling Passions*. Clarendon Press, Oxford, 1998.
- Daniel Boisvert. Expressive-assertivism. *Pacific Philosophical Quarterly*, 89:169–203, 2008.
- Michael Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.
- Ross Cameron. Turtles all the way down: Regress, priority, and fundamentality. *The Philosophical Quarterly*, 58:1–14, 2008.
- Alejandro Pérez Carballo and Paolo Santorio. Communication for expressivists. *Ethics*, 126(3): 607–635, 2016.
- J. Adam Carter and Matthew Chrisman. Is epistemic expressivism incompatible with inquiry? *Philosophical Studies*, 159:323–339, 2012.
- Matthew Chrisman. From epistemic contextualism to epistemic expressivism. *Philosophical Studies*, 135(2):225–254, 2007.
- Matthew Chrisman. Expressivism, inferentialism, and saving the debate. *Philosophy and Phenomenological Research*, 77(2):334–358, 2008.
- Matthew Chrisman. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press, Oxford, 2015.
- Stephen Darwall, Allan Gibbard, and Peter Railton. Toward *Fin de siècle* ethics: Some trends. *The Philosophical Review*, 101(1):115–189, 1992.
- Cian Dorr. Non-cognitivism and wishful thinking. *Noûs*, 36:97–103, 2002.
- Jamie Dreier. Internalism and speaker relativism. *Ethics*, 101(1):6–26, 1990.
- Jamie Dreier. Negation for expressivists: A collection of problems with a suggestion for their solution. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1, pages 217–233. Oxford University Press, New York, 2006.

- Jamie Dreier. Moral relativism and moral nihilism. In Copp, editor, *The Oxford Handbook of Ethical Theory*. Oxford University Press, New York, 2007.
- Jamie Dreier. Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives*, 23(1):79–110, 2009.
- Andy Egan. Secondary qualities and self-location. *Philosophy and Phenomenological Research*, 72:97–119, 2006.
- Andy Egan. Disputing about taste. In Feldman and Warfield, editors, *Disagreement*. Oxford University Press, Oxford, 2010.
- Andy Egan. Relativist dispositional theories of value. *Southern Journal of Philosophy*, 50(4):557–582, 2012.
- David Enoch. How noncognitivists can avoid wishful thinking. *Southern Journal of Philosophy*, 41:527–545, 2003.
- Hartry Field. Epistemology without metaphysics. *Philosophical Studies*, 143(2):249–290, 2009.
- Hartry Field. Epistemology from an evaluativist perspective. *Philosophers' Imprint*, 18(12):1–23, 2018.
- Kit Fine. The question of realism. *Philosophers' Imprint*, 1(1):1–30, 2001.
- Kit Fine. Guide to ground. In Correia and Schneider, editors, *Metaphysical Grounding*. Cambridge University Press, Cambridge, 2012.
- Guy Fletcher and Michael Ridge. *Having it Both Ways: Hybrid Theories and Modern Metaethics*. Oxford University Press, New York, 2014.
- Peter Geach. Assertion. *Philosophical Review*, 74:449–465, 1964.
- Allan Gibbard. Contingent identity. *Journal of Philosophical Logic*, 4:187–221, 1975.
- Allan Gibbard. *Wise Choices, Apt Feelings*. Harvard University Press, Cambridge, MA, 1990.
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, MA, 2003.
- Allan Gibbard. *Meaning and Normativity*. Oxford University Press, Oxford, 2013.
- Martin Grajner. Hybrid expressivism and epistemic justification. *Philosophical Studies*, 172(9):2349–2369, 2015.
- Alex Gregory. Changing direction on direction of fit. *Ethical Theory and Moral Practice*, 15(5):603–614, 2012.
- Terry Horgan and Mark Timmons. Cognitivist expressivism. In Horgan and Timmons, editors, *Metaethics After Moore*. Oxford University Press, Oxford, 2006.
- I.L. Humberstone. Direction of fit. *Mind*, 100(400):277–281, 1992.
- David Kaplan. Demonstratives. In Almog, Perry, and Wettstein, editors, *Themes from Kaplan*. Oxford University Press, New York, 1989.
- Klemens Kappel. Expressivism about knowledge and the value of knowledge. *Acta Analytica*, 25(2):175–194, 2010.
- Antti Kauppinen. What makes a sentiment moral? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1. Oxford University Press, New York, 2010.
- Sebastian Köhler. Do expressivists have an attitude problem? *Ethics*, 123(3):479–507, 2013.
- Max Kölbel. Faultless disagreement. *Proceedings of the Aristotelian Society*, 104:53–73, 2003.
- Peter Lasersohn. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28:643–686, 2005.
- Nicholas Laskowski. The sense of incredibility in ethics. *Philosophical Studies*, forthcoming.
- Brian Loar. *Mind and Meaning*. Cambridge University Press, Cambridge, 1981.
- John MacFarlane. *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford, 2014.
- Alexander Miller. *Introduction to Contemporary Metaethics*. Polity, Cambridge, 2003.

- Sarah Moss. Epistemology formalized. *Philosophical Review*, 122(1):1–43, 2013.
- Mark Platts. *Ways of Meaning*. Routledge and Kegan Paul, Boston, 1979.
- David Plunkett and Timothy Sundell. Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13:1–37, 2013.
- James Pryor. The skeptic and the dogmatist. *Noûs*, 34(4):517–549, 2002.
- W.V. Quine. Reference and modality. In *From a Logical Point of View*, pages 145–171. Harper, New York, 1961.
- Peter Railton. Reliance, trust, and belief. *Inquiry*, 57(1):122–150, 2014.
- Michael Ridge. Saving the ethical appearances. *Mind*, 115(459):633–650, 2006a.
- Michael Ridge. Ecumenical expressivism: Finessing Frege. *Ethics*, 116:302–336, 2006b.
- Michael Ridge. Expressivism and epistemology: Epistemology for ecumenical expressivists. *Aristotelian Society Supplementary Volume*, 81(1):83–108, 2007a.
- Michael Ridge. Ecumenical expressivism: The best of both worlds? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 2. Oxford University Press, 2007b.
- Michael Ridge. The truth in ecumenical expressivism. In Sobel and Wall, editors, *Reasons for Action*. Cambridge University Press, 2009.
- Michael Ridge. *Impassioned Belief*. Oxford University Press, New York, 2014.
- Gideon Rosen. Metaphysical dependence: Grounding and reduction. In *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press, Oxford, 2010.
- Jacob Ross. How to be a cognitivist about practical reason. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 4. Oxford University Press, 2009.
- Bertrand Russell. *Religion and Science*. Oxford University Press, New York, 1935.
- Jonathan Schaffer. On what grounds what. In Chalmers and Wasserman, editors, *Metametaphysics*. Oxford University Press, Oxford, 2009.
- Mark Schroeder. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford University Press, Oxford, 2008a.
- Mark Schroeder. What is the Frege-Geach problem? *Philosophy Compass*, 10:703–720, 2008b.
- Mark Schroeder. Hybrid expressivism: Virtues and vices. *Ethics*, 119:257–309, 2009.
- Mark Schroeder. *Noncognitivism in Ethics*. Routledge, New York, 2010.
- Mark Schroeder. Tempered expressivism. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press, New York, 2013.
- Mark Schroeder. What makes reasons sufficient? *American Philosophical Quarterly*, 52(2):159–170, 2015.
- Eric Schwitzgebel. A phenomenal, dispositional account of belief. *Noûs*, 36:249–275, 2002.
- Russ Shafer-Landau. *Moral Realism: A Defence*. Oxford University Press, New York, 2003.
- Nishi Shah and David Velleman. Doxastic deliberation. *The Philosophical Review*, 114(4):497–534, 2005.
- Alex Silk. How to be an ethical expressivist. *Philosophy and Phenomenological Research*, 91(1):47–81, 2015.
- Michael Smith. The Humean theory of motivation. *Mind*, 96:36–61, 1987.
- Michael Smith. *The Moral Problem*. Blackwell, Oxford, 1994.
- Robert Stalnaker. Assertion. In *Syntax and Semantics*, pages 315–332. New York Academic Press, New York, 1978.
- Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, MA, 1984.
- Dennis Stampe. Defining desire. In Marks, editor, *The Ways of Desire*. Precedent, Chicago, 1987.
- Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.

- Charles L. Stevenson. The emotive meaning of ethical terms. *Mind*, 46:14–31, 1937.
- C.L. Stevenson. *Ethics and Language*. Yale University Press, New Haven, 1944.
- Bart Streumer. Do normative judgements aim to represent the world? *Ratio*, 26(4):450–470, 2013.
- Sigrun Svavarsdóttir. Moral cognitivism and motivation. *The Philosophical Review*, 108:161–219, 1999.
- Teemu Toppinen. Believing in expressivism. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press, New York, 2013.
- Kelly Trogon. Grounding, necessary or contingent? *Pacific Philosophical Quarterly*, 94(4):465–485, 2013.
- David Velleman. *The Possibility of Practical Reason*. Oxford University Press, New York, 2000.
- Seth Yalcin. Nonfactualism about epistemic modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012.
- Nick Zangwill. Non-cognitivism and consistency. *Zeitschrift für Philosophische Forschung*, 65(4): 465–484, 2011.