Moral and Epistemic Evaluations: A Unified Treatment

Bob Beddor

Philosophical Perspectives, forthcoming

1 Introduction

A rich tradition in metaethics seeks to explain the meaning of moral language in terms of desire-like attitudes. This approach can be implemented in different ways. On a contextualist implementation, moral discourse describes the desire-like attitudes of some agent(s), for example, the speaker or the speaker's community. On an expressivist implementation, moral discourse does not describe desire-like attitudes; it expresses them. On a relativist implementation, the truth-value of a moral assertion at a context of assessment depends on the desire-like attitudes of the assessor—that is, someone assessing the utterance for truth or falsity. Despite these differences, all such "attitudinal metaethics" are bound by a common thread: they analyze moral discourse in terms of conative states.

Attitudinal metaethicists sometimes propose extending their approach to other varieties of normative discourse, including epistemic discourse.[1] A generalized attitudinal semantics along these lines carries obvious attractions. But it also faces important challenges. An initial challenge concerns how to even spell out a generalized attitudinal semantics. While much ink has been spilled in pursuit of a precise attitudinal semantics for moral discourse, the extension to other fragments of normative language has not received a comparable degree of attention. Second, and more worrisome, some philosophers have argued that there are principled obstacles to the very idea of a generalized attitudinal semantics. For example, Boult and Köhler 2020 argue that a generalized attitudinal semantics is under-motivated, since the primary arguments for an attitudinal metaethics do not carry over to the epistemic domain. And Wodak 2017 argues that a generalized attitudinal semantics over-predicts disagreements across normative domains.

These obstacles can be used to frame a dilemma for any attempt to generalize an attitudinal semantics. On the one hand, generalizers need to show that there are sufficient commonalities between different normative domains to warrant a unified treatment. On the other hand, generalizers had better not erase the obvious differences between different flavors of normative judgment.

[1] See e.g., Blackburn 1998; Gibbard 2003; Chrisman 2007; Ridge 2007, 2018; Field 2009; Greco 2014.

This paper develops a generalized attitudinal semantics that steers through the horns of this dilemma. I start by pointing out some commonalities between the moral and epistemic domains. These commonalities cry out for a unified explanation. The rest of the paper develops a unified explanation and explores its consequences. On the view developed here, there is a variety of distinct pro-attitudes, all of which share a functional core. Moral discourse is used to voice a distinctly moral form of approval. Epistemic discourse is used to voice a distinctly epistemic form of approval. The functional similarities between the different species of approval explain the commonalities between different flavors of normative discourse. And the differences between the former explain the differences between the latter.

The view developed here has a number of further advantages. First, it is naturalistically respectable: each species of approval is explained in non-normative language. Second, it is developed with attention to the semantic details. In particular, I show how my semantics can be derived from a single lexical entry for deontic modals—a lexical entry that is fully consistent with leading analyses in the linguistics literature. Finally, it yields a simple solution to a seemingly independent problem for metaethical expressivism (Woods 2014). In an appendix, I highlight some advantages of my approach over an alternative way of developing a generalized attitudinal semantics, due to Gibbard 2003.

2 Commonalities

2.1 Linguistic commonalities

We use many of the same lexical items in both moral and epistemic evaluations. For example, we frequently use deontic modals (*may*, *must*, *ought*, *should*) to make both sorts of appraisals:

(1) a. You ought to [/should, must, may] help those in need. (**moral**)
b. You ought to [/should, must, may] believe in evolution. (**epistemic**)

It is usually left to context to determine whether a deontic modal has moral or epistemic flavor. However, we can use *in view of*-phrases to make our intended reading explicit (cf. Kratzer 1977):

(2) a. In view of the moral requirements, you ought to [/should, must, may] help those in need.
b. In view of the epistemic requirements, you ought to [/should, must, may] believe in evolution.

This lexical overlap is cross-linguistically robust. For example, (3)-(7) render the contrast in (1) in Italian, German, Serbian, Afrikaans, Filipino, and Japanese. In each of these languages, the same deontic modal (highlighted in blue) is used to make both the moral and the epistemic evaluation, just as in English.

(3) a. Devi [/dovresti, potresti] aiutare quelli che ne hanno bisogno. (**Italian**)
b. Devi [/dovresti, potresti] credere nell'evoluzione.

(4) a. Man sollte [/muss] Notleidenden helfen. (**German**)
b. Man sollte [/muss] an die Evolution glauben.

(5) a. Covek treba [/mora] da pomogne ljudima u nevolji. (**Serbian**)
b. Covek treba [/mora] da veruje u evoluciju.

(6) a. Jy moet hulle wat nood het help. (**Afrikaans**)
b. Jy moet in evolusie glo.

(7) a. Kailangan [/Dapat] mong tulungan ang mga nangangailangan. (**Filipino**)
b. Kailangan mong maniwala sa ebolusyon/Dapat maniwala ka sa ebolusyon.

(8) a. Komatteiru hito-o tasukeru beki da. (**Japanese**)
b. Shinkaron-o shinjiru beki da.[2]

This linguistic overlap is not limited to modals either. The language of *justification*, *rationality*, and *reasons* is also used in both moral and epistemic contexts, e.g.:

(9) a. You were justified in acting as you did. (**moral**)
b. Poirot is justified in believing the butler did it. (**epistemic**)

Our semantics for normative language should explain these commonalities.

2.2 The connection with reactive attitudes

We are not usually indifferent to perceived wrongdoings. When we judge that someone acted wrongly, we are disposed to *blame* them, unless they have a good excuse. And when we judge that someone acted rightly, we are disposed to praise them, particularly if they did so at some cost to themselves. More generally, our moral judgments are closely tied to 'reactive attitudes' (Strawson 1974).[3]

Epistemic judgments are also closely bound up with reactive attitudes. Consider the Flat Earther, who maintains that the earth is flat. They are aware of the evidence to the contrary, but they ignore it or try to explain it away. The Flat Earther seems to be blameworthy—in some sense—for holding this belief. Similarly, people are praiseworthy when their beliefs meet a particularly exalted epistemic standard, for example, when they

[2]Speakers I have consulted inform me that the same phenomenon occurs in Thai, Korean, and Mandarin Chinese. Thanks to Carlotta Pavese, Jelena Krivokapic, Andries Coetzee, Chautamanee Onsuwan, Harim Kwan, Mitcho Erlewine, and Wilkinson Daniel Wong Gonzales for these examples and their linguistic judgments.

[3]The idea that certain moral judgments are closely connected to reactive attitudes traces back to Mill. According to Mill, to judge that someone acted wrongly is to judge that they ought to be punished, 'if not by law, by the opinion of his fellow-creatures; if not by opinion, by the reproaches of his own conscience.' (Mill 1861: chp. V, par. 4). For contemporary developments of the idea that moral judgments are tied to reactive attitudes, see Gibbard 1990; Schroeder 2008; Kauppinen 2010; Sepielli 2012; Björnsson and McPherson 2014.

scrutinize their beliefs for subtle tensions, or tease out the logical implications of their beliefs through a particularly grueling deduction. This idea that beliefs can be praiseworthy and blameworthy also has epistemological payoff. For example, one prominent strategy for defusing resistance to externalist norms on belief is to insist that there is a difference between violating some epistemic norm and being blameworthy for doing so. According to this strategy, the victim of an evil demon violates an epistemic norm, but in a blameless fashion.[4]

Now, we should not deny that there are important differences between moral blame and epistemic blame. Irrational beliefs rarely rouse intense feelings of anger or righteous indignation in the same way that moral infractions do. But this is not a reason to doubt the existence of epistemic blame. Thankfully, there is by now a burgeoning literature devoted to explaining the nature of epistemic blame and how it differs from moral blame (e.g., Rettler 2018; Brown 2020; Boult forthcoming; Pichovy forthcoming). For the moment, there is no need to take a stand on how to draw the distinction; I will return to this issue in §3. The important point is that both moral and epistemic judgments are connected with reactive attitudes. And this is something that deserves explanation.

2.3 The connection with motivation

Many metaethicists maintain that there is a special connection between moral judgment and motivation.[5] To see the appeal of this thesis, imagine that Ari insists that everyone is morally obligated to worship God. But, she hastens to add, she is not at all inclined to worship God. Most people would find Ari's position puzzling. They might even doubt the sincerity of Ari's professed moral convictions.

We should be careful not to overstate the connection between moral judgment and motivation. Akrasia is a fact of life: many people (your author included) sometimes fail to act in accordance with their moral judgments. Our formulation of the connection between moral judgment and motivation had better allow for this.[6] Here's one suitably weak formulation of the connection:

Moral Judgment-Motivation Connection (MJC) If someone believes that they morally ought to ϕ, then they will be at least somewhat disposed to ϕ.

This formulation leaves room for akrasia. Ari might believe that she is morally obligated to worship God but still fail to do so, since her disposition to worship might be masked

[4]See e.g., Sutton 2007; Hawthorne and Srinivasan 2013; Williamson forthcoming. For discussion of this distinction, see e.g., Kelp and Simion 2017; Madison 2018; Greco forthcoming.

[5]Moral judgment internalists hold that there is a necessary connection between moral judgment and motivation; see Dreier 1990; Smith 1994; Korsgaard 1996, Blackburn 1998; Gibbard 2003; Wedgwood 2007; van Roojen 2010; Egan 2012, among others. Even moral judgment externalists typically grant that there is a close connection between moral judgment and motivation; they just deny that it is a necessary connection.

[6]For relevant discussion, see Stocker 1979; Smith 1994; Mele 1996; Svavarsdóttir 1999.

by various factors—listlessness, inattention, competing preferences, and the like.[7]

I want to argue that an analogous connection holds between epistemic judgment and motivation. Start with judgments about our epistemic obligations to inquire. Suppose Poirot is investigating a murder. The initial evidence indicates that the butler is the culprit. But Poirot has not yet inquired into the butler's whereabouts on the night of the crime. It's easy to fill in the details so that Poirot epistemically ought to engage in such inquiries. Now, suppose Poirot agrees with this judgment. That is, suppose he thinks: *I (epistemically) ought to inquire into the butler's whereabouts.* We would expect Poirot to pursue such inquiries.

Of course, we can imagine circumstances in which Poirot makes this judgment but does not follow through; perhaps he is lazy or distracted by other cases. But this does not show that there is *no* connection between epistemic judgment and motivation. As with moral judgment, this only shows that the connection had better be suitably qualified: if someone judges that they epistemically ought to pursue some inquiry, they will be at least somewhat disposed to do so.[8]

Next, consider our epistemic evaluations of doxastic states. Imagine that Ari claims that theism is epistemically irrational. But, she hastens to add, she herself is a committed theist. Most people would find Ari's position puzzling. They might even question the sincerity of her professed epistemic convictions.

Here too, we should not overstate the connection between epistemic judgment and motivation. Epistemic akrasia is also a fact of life. Consider the parent whose child has been arrested. Given the evidence, the parent is convinced that they epistemically ought to believe their child is guilty. But they can't bring themselves to hold this belief. Their position may be irrational, but it is psychologically possible.

These observations suggest a connection between epistemic judgment and motivation along the following lines:

Epistemic Judgment-Motivation Connection (EJC) If someone believes that they epistemically ought to ϕ, then they will be at least somewhat disposed to ϕ.

Here the variable ϕ ranges over both acts (e.g., *inquiring into the detective's whereabouts*) and doxastic states (e.g., *believing the butler did it*, *having a .8 credence in God's existence*).

[7] I'll remain noncommittal on whether to construe MJC as a universally quantified claim ranging over all agents or a generic claim restricted to normal agents. Which of these options we prefer will depend on what we think about the possibility of amoralists, who hold a moral belief while exhibiting no disposition to act in accordance with it. For discussion, see Dreier 1990; Blackburn 1998; Roskies 2003; Leary 2017.

[8] Kappel and Moeller 2014 also argue that epistemic judgments are connected to inquiry, and they also appeal to this connection in motivating an attitudinal semantics. But their formulation of the connection is importantly different. They focus on knowledge attributions; according to them, if A judges, *S knows that p*, A will be motivated to terminate inquiry into p. One worry for their argument is that this can be explained by the factivity of knowledge attributions: if A judges that *S knows that p*, A thereby judges that p, which in many cases will dispose A to cease inquiry into p (Ridge 2018). But then there is no need to appeal to an attitudinal semantics to explain Kappel and Moeller's *explanandum*. To motivate this worry, note that if we modify A's judgment with the non-factive construction, *S knows whether p*, we would not necessarily expect A to be motivated to cease inquiry into whether p.

Before moving on, let me head off a potential objection. I've been talking freely of 'motivation' to believe. Some might balk at this. In the ordinary sense of the word, we can only be motivated to *act*. I will discuss this issue in more depth in §5. For now I'll just note that nothing of substance hinges on this terminology. EJC (much like MJC) is formulated as a claim about the connection between judgments and *dispositions*. Surely we have dispositions to be in doxastic states. Readers are welcome to regard 'motivation' talk as shorthand for the relevant dispositional connections.

Other species of normative judgment also carry motivational pull. Suppose someone judges that they prudentially ought to look both ways before crossing the street. Here too, we would expect them to be at least somewhat disposed to look both ways before crossing. More generally, we might hypothesize that every genuinely normative judgment has a defeasible connection with motivation.[9]

Historically, one of the main arguments in favor of an attitudinal metaethics is that it explains MJC.[10] If every normative judgment is connected with motivation, then any explanation that is tailored to moral judgments will be insufficiently general. Rather, we should seek out a general explanation for why normative judgment—in all its forms—carries motivational oomph.

2.4 Descriptively enlightened disagreement

People disagree over moral matters. And these disagreements often persist even once all parties have been appraised of the pertinent descriptive (i.e., non-normative) facts. The existence of such 'descriptively enlightened' moral disagreements will be familiar to anyone who has taught Ethics 101. Present your students with a standard puzzle case, e.g., A has the opportunity to kill B to save a hundred. You'll quickly discover that students disagree about whether A is morally permitted to kill B. Even once you clarify all of the descriptive facts about A's motives, B's life history, and exactly how much utility B's death will bring about, the disagreement will likely continue.

Many philosophers maintain that our metaethics should explain the presence and persistence of descriptively enlightened moral disagreements. Of course, it is a matter of controversy how best to do so. For the moment, I won't take a stand on this question; in §6, I'll consider one promising explanation in detail. For now, I want to point out that the very same sort of descriptively enlightened disagreements arise in epistemology.

Here's a familiar experience for anyone who has taught Epistemology 101. It's the first day of class. You want to engage your students, so you present them with the skeptical paradox. If your experience is anything like mine, you'll find some students strongly

[9] It would be implausible to claim that every *ought*-judgment is tied to motivation. Someone might believe that, in view of the rules of etiquette, they ought to start with the outside fork, without having any inclination to do so. But this does not cast doubt on the idea that every genuinely normative judgment is tied to motivation; it just shows that not every *ought*-judgment is genuinely normative.

[10] This motivation is particularly explicit in the writings of metaethical expressivists; see e.g., Stevenson 1963; Blackburn 1998; Gibbard 2003. But it also crops up in the work of contextualists and relativists; see e.g. Dreier 1990; Egan 2012.

sympathize with the skeptical conclusion whereas others deem it absurd. Moreover, this disagreement will persist even once you have clarified all of the pertinent descriptive facts. Even once you've made clear that a particular agent A is not envatted, or deceived by a demon, etc., some students will maintain that A ought to suspend belief in the external world, or at least regard the skeptical hypothesis as a live contender. Others will disagree. Similar points can be made using other recalcitrant disagreements in epistemology. Just consider disputes between internalists and externalists about whether the brain in a vat is justified in believing in the external world, or about whether the unwitting clairvoyant ought to believe the outputs of their clairvoyance.

2.5 Looking Ahead

I've identified four commonalities between moral and epistemic evaluations. These commonalities cry out for a unified explanation. The rest of this paper develops a unified attitudinal semantics that does the job.[11]

3 Varieties of approval

3.1 Distinguishing different forms of approval

Philosophers commonly distinguish between representational mental states and non-representational mental states. Within the latter class, they frequently distinguish between various pro-attitudes: desires, preferences, intentions, approval and disapproval.

Around here, most philosophers stop subdividing mental states. But we can go further: pretty much any pro-attitude comes in different varieties. Take *approval*. Suppose Beth is a business mogul who has grown accustomed to the finer things in life. One day she happens across Singer's 'Famine, Affluence, and Morality.' Try as she might, she cannot find any mistakes in Singer's argument. She begrudgingly admits that she is morally obligated to donate a hefty chunk of her fortune to charity. Alas, doing so would prevent her from buying the new jet that she has been coveting. *Question:* Does Beth approve of giving her money to charity? There is a sense in which the answer is 'Yes', and a sense in which the answer in 'No.' From the moral point of view, she approves of it. But from a prudential point of view, she disapproves.

There also seems to be an epistemic form of approval, which can come apart from both moral and prudential approval. Consider the lonely researcher, who spends all of their time devoted to scholarly pursuits, neglecting family and friends. If their research

[11] I make no claim that these are the only interesting parallels between moral and epistemic evaluations. Another intriguing commonality is suggested by Greco 2015, who argues that both moral and epistemic judgments give rise to the 'Open Question Phenomenon' (Moore 1903). I think this is right; in fact, this is exactly what we should expect, given that both give rise to descriptively enlightened disagreements. For my purposes, I won't explore the Open Question Argument in any detail. However, the sort of unified attitudinal framework developed here offers one way of explaining why all normative questions are open in this sense.

yields important discoveries, we might think there is something epistemically commendable about their pursuits. But we might also think there is something regrettable about their single-minded devotion to their studies. We epistemically approve of their tireless dedication to research, while disapproving of it from the moral, prudential, and (perhaps) all-things-considered perspectives.

Some might worry that we can only understand the differences between these forms of approval using normatively charged vocabulary: we can only understand 'moral approval' in terms of the 'moral point of view', and can only understand 'epistemic approval' in terms of the 'epistemic point of view.' This would be unwelcome, given the explanatory ambitions associated with an attitudinal semantics. After all, attitudinal semanticists aim to explain normative thought and talk in non-normative, naturalistically respectable terms. Call this 'the Individuation Problem.'[12]

While this is an important challenge, we should not despair of meeting it. In the rest of this section, I introduce two strategies for developing a fully naturalistic account of what distinguishes different species of approval: *the basis strategy* and *the functional role strategy*. I should note at the outset that nothing I say hinges on the relevant attitude being *approval*, rather than, say, preference or intention. The strategies I introduce could be just as well used to distinguish between moral and epistemic preferences, or between moral and epistemic intentions.

3.2 The basis strategy

Start with all-things-considered approval. Often when we all-thing-considered approve of some course of action, it's because we believe it to be conducive to some further ends that we non-instrumentally desire. All-things-considered, I approve of ordering Thai for dinner, since I believe that ordering Thai will be most conducive to the satisfaction of my gustatory desires.

According to the basis strategy, we can analyze different varieties of approval as dispositions towards all-things-considered approval—dispositions that are based in different non-instrumental desires. In the case of moral approval, the relevant non-instrumental desire might be a *pro tanto* desire for fairness, social harmony, and wellbeing. In the case of prudential approval, the relevant non-instrumental desire might be a *pro tanto* desire for the wellbeing of a particular agent. To illustrate, go back to Beth, our Singer-reading mogul. She has some disposition to all-things-considered approve of donating her income to charity—a disposition that's based in a *pro tanto* desire to alleviate suffering, together with a means-ends belief that donating her income will be conducive to that end. However, she also has a countervailing disposition to all-things-considered approve of keep-

[12]The problem of characterizing the 'moral attitude'—that is, the distinctive conative attitude involved in moral judgment—has received a bit of attention; see, e.g., Miller 2003; Kauppinen 2010; Köhler 2013; Björnsson and McPherson 2014. The problem of characterizing the 'epistemic attitude'—that is, the distinct conative attitude involved in epistemic judgment—has received less discussion. Indeed, some proponents of an attitudinal metaepistemology have acknowledged the problem but despaired of providing a substantive answer; see e.g., Field 2009: 259-260.

ing her fortune—a disposition that is based in a *pro tanto* desire for her own wellbeing, together with a belief that donating is inimical to that end. According to the present proposal, the former disposition constitutes her moral approval of donating her fortune; the latter disposition constitutes her prudential disapproval of donating.[13]

In the case of epistemic approval, the basis might be a non-instrumental *pro tanto* desire for the attainment of truth and avoidance of error. According to this proposal, we are disposed to all-things-considered approve of the lonely scholar's single-minded pursuit of their research—a disposition that's based in a *pro tanto* desire for our scholar to attain accurate doxastic states. However, we also have a countervailing disposition to disapprove of their myopic pursuits—a disposition that is based in other considerations (say, the goods of family and friendship).

3.3 The functional role strategy

According to the functional role strategy, we can analyze different varieties of approval in terms of their causal relations with other attitudes and behaviors. Now, since all varieties of approval are conative attitudes, we should expect them all to have broadly similar functional profiles. But beneath these broad similarities may lurk important differences.

What are these differences? Our discussion of the reactive attitudes (§2.2) suggests a plausible hypothesis. Perhaps moral approval is closely tied to some distinctly moral reactive attitudes. On this view, for Beth to morally approve of giving to charity is for her to be disposed to resent and blame those who fail to give, when they are in a position to do so, and to feel guilt if she herself fails to give. Perhaps epistemic approval is also close tied to the reactive attitudes, only distinctly epistemic versions thereof.

Of course, going this route pushes the Individuation Problem back a step: we are now stuck with the problem of distinguishing the moral reactive attitudes from the epistemic reactive attitudes. However, the literature already contains some promising suggestions for how to draw this distinction. For example, Boult forthcoming develops a proposal that builds on Scanlon's 2008 account of blame. According to Scanlon, all blame involves an intention to modify one's relationships with the blamed individual. According to Boult, what distinguishes epistemic blame is that the relevant relationship is one of trust: when we epistemically blame someone, we form an intention to suspend our default trust in that person, at least within a restricted domain.[14] Another option would be to combine the functional role strategy with the basis strategy: perhaps what distinguishes epistemic blame from moral blame is that the former is based in epistemic considerations (e.g., accuracy) whereas the latter is based in moral considerations (e.g., fairness, wellbeing).[15]

[13]Does this mean that every agent who morally approves of some course of action has a *pro tanto* desire for fairness, harmony, and wellbeing? This might seem like a stretch. However, we can finesse this point by identifying moral approval with a disposition to all-things-considered approval that *in normal agents* is based in a non-instrumental desire for such ends. Cf. Björnsson and McPherson 2014 for a related view on which the moral attitude is distinguished, in part, by the grounds that paradigmatically give rise to it.

[14]Cf. Kauppinen 2018 for a related story about the distinction between epistemic and non-epistemic norms.

[15]In developing the functional role strategy, I have focused on the different causal roles these attitudes

3.4 Taking stock

I've sketched two strategies for individuating varieties of approval. Both strike me as promising. We have also seen that they are not mutually exclusive. Indeed, I'm inclined to think that a complete answer to the Individuation Problem will synthesize elements of each. For our purposes, we need not take a stand on the final answer to the Individuation Problem. It's enough that we have two plausible paths forward.

Some might worry that even if the Individuation Problem can be solved in this fashion, the picture of the mind that emerges is unduly complex. Why saddle the mind with all these different pro-attitudes? Surely parsimony favors a cleaner view, according to which there is just a single attitude of approval.

However, we should only invoke simplicity to choose between theories that explain the data equally well. And the psychological data here are complex. Go back to our Singer-reading business mogul, Beth. We noted that there is some sense in which she approves of donating her fortune. But there is also a sense in which she disapproves of doing so. These two senses will reveal themselves in her psychological and behavioral dispositions. Even if she ultimately chooses to continue funding her lavish lifestyle, we would expect her to hesitate before doing so, and to feel occasional pangs of remorse. This requires explanation. Similarly, if she instead chooses to donate her funds, we would expect her to occasionally pine after that never-purchased jet—a fact that also requires explanation. By positing different varieties of approval, we explain these facts.

Moreover, even if attitudinal metaethicists had no interest in generalizing their account to encompass other species of normativity, they would still face the Individuation Problem. A familiar challenge for any attitudinal metaethics comes from the fact that not every pro-attitude makes for a genuinely moral judgment. I might desire to eat foie gras without thinking that I am *morally* required to do so. So we need to say something about what distinguishes the pro-attitude that constitutes moral judgment from other pro-attitudes. To do so, we will need to engage in the project I have undertaken here.

4 Putting the attitudes to work

The next step is to put these varieties of approval to work in our semantics. This section develops a 'proof of concept' of what this semantics might look like. I start (§4.1) by developing a general attitudinal semantics for normative appraisals made using deontic modals. I proceed (§4.2) to extend this semantics to other normative language, focusing on justification ascriptions.

have in our psychologies. Another approach would be to distinguish different species of approval in terms of the social functions that these attitudes serve. Borrowing from Gibbard's 1990 account of 'accepting a norm', one might suggest that the function of moral approval is to foster coordination in our behavior. Similarly, one might follow Dogramaci 2012 and hold that the function of epistemic approval is to foster coordination in our belief-forming methods.

4.1 Deontic modals

The standard semantics for modals in the linguistics literature is due to Kratzer 1981, 1991, 2012. On Kratzer's analysis, the extension of any modal depends on two parameters. The first parameter is a modal base f, a function from a world to a set of propositions that delivers an accessibility relation over worlds. The second parameter is an ordering source g, a function from a world to a set of propositions that induces a ranking over worlds. A necessity modal such as *ought* universally quantifies over the g-best of the f-accessible worlds. Formally:

KRATZERIAN *OUGHT* $[\![\text{Ought } \phi]\!]^{c,f,g,w} = 1$ iff $\forall v \in \textsc{Best}_{g(w),f(w)} : [\![\phi]\!]^{c,f,g,v} = 1$,
where $\textsc{Best}_{g(w),f(w)}$ is the set of worlds in $\bigcap f(w)$ ranked highest by $g(w)$.

One of the main advantages of Kratzer's semantics is that it provides a uniform semantics for all modals. This makes it well-suited to capture the differences between different types of deontic modals, which is our focus here. Consider (2), repeated here as (10):

(10) a. You ought to [/should, must] help those in need.
b. You ought to [/should, must] believe in evolution.

As we noted earlier, (10a) is naturally read as making a moral evaluation. On Kratzer's semantics, this reading is captured using a moral ordering source—an ordering source that ranks worlds based on the extent to which they conform with the moral requirements. On this reading, (10a) is true if and only if all of the morally best worlds in the modal base are worlds where the addressee helps those in need. By contrast, (10b) is naturally read as making an epistemic evaluation. This reading is captured using an epistemic ordering source—that is, an ordering source that ranks worlds based on the extent to which the conform with the epistemic requirements.[16]

As it stands, Kratzer's semantics does not tell us what it means for a world to conform to the moral requirements or the epistemic requirements—this is treated as a black box. This silence is well and good, for Kratzer's purposes: she aims to give a semantics that prescinds from metanormative controversies. However, philosophers who seek a semantic resolution of such controversies will want something more. They will want to an informative account of what it is for a world to be morally or epistemically best.

Here is where varieties of approval can earn their keep. The job of an ordering source is to induce a ranking over worlds. Varieties of approval are well-suited for this task. A toy example: suppose the only thing I morally approve of is utility maximization. Then my state of moral approval will rank a world w over a world v just in case w contains

[16]Epistemic ordering sources are not widely discussed in the literature, and should be distinguished from the more common notion of an epistemic modal base. An epistemic modal base is used to evaluate the meanings of epistemic modals (e.g., *The butler might have done it*). These epistemic modals are not deontic, and consequently are often thought to be evaluated using an empty ordering source. By contrast, we are invoking an epistemic ordering source to evaluate a specific type of deontic modal: a modal evaluation of what someone epistemically ought to do or believe, as in (10b).

more overall utility than v. And suppose the only thing I epistemically approve of is credal accuracy. Then my state of epistemic approval will rank a world w over a world v just in case w contains more credal accuracy than v.

One advantage of this approach is that it allows us to accept the Kratzerian semantics without any *ad hoc* revisions. To see how this would work, let a moral approval function m be a function from a world to a set of propositions representing what some agent (or group of agents) morally approves of. Similarly, let an epistemic approval function e be a function from a world to a set of propositions representing what some agent (or group of agents) epistemically approves of. Then moral and epistemic uses of deontic modals are just special instances of Kratzer's semantics for *ought*; they are the special instances where the ordering source is given by m and e respectively.[17]

My strategy for developing a generalized attitudinal semantics faces various choice points. Are the ordering sources determined by the context of utterance, as suggested by Kratzer 1981? If so, we get a form of contextualism. Relative to any context of utterance, the content of e.g., (10a) will be a classical proposition (a set of worlds). But which proposition (10a) asserts will depend on the context. Or is the ordering source an independent feature of the circumstance of evaluation, which floats free from the context? If so, then the view that emerges has good claim to be considered a form of expressivism, similar in structure to those defended by Gibbard 1990, 2003; Yalcin 2012 and Silk 2014. Relative to any context, the content of e.g., (10a) will not be a set of worlds, but rather a set of world, moral approval function pairs. It's also easy to give our semantics a relativist spin. All that is needed is to let our circumstances of evaluation be centered worlds, and let the moral and epistemic approval functions be functions from a centered world $\langle w, a\rangle$ to a set of propositions reflecting what a morally/epistemically approves of at w.

Each of these options faces further choice points. Within the contextualist camp, we can ask: whose conative attitudes are relevant? Does context always select the speaker's moral/epistemic approval function? Or does it select the moral/epistemic approval function of some group that includes the speaker—say, the speaker's linguistic community, or the conversational interlocutors? While these are important questions, for our purposes we can avoid taking a stand on them. A virtue of my approach is that it is compatible with a variety of implementations, contextualist, expressivist, and relativist.[18]

4.2 Beyond modals

So far I have focused on moral and epistemic evaluations made using deontic modals. However, my framework can be extended to encompass other normative expressions.

Earlier we noted that *justification* talk is used to make both moral and epistemic appraisals. Recall (9), repeated here as (11):

[17] Is every deontic ordering source supplied by a desire-like state of mind? Presumably not: we can use deontic modals to convey what is required by a code of etiquette, or the rules of some arcane game, without having any *pro tanto* desire for one to act accordingly. But we might hypothesize that whenever a deontic modal is genuinely normative, the ordering source reflects some desire-like attitudes.

[18] For some considerations bearing on the choice between these frameworks, see Beddor 2019c.

(11) a. You were justified in acting as you did.
b. Poirot is justified in believing the butler did it.

From the perspective of this paper, a natural option is to analyze *justified* in deontic terms.[19] According to what is perhaps the most popular version of this approach, justification is a type of permission: to say someone is justified in doing (or believing) such-and-such is to say that they are permitted to do (or believe) such-and-such. Given the assumption that expressions of permission are the duals of necessity modals, we get the following analysis:

DEONTOLOGICAL *JUSTIFIED* $[\![\text{S is justified in } \psi\text{ing}]\!]^{c,f,g,w} = [\![\text{S is permitted to } \psi]\!]^{c,f,g,w} =$ 1 iff $\exists v \in \textsc{Best}_{g(w),f(w)}$: S ψs at v.

Two advantages of this approach are worth highlighting. First, it immediately validates some plausible entailments between *justified* and deontic modals, e.g.:

(12) a. Given his evidence, Poirot ought to believe the butler did it. ⇒
b. Given his evidence, Poirot is justified in believing the butler did it.

(13) a. You weren't justified in acting as you did. ⇒
b. You shouldn't have acted as you did.

Second, this analysis allows us to capture the moral and epistemic readings of justification ascriptions using the same resources that we used to capture the moral and epistemic readings of deontic modals. As with deontic modals, justification ascriptions are not ambiguous; we have provided a uniform semantic entry for all uses of *justified*, both moral and epistemic. As with deontic modals, different readings of justification ascriptions are captured using different ordering sources. Moral justification ascriptions use a moral ordering source, which is a moral approval profile. Epistemic justification ascriptions use an epistemic ordering source, which is an epistemic approval profile. Thus (11a) communicates that in at least one of the accessible worlds that comes closest to conforming to the relevant moral approval profile, the addressee acted as they did at the actual world. And (11b) communicates that in at least one of the accessible worlds that comes closest to conforming the the relevant epistemic approval profile, Poirot believes the butler did it.[20]

[19] By now there's a large literature on deontological approaches to justification. See e.g. Alston 1988; Plantinga 1993: chp.1; the papers in Steup 2001; Littlejohn 2012: chp.1.

[20] Deontological *Justified* glosses over some important complications. One complication comes from the distinction between weak and strong necessity modals. If I say someone *must* or *has to* do something, I make a stronger claim than if I say they *ought* to do it. Drawing this distinction expands the space of options for a deontological account: one could maintain that *justified* is the dual of a strong necessity modal, or one could maintain that it is the dual of a weak necessity modal. For our purposes, we can set aside this issue; for discussion, see Beddor 2017; for general discussion of the distinction between weak and strong necessity, see, a.o., Sloman 1970; Horn 1989; McNamara 1996; von Fintel and Iatridou 2008. Another complication comes from the gradability of *justified*: we are happy to say things like, *A's belief is more justified than B's belief.*

4.3 Taking stock and looking ahead

Our main rationale for developing a unified treatment moral and epistemic evaluations was to explain their commonalities. How does the framework developed here fare on this front?

The explanation of the linguistic commonalities falls out immediately from our semantics. To recap: our semantics for deontic modals is just a specific implementation of Kratzer's semantics (§4.1). One of signal achievements of Kratzer's semantics is that it accounts for the different readings of modals without positing rampant ambiguity. Given a deontological analysis of justification ascriptions (§4.2), we can extend this style of explanation to capture the different readings of justification ascriptions.

The explanation of the shared connection with the reactive attitudes is also straightforward, at least if we accept some version of the functional role strategy (§3.3). On the framework developed here, moral evaluations voice moral approval. By the functional role strategy, moral approval is characterized in terms of its connection with moral reactive attitudes. Epistemic evaluations voice epistemic approval, which is characterized in terms of its connection with epistemic reactive attitudes.

The explanation of the remaining two commonalities is a little less straightforward, so I will devote a section to each. The next section considers the shared connection with motivation; §6 revisits descriptively enlightened disagreements.

5 Explaining the motivational commonalities

5.1 The general idea

One of the most commonly cited reasons for embracing an attitudinal metaethics is to explain the connection between moral judgment and motivation (MJC). The details of the explanation will depend on the details of one's attitudinal metaethics. However, one common version of the explanation goes like this. For A to believe that they morally ought to ϕ is for A to be in a desire-like state towards the proposition: A ϕs. Now, it's part of the functional role of desire that someone who desires p will be at least somewhat disposed to bring about p. From these two ingredients, it follows that if A believes they morally ought to ϕ, A will be somewhat disposed to ϕ.[21]

One way to accommodate this gradability is to offer a scalar semantics for *justified* and classify *justified* as an absolute gradable adjective (cf. Hawthorne and Logins 2021). We could then impose constraints that relate degrees of justification to the ordering source. For example, we could propose that S is maximally justified in ϕing iff S ϕs in at least one of the best of worlds in the modal base. For relevant discussion of the connection between gradability and modality, see Lassiter 2016; Portner and Rubinstein 2016.

[21] This style of explanation is championed by expressivists; versions of it can be found in Stevenson 1944; Blackburn 1998; Gibbard 1990, 2003, among others. A different style of explanation takes moral beliefs to be representational states of mind *about* one's desire-like states of mind. For example, Egan 2012 defends a relativist view on which moral beliefs are *de se* beliefs about the conditions under which one would have certain desires. According to Egan, this view also can be used to underwrite MJC.

By going in for a generalized attitudinal semantics, we can generalize this style of explanation. On the view that emerges, every *ought*-belief consists in some desire-like state of mind. When it comes to moral *ought*s, this state is moral approval; when it comes to prudential *ought*s, this state is prudential approval; when it comes to epistemic *ought*s, this state is epistemic approval. While there are important differences between these varieties of approval, they share a common motivational core. So whenever someone believes that they (morally, prudentially, epistemically) ought to ϕ, they will be in a state whose functional role disposes them to ϕ. The specific connections between moral judgment and motivation (MJC) and epistemic judgment and motivation (EJC) are just special instances of this more general explanation.

Having laid out the basic idea, the rest of this section defends this style of explanation in response to two concerns.

5.2 Differences between moral and epistemic motivation?

A first objection is that our explanation for MJC does not actually carry over to explain EJC. This objection has been pressed forcefully by Boult and Köhler 2020. First, a quick reminder of what EJC claims: whenever someone believes that they epistemically ought to ϕ, they will be somewhat disposed to ϕ. As we noted in §2.2, here the variable ϕ ranges over both actions (say, engaging in inquiry) and doxastic attitudes.

On the story sketched above (§5.1), moral *ought*-beliefs are constituted by desire-like states. So to explain MJC we need only appeal to the familiar fact that desires motivate action. Does this explanation carry over to explain EJC? When EJC concerns dispositions to act, the answer is obviously 'Yes.' But when EJC concerns our dispositions to form and retain doxastic attitudes, this is debatable. As Boult and Köhler 2020 emphasize, it's not clear that we have a grip on how desires cause doxastic states.

To illustrate the concern, consider a concrete case. Suppose Poirot believes, *I epistemically ought to believe that the butler is guilty.* According to our proposal, this belief is constituted by his epistemic approval of believing the butler is guilty. This epistemic approval disposes Poirot to believe, *the butler is guilty* (call this proposition 'b'). But how does it do so? The worry is that we lack an account of the mechanism whereby this desire-like state disposes our mustachioed detective to believe b.

While this is an important challenge, I think it can be met. Return to our account of epistemic approval (§3). Suppose the basis strategy is on the right track, at least as a partial explication of what distinguishes epistemic from moral approval. On this view, Poirot's epistemic approval consists in a disposition to all-things-considered approve of believing b—a disposition that is grounded in a desire to attain truth and avoid error. Presumably, then, there are some considerations in virtue of which Poirot thinks that believing b will be conducive to the goal of attaining truth and avoiding error. What might these considerations be? The answer depends on the details of his situation. But on a plausible way of filling in the details, the relevant considerations are just parts of his evidence which speak in favor of the butler's guilt—the witness testimony, the butler's

motive and flimsy alibi, etc. If these considerations are strong enough to lead Poirot to think that believing b will be conducive to the goal of attaining truth, presumably they will also be strong enough to generate at least some *pro tanto* disposition for Poirot to believe b.

Some might worry that this explanation renders epistemic approval explanatorily idle. If Poirot's evidence is sufficient to explain his disposition to believe b, why do we also need to invoke his epistemic approval of so believing? However, this objection is too quick, for a couple of reasons.

First, just because Poirot's evidence supports believing b does not mean that Poirot will be disposed to believe b on this basis. After all, Poirot might not pay attention to his evidence, or recognize that his evidence counts in favor of believing b. Here's where epistemic approval comes in. As various philosophers have noted, part of the functional role of desire is to direct our attention to considerations that we take to count in favor of the desire's content.[22] If Ari desires to ski in Chamonix this winter, this desire will dispose her to attend to the considerations that count in favor of skiing in Chamonix this winter—the wide open runs, the stunning views of the Alps, etc. Since epistemic approval is a type of desire, we should expect it to play a similar attention-directing role. Thus Poirot's epistemic approval of believing b will dispose him to attend to those features of the evidence that count in favor of believing b. Attending to these features will, in turn, dispose him to form and retain a belief in b.[23]

Some desire-like states have a further connection with attention: they dispose us to avoid considering alternative possibilities. According to Bratman 1987, this is one of the hallmarks of intention. To illustrate with Bratman's example, suppose someone intends to spend the afternoon at the library. According to Bratman, they have thereby settled for themselves the question, *What to do this afternoon?* Consequently, they will be disposed to avoid seriously considering alternative answers to this question—say, spending the afternoon at the movies instead (1987: 18–19). Suppose we accept this view of intention. Now, so far we have been using the notion of 'approval' as a placeholder, remaining noncommittal on exactly what sort of conative state is involved. Suppose that approval is a species of intention, or at least that it is intention-like in this respect.[24] Then if Poirot

[22]See esp. Scanlon 1998; for related ideas, see Sinhababu 2013, 2017.

[23]Here's a potential worry. It's generally agreed that one cannot form a belief merely because one desires to do so. A familiar style of example: someone offers Beth a million dollars to believe she is Julius Caesar. Even if Beth cares about money more than truth, it seems she cannot believe she is Caesar on this basis. On my account, won't her desire to believe she is Caesar direct her attention to the pecuniary considerations in favor of so believing? If so, why doesn't attending to these considerations dispose Beth to believe she is Caesar? In response, we can borrow a move from the literature on practical reasons for belief. Many philosophers have taken this sort of case to reveal something important about the nature of belief: it's part of the functional role of belief that beliefs can only be held on the basis of considerations that the agent takes to provide evidence in favor of the belief's content (e.g., Adler 2002; Shah 2006; Nolfi 2018). This would explain why attending to the practical considerations favoring a belief cannot generate a disposition to believe, but attending to the epistemic considerations can generate such a disposition.

[24]The hypothesis that moral and epistemic approval are intention-like can be motivated on independent grounds. One challenge for an attitudinal semantics is to explain why normative belief and descriptive belief

epistemically approves of believing b, he will take himself to have settled the question, *Whether to believe* b*?* As a result, he will be disposed to avoid seriously considering other hypotheses about who committed the crime. And this disposition will causally contribute to him forming and retaining a belief in b.[25]

Taking stock: I've suggested two mechanisms whereby epistemic approval can lead one to form and retain some doxastic attitude. First, epistemic approval of holding some doxastic state s will dispose one to attend to the considerations in virtue of which one thinks that holding s is accuracy-conducive. Attending to these considerations will typically generate some disposition to hold s. Second, if epistemic approval is intention-like, then epistemically approving of holding s will dispose one to avoid seriously considering adopting incompatible doxastic states, which will increase the probability that one adopts (or retains) s. If epistemic judgments are constituted by states of epistemic approval, we can appeal to these mechanisms to explain EJC—an explanation that closely parallels our account of the way moral judgment motivates action.[26]

5.3 Approval and normative belief: Explaining the connection

A different worry for my explanation of the motivational commonalities is that I have not yet explained how some form of approval constitutes a normative belief. I've said that states of moral approval constitute *ought*-beliefs, and that states of epistemic approval constitute epistemic *ought*-beliefs. But presumably a complete explanation will not just stipulate this constitution thesis. It will explain why this thesis holds.

In response, we should start by noting that all attitudinal metaethicists already face this challenge—at least, all attitudinal metaethicists who maintain that normative beliefs are constituted by desire-like attitudes. Nothing in the challenge specifically targets the unified framework developed here.

have so much in common. For example, they are subject to the same coherence constraints; they also display functional similarities. As Schroeder notes, this is *prima facie* surprising if normative belief has a very different functional profile than descriptive belief (2010: 96–97). One response is to identify normative belief with a special type of desire-like state—one that is belief-like in certain respects. Intention is a promising candidate for such a state (cf. Gibbard 2003). After all, intention is subject to similar coherence constraints: just as it is irrational to simultaneously believe p and believe $\neg p$, it's irrational to simultaneously intend to ϕ and intend to not ϕ. If Bratman is right, intention also settles questions in much the same manner as belief. For further discussion of how these analogies can be used to defend a noncognitivist view, see Beddor 2020. For a different approach to this 'Commonality Challenge', see Beddor 2019b.

[25]To put it another way, the suggestion is that epistemically approving of holding some belief will help the agent achieve 'cognitive closure' on the relevant question, allowing them to redirect their cognitive resources to other questions. For psychological research on cognitive closure, see e.g., Kruglanski 1989; Kruglanski et al. 1993. For philosophical discussion of this research, see e.g., Nagel 2008; Weisberg 2020.

[26]Some might worry that this discussion assumes that states of epistemic approval always *precede* the formation of the corresponding beliefs. But, the objection runs, this is surely too strong: often we come to hold some doxastic state at the very same time that we come to approve of doing so. Fortunately the account developed here is not committed to this assumption. Frequently, the evidence that counts in favor of holding some doxastic state s is sufficient to cause us to hold s; at the same time, this evidence also causes us to epistemically approve of doing so. But note that even in such cases, the our epistemic approval of holding s may causally contribute to our *retention* of s, thanks to the two mechanisms outlined above.

That said, I want to suggest one way of answering the challenge. The strategy is to provide a semantics for belief reports that, when combined with our attitudinal semantics for normative language, predicts that all species of normative belief are constituted by a corresponding variety of approval. Here I'll sketch one implementation of this strategy, building on ideas in Beddor 2021. For a somewhat different attitudinal semantics for belief reports, see Yalcin 2012. (Readers uninterested in the formal nuts-and-bolts can skip ahead without loss.)

In the semantics literature, belief reports are usually analyzed in modal terms. Specifically, *believes* universally quantifies over the believer's doxastic alternatives—that is, all of the worlds consistent with what they believe (Hintikka 1962). Formally:

HINTIKKA *BELIEVES* $[\![A \text{ believes } \phi]\!]^{c,f,g,w} = 1$ iff $\forall v \in Dox_A^w : [\![\phi]\!]^{c,f,g,v} = 1$,
where $Dox_A^w = \{v \mid v \text{ is compatible with what } A \text{ believes at } w\}$.

In order to validate our constitution thesis (i.e., that normative beliefs are constituted by species of approval), I propose a simple twist on this semantics. Some notation: let m_w^A be A's world-indexed moral approval function at w. This is a constant function from an arbitrary world v to a set of propositions representing the things that A morally approves of at w. Then as a first pass we could propose that *believes* shifts the ordering source in the index to the believer's world-indexed moral approval function:

SHIFTY *BELIEVES* (FIRST PASS) $[\![A \text{ believes } \phi]\!]^{c,f,g,w} = 1$ iff $\forall v \in Dox_A^w : [\![\phi]\!]^{c,f,m_w^A,v} = 1$.[27]

This revision of a Hintikka semantics is conservative, in the sense that it equivalent to Hintikka *Believes* when the complement clause of the belief report is descriptive. The differences only surface when the complement clause contains normative vocabulary. Given our treatment of deontic modals, this semantics predicts that moral *ought*-beliefs are constituted by desire-like states. Consider a belief report such as:

(14) Ari believes she (morally) ought to worship God.

According to Shifty *Believes* (First Pass), (14) is true at w if and only if, for every world v consistent with Ari's beliefs at w, all of the v-accessible worlds that are ranked highest by Ari's w-indexed moral approval profile are worlds where Ari worships God. So (14)

[27] The idea that belief reports shift the values of parameters in the index has been explored independently. Usually such shiftiness is explored in connection with epistemic modals; see e.g., Yalcin 2007; Stephenson 2007; Hacquard 2010; Silk 2017; Ninan 2018 for views on which *believes* shifts the modal base or information state parameter. (Ninan's view is a particularly close formal analogue of the semantics developed here.) See Yalcin 2012 for a view on which *believes* also shifts a hyperplan parameter—a parameter that, on Yalcin's framework, plays a role analogous to an ordering source.

is true if and only if Ari stands in a particular desire-like relation to the proposition that she worships God.[28]

Shifty *Believes* (First Pass) is tailored to deal with moral beliefs; it says nothing about other varieties of normativity. However, generalizing this semantics is straightforward. On our Kratzerian semantics, the ordering source g supplies either a moral, epistemic, or prudential approval function. For any ordering source g, let g_w^A be the set of propositions represents A's corresponding world-indexed approval function at w. Our final semantics for belief reports goes like this:

SHIFTY *BELIEVES* (FINAL) $[\![A \text{ believes } \phi]\!]^{c,f,g,w} = 1$ iff $\forall v \in Dox_A^w : [\![\phi]\!]^{c,f,g_w^A,v} = 1$.

This predicts that every normative belief report ascribes some corresponding species of approval.

Where does this leave us? Our explanation of the motivational commonalities relied on a constitution thesis: every normative belief is constituted by some species of approval (§5.1). Our objector rightly pointed out that we need some explanation of why this constitution thesis holds. Here I've sketched a candidate explanation. The explanation proceeded by way of belief ascriptions: I offered a semantics for *believes* that, when combined with our semantics for *ought*, predicted that every *ought*-belief is constituted by a corresponding flavor of approval.

Some readers might be suspicious of any explanation that proceeds by way of belief *ascriptions*, rather than by way of the functional role of belief. In response, we should make two points. First, this style of explanation is in keeping with the attitudinal metaethicist's *modus operandi*. Suppose you ask an attitudinal metaethicist to explain the nature of moral obligation. The attitudinal metaethicist will not tackle your question directly; rather, they will reframe it as a question about normative language, and proceed to give you a semantic analysis of moral obligation *ascriptions*. Second, there is nothing stopping us from supplementing this approach with a theory of the functional role of belief that—when combined with our attitudinal semantics—delivers the same predictions. For one candidate theory, see Beddor 2019b.

6 Explaining disagreement

6.1 Disagreement in attitude, generalized

Turn now to the final commonality between moral and epistemic evaluations: both give rise to descriptively enlightened disagreements.

[28]Just as our semantics for *ought* was compatible with both contextualism and expressivism, so too is Shifty *Believes*. This has an interesting upshot. Define noncognitivism as the thesis that normative beliefs are desire-like states. It is often thought that noncognitivism is committed to expressivism. The semantics developed here shows that this is mistaken. The semantics developed here entails noncognitivism, while remaining neutral on expressivism. For further discussion of this issue and its metanormative implications, see Beddor 2021.

According to many attitudinal metaethicists, the persistence of descriptively informed moral disagreements shows that not all moral disagreements concern the way the world is. At least some moral disagreements arise from clashes in desire-like states. This notion of 'disagreement in attitude' originates in Stevenson 1944, who offers the following example:

> Suppose that two people have decided to dine together. One suggests a restaurant where there is music; another expresses his disinclination to hear music and suggests some other restaurant. ... The disagreement springs more from divergent preferences than from divergent beliefs, and will end when they both *wish* to go to the same place... (1944: 3)

If moral disagreements have an attitudinal dimension, we can explain why people can agree on the descriptive facts while disagreeing over what ought to be done. After all, people can agree on the descriptive facts while disagreeing in their desire-like attitudes.

To illustrate, go back to our earlier example (§2.3): you present your Ethics 101 class a case where A has the opportunity to kill B to save a hundred. A heated debate ensues. When one student maintains that A morally ought to kill B, they are voicing a consequentialist moral approval function. When another student insists that A is morally prohibited from killing B, they are voicing a deontological moral approval function. These moral approval functions conflict with one another, hence the students disagree in attitude.

Suppose we accept an account along these lines. The framework developed here allows us to generalize this story to explain the persistence of descriptively informed disagreements in other normative domains (cf. Chrisman 2007). When your Epistemology 101 students disagree over whether someone epistemically ought to believe in the external world, their disagreement also springs from conflicting conative attitudes—specifically, conflicting epistemic approval functions. Similar remarks apply to other longstanding epistemological disputes, for example, between those with internalist sympathies and those with externalist commitments. What emerges is a general, unified account of why different varieties of normative discourse lead to descriptively enlightened disagreements.[29]

Let me now consider two objections to this account.

[29] Schafer 2014, 2018 provides a useful foil for my approach. Schafer also advocates a type of attitudinal metaepistemology, inspired by Gibbard 2003. According to Schafer, judgments of epistemic rationality are plans to hold certain beliefs in certain situations. Schafer uses this view to argue for epistemic internalism. If Schafer is right, then properly understanding certain descriptive facts about the nature of epistemic judgment will help settle first-order epistemological disputes. By contrast, my attitudinal metaepistemology is neutral on first-order epistemological questions. According to my diagnosis, the reason why disputes between e.g., internalists and externalists are so intractable is that they no amount of descriptive information will decisively settle the first-order epistemic questions. For further discussion of the plan-based model of epistemic judgment, see the Appendix.

6.2 Too much disagreement?

In one of the few extended discussions of whether an attitudinal semantics can cover the full spectrum of normativity, Wodak 2017 considers a generalized account of disagreement in attitude similar to that sketched above. According to Wodak, this account succumbs to a fatal problem: it overgenerates disagreement.[30] To see the concern, go back to our lonely scholar, who devotes all of their time to research, neglecting family and friends. Suppose A and B our discussing our scholar (C):

(15) A: From the epistemic point of view, C ought to continue devoting all of their time to research.

(16) B: From the moral point of view, C ought not continue devoting all of their time to research.

Intuitively, A and B do not disagree. But, Wodak argues, the disagreement in attitude account predicts they do. After all, A's epistemic approval function and B's moral approval function recommend incompatible courses of action.

This objection is helpful, since it forces us to clarify the conditions under which disagreement in attitude occurs. As the example shows, just because two agents' desire-like attitudes issue incompatible recommendations does not mean that these agents disagree. In order for them to disagree, *the attitudes need to be of the same type.* This requirement explains why A and B do not disagree in the foregoing exchange. By contrast, A and B would disagree if we replaced B's utterance with an epistemic evaluation:

(17) B: From the epistemic point of view, C ought not continue devoting all of their time to research.

Wodak discusses a related solution, which holds that in order for two token attitudes to disagree, they must have the same grounds (2017: 281-282). Wodak objects that this solution still overgenerates disagreement. To illustrate the concern, suppose that A morally disapproves of scientific inquiry into human life extension. The fact that a particular experiment will produce accurate beliefs about life extension might ground A's moral disapproval of conducting the experiment, while simultaneously grounding B's epistemic approval of conducting the experiment. It does not follow that A and B disagree.

This example would only undermine the solution advocated here if we assumed a particular version of the basis strategy, according to which two token attitudes of approval belong to the same type provided they are held on the same basis. But this way individuating types of approval is implausible, as shown by the example under discussion. (A's token attitude of moral disapproval and B's token attitude of epistemic approval have the same basis, but these attitudes are not of the same type.) The example poses no trouble for the view that types of approval are individuated by the bases that normally give rise

[30] In formulating the problem, Wodak focuses on the relation between moral *ought*s and legal *ought*s. But Wodak's points easily extend to the interaction between moral *ought*s and epistemic *ought*s, which is our focus here.

to them in the general population (§3.2). It also causes no trouble for a view on which we individuate varieties of approval by their functional roles (§3.3). By adopting either of these accounts, we avoid overgenerating normative disagreements.

6.3 Towards a general account of disagreement ascriptions

A further worry remains: I seem to be stipulating that certain combinations of attitudes constitute disagreements. But am I entitled to do so? Shouldn't a complete theory tell us why these combinations of attitudes deserve to be called 'disagreements'?

Now, this is really a problem for the very idea of 'disagreement in attitude.' It is not specifically a problem for my attempt to generalize the notion of disagreement in attitude beyond the moral domain. That said, I want to briefly sketch one possible solution, building on earlier treatment of normative belief.

Ask your average philosopher, 'What does it take for two people to disagree?' Chances are, they'll say something along the lines of 'Two people disagree when one of them believes something, and the other believes its negation.' This idea has considerable plausibility. Indeed, I suspect many philosophers' resistance to 'disagreement in attitude' springs from this idea that disagreement is always a doxastic phenomenon.

Now, suppose we were to translate this idea into a semantics for disagreement ascriptions. Here is the obvious way to do so:

Disagreement Ascriptions $[\![A \text{ and } B \text{ disagree over whether } \phi]\!]^{c,f,g,w} = 1$ iff both:

1. $[\![\alpha \text{ believes } \phi]\!]^{c,f,g,w} = 1$,
2. $[\![\beta \text{ believes } \neg\phi]\!]^{c,f,g,w} = 1$,

or *vice versa*.[31]

This analysis is perfectly general: it applies to all disagreements, both descriptive and normative.

The next step is to observe that Disagreement Ascriptions is actually compatible with an attitudinal semantics. After all, attitudinal semanticists need an analysis of belief reports. §5.3 offered one analysis on their behalf. Whatever analysis they end up giving, they are free to combine it with Disagreement Ascriptions. This will yield an account of normative disagreement as a special case.

It may help to walk through an example. Suppose we want to analyze:

(18) A and B disagree over whether C ought to devote all their time to research.

Since (18) contains a deontic modal, its interpretation will depend on the ordering source parameter g. According to Disagreement Ascriptions, (18) will be true relative to g provided the corresponding belief reports are both true, as evaluated relative to g:

[31] I defend this semantics for disagreement ascriptions in Beddor 2019a, where I discuss how it applies to disputes involving taste predicates and epistemic modals.

(19) A believes C ought to devote all their time to research.

(20) B believes it's not the case that C ought to devote all their time to research.

For example, suppose we are evaluating (18)—and consequently (19) and (20)—using a moral ordering source m. By our semantics for *believes* (§5.3), (19) is true, relative to m, if and only if A is in a particular state of moral approval. (More precisely, this reading of (19) is true just in case for every world v in A's doxastic alternatives, all of the v-accessible worlds ranked highest by A's moral approval function are worlds where C devotes all of their time to research.) And (20) is true, relative to m, just in case B has a conflicting state of moral approval. (More precisely, this reading of (20) is true just in case for every world v in B's doxastic alternatives, not all of the v-accessible worlds ranked highest by B's moral approval function are worlds where C devotes all of their time to research.) Now, suppose instead we evaluate (18)—and consequently (19) and (20)—using an epistemic ordering source. On this reading, (18) is true just in case A and B have conflicting epistemic approval functions.

This implementation respects the idea that normative disagreement arises from conflicting desire-like attitudes, but only when the conflicting attitudes are of the same type. It does not merely stipulate this. Rather, we have derived this from a general semantics for disagreement ascriptions, together with our semantics for *believes*.[32]

7 Bonus: defusing Woods' challenge

I've argued that we should distinguish between different varieties of approval. Once we do, we can enlist these varieties to serve as ordering sources for different sorts of normative evaluations. This enables us to explain the commonalities between moral and epistemic evaluations, while also accounting for their differences. But it also carries additional benefits. In this section, I advertise one such benefit: it solves an important challenge to an attitudinal metaethics, due to Woods 2014.

7.1 The challenge

Woods starts by reminding us of Moore-paradoxical sentences such as:

(21) ?? It's raining, but I don't believe it's raining.

A standard diagnosis of the infelicity of (21) goes like this. In uttering the first conjunction (*It's raining*), the speaker expresses a belief that it's raining. But the second

[32]While Disagreement Ascriptions plays nicely with our semantics for *believes* from §5.3, it is not wedded to this particular semantics, or even to the idea of disagreement in attitude. As noted in fn.21, some proponents of an attitudinal semantics take normative beliefs to be representational beliefs about one's desire-like attitudes. Disagreement Ascriptions is perfectly compatible with such an approach. For further discussion, see Beddor 2019a.

conjunct goes on to deny that the speaker has any such belief. More generally, if an utterance expresses some mental state m, it is incoherent for the speaker to deny that they are in m.

Now for the problem. Woods observes that it sounds fine to say:

(22) Eating meat is wrong, but I don't disapprove of it.

But, Woods argues, if expressivism is true, (22) has the same defect as (21). By the expressivist's lights, the first conjunct (*Eating meat is wrong*) expresses disapproval of eating meat. The second conjunct denies that the speaker disapproves of eating meat. So the speaker expresses a mental state while disavowing it, exactly as in (21).

Before offering a solution, two preliminary remarks are in order. First, Woods' example involves *wrongness*-talk, whereas we have focused on deontic modals. This difference is unimportant, since we can easily adapt his example:

(23) Morally speaking, we ought not eat meat. But I don't disapprove of eating meat.

The other point is that Woods focuses on a specific version of an attitudinal semantics: expressivism. As noted in §4, the framework developed in this paper can be implemented in a contextualist, expressivist, or relativist setting. So some might think that we could escape the challenge by retreating to a contextualist or relativist implementation. However, I won't pursue this escape route for a couple of reasons. First, Woods' challenge extends to standard versions of contextualism and relativism. On many contextualist and relativist views, the first conjunct of (22)/(23) commits the speaker to some form of disapproval of meat-eating. Second, as we will see momentarily, the framework developed here allows us to give a more direct response to Woods' challenge—a response that is perfectly compatible with expressivism.

7.2 The solution

According to the view put forward here, different normative evaluations voice different types of approval. This yields a simple solution to Woods' challenge.

On my view, the first conjunct of (22)/(23) does not voice just any old state of disapproval. It voices moral disapproval. Someone can morally disapprove of eating meat, even though they do not all-things-considered disapprove of it. How should we understand this state of mind? The answer depends on our preferred solution to the Individuation Problem. If we adopt the basis strategy (§3.2), this agent is disposed to all-things-considered disapprove of eating meat—a disposition that is held on the basis of moral considerations (e.g., the wellbeing of sentient creatures). It's just that this disposition is masked by countervailing considerations (e.g., taste and convenience). If we adopt the functional role strategy (§3.3), the story is similar. Our agent is disposed to blame meat-eaters. It's just that this disposition does not lead them to all-things-considered disapprove of eating meat, perhaps because it is outweighed by the sorts of considerations mentioned above. However we spell out the details, this mental state is coherent.

Some might wonder whether there is a way of reviving Woods' challenge. Can't we modify the second conjunct of (22)/(23) so that the speaker denies *morally disapproving* of eating meat? Indeed we can. But doing so degrades the coherence of these sentences:

(24) ?? Eating meat is wrong, but I don't morally disapprove of eating meat.

(25) ?? Morally speaking, we ought not eat meat. But from the moral point of view, I don't disapprove of doing so.

These conjunctions strike me as absurd. Indeed, they seem just as absurd as canonical Moore-paradoxical sentences. So far from being a problem for my proposed solution, these data provide strong reason to think that it is on the right track.[33]

8 Conclusion

A long-standing tradition in metaethics takes moral discourse to communicate desire-like attitudes. But moral discourse is just one species of normative language. Moreover, the different varieties of normative language have much in common. These commonalities motivate the development of a unified attitudinal semantics.

According to the framework developed here, different flavors of normative language voice different varieties of approval. These varieties are unified by a functional core, allowing us to explain the commonalities between the different types of normative judgment. But there are also important differences between the varieties of approval, allowing us to explain the differences between moral, epistemic, and prudential evaluations.

The resulting view is thoroughly naturalistic. It is also compositionally serious, allowing us to take on board the leading linguistic analysis of deontic modals. And it delivers downstream benefits. For example, it yields a simple solution to Wood's 2014 objection to expressivism.

Appendix: Comparison With Gibbard

I am not the first to suggest that we should extend an attitudinal semantics to epistemic discourse.[34] However, my framework differs from other unificationist proposals in important respects. In this appendix, I compare my approach to one particularly influential unificationist strategy, due to Gibbard 2003.

On Gibbard's view, all normative judgments are plan-laden states of mind, where a plan is a conative attitude akin to intention. While Gibbard primarily focuses on practical normativity, he also devotes a chapter to knowledge (2003: chp.11). According to Gibbard, for A to judge that B knows whether p is for A to plan to defer to B's judgment about p in similar situations (2003: 228–229).[35]

[33] For other responses to Woods' challenge, see Toppinen 2014; Raskoff 2018; Franzén 2019.

[34] For references, see fn. 1.

[35] See Schafer 2014, 2018; Greco and Hedden 2016 for related analyses of rationality judgments.

Gibbard's account has some obvious affinities with my approach: we both take epistemic evaluations to voice conative attitudes. But there are also major differences. One difference concerns the level of semantic detail. Gibbard does not develop a precise semantics for knowledge ascriptions; by his own admission, his account is "rough and inexact" (2003: 229). By contrast, one of my main goals was to provide a precise compositional analysis of my target constructions.

There are also more principled differences between our accounts. On Gibbard's account, epistemic judgments are what we might call 'deference plans': plans to defer to someone's judgment. For an initial counterexample to this view, suppose Tess the teacher is grading Stu the student's math quiz. Tess judges that Stu knows the answer to a particular arithmetic problem. But Tess does not plan to defer to Stu's arithmetic judgments, since she takes herself to be a better judge of such matters than Stu (Schafer 2014: 2576).

It may seem that there is an obvious reply. We should think of Tess' deference plan as a contingency plan: she is planning to defer to Stu's judgment in counterfactual situations where she is less informed (Gibbard 2003: 228). However, let's add a wrinkle to the case. Suppose Tess is a 'proof chauvinist' (Barnett forthcoming): she thinks one is only justified in holding a mathematical belief if one has worked out a proof of that belief on one's own. Consequently, her contingency plan for her less informed counterfactual self is to work out the proof rather than rely on Stu. So Tess judges that Stu knows the answer, even though she plans not to defer to his judgment in situations where she is less knowledgeable. On Gibbard's view, Tess' mental state is incoherent, which seems wrong. My view has no trouble here. Tess can epistemically approve of Stu's belief without epistemically approving of deferring to his belief.

A related problem for Gibbard's account comes from our epistemic self-evaluations. I judge that I am sitting in front of a computer; I also judge that I know that I am sitting in front of a computer. On Gibbard's account, the latter judgment consists in a plan to defer to myself when I make the former judgment. This seems wrong. In order to defer to someone's judgment that p, I must judge p on the basis of my belief that they judge p. So to defer to my own judgment about p would be to judge p on the basis of my belief that I judge p. This is a very odd state to be in, and an even odder state to plan to be in. My view avoids this problem. I can epistemically approve of my current beliefs without planning to form those beliefs on the basis of the fact that I hold them.

A final difference is that Gibbard takes all normative judgments to involve a single conative attitude, *planning*. Different types of plans are distinguished only by their objects—i.e., what they are plans to do. Epistemic plans are plans to believe in certain ways; moral plans are plans to have certain reactive attitudes. The examples in this paper suggest that this view has trouble distinguishing the full range of normative judgments.

Go back to our reclusive scholar. As we noted in §3, here are three distinct judgments that someone might make about them:

(26) a. They morally ought not continue devoting all their time to research.
b. They epistemically ought to continue devoting all their time to research.

c. They all-things-considered ought not continue devoting all their time to research.

For Gibbard, (26a) expresses a plan to blame the scholar for devoting all their time to research (or a plan to approve of so blaming). While Gibbard does not explicitly discuss (26b), a natural way of applying his view would be to say it expresses a plan to continue devoting all one's time to research, in the contingency of being in the scholar's shoes. This already seems rather implausible: it means that anyone who judges both (26a) and (26b) is planning to ϕ in a particular contingency, and also planning to blame themselves for ϕ-ing in that contingency—which seems wrong. And things get worse when we ask: what does (26c) express? One possibility is that it expresses the same plan as (26a). But then we fail to distinguish between these two judgments. Another possibility is that it expresses a plan to **not** devote all of one's time to research, in the contingency of being in the scholar's shoes. But then we incorrectly predict that the mental state expressed by (26b) disagrees with the mental state expressed by (26c).[36]

From the perspective of this paper, the solution is to distinguish between different types of pro-attitude: epistemic, moral, and all-things-considered. These attitudes are not distinguished by their objects, but rather by their bases or functional roles. Once we take this line, we correctly predict that (26b)-(26c) voice distinct but consistent judgments.[37,38]

References

Jonathan Adler. *Belief's Own Ethics.* MIT Press, Cambridge, MA, 2002.

William Alston. The deontological conception of epistemic justification. *Philosophical Perspectives*, 2:115–152, 1988.

Zachary Barnett. Rational moral ignorance. *Philosophy and Phenomenological Research*, forthcoming.

Bob Beddor. *Reduction in Epistemology.* PhD thesis, Rutgers University, 2016.

Bob Beddor. Justification as Faultlessness. *Philosophical Studies*, 174:901–926, 2017.

Bob Beddor. Subjective Disagreement. *Noûs*, 53(4):819–851, 2019a.

Bob Beddor. Noncognitivism and epistemic evaluations. *Philosophers' Imprint*, 19(10):1–27, 2019b.

[36]This is an intrapersonal instance of the problem of overgenerating disagreement (§6.2). See Wodak 2017, §11 for related concerns about Gibbard's view. It is worth noting that this problem also arises for a number of views in the ballpark. For example, Schafer 2014 criticizes the details of Gibbard's account, and proposes an alternative plan-based metaepistemology that avoids some of the problems raised earlier. But like Gibbard, Schafer does not distinguish epistemic plans from moral plans, except by their objects. So his approach is also susceptible to a version of this challenge. Similar remarks apply to the plan-based metaepistemology defended in Greco and Hedden 2016.

[37]As we've seen in §3, this is compatible with Gibbard's insight that epistemic judgments are closely tied to belief-regulation and that moral judgments are closely tied to reactive attitudes. According to the functional role strategy, moral approval is a state whose functional role makes reference to moral blame. Epistemic approval is a state whose functional role makes reference to epistemic blame, which can be understood in terms of its connections with habits of belief formation.

[38]An early form of some of the ideas in this paper appeared in chp.2 of my dissertation (Beddor 2016). For helpful feedback on that material, I am grateful to Andy Egan, Alvin Goldman, and participants in the Rutgers dissertation seminar. Thanks also to Cameron Boult, Simon Goldstein, and Daniel Wodak for detailed comments on an earlier draft of this paper.

Bob Beddor. Relativism and expressivism. In Kusch, editor, *Routledge Handbook to the Philosophy of Relativism*, pages 528–540. Routledge, 2019c.

Bob Beddor. Fallibility for Expressivists. *Australasian Journal of Philosophy*, 98(4):767–777, 2020.

Bob Beddor. Noncognitivism without expressivism. Manuscript, 2021.

Gunnar Björnsson and Tristram McPherson. Moral Attitudes for Non-Cognitivists: Solving the Specification Problem. *Mind*, 123(489):1–38, 2014.

Simon Blackburn. *Ruling Passions*. Clarendon Press, Oxford, 1998.

Cameron Boult. There is a distinctively epistemic kind of blame. *Philosophy and Phenomenological Research*, forthcoming.

Cameron Boult and Sebastian Köhler. Epistemic judgement and motivation. *Philosophical Quarterly*, 70(281): 738–758, 2020.

Michael Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

Jessica Brown. What is epistemic blame? *Noûs*, 54(2):389–407, 2020.

Matthew Chrisman. From epistemic contextualism to epistemic expressivism. *Philosophical Studies*, 135(2): 225–254, 2007.

Sinan Dogramaci. Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84 (3):513–530, 2012.

James Dreier. Internalism and Speaker Relativism. *Ethics*, 101(1):6–26, 1990.

Andy Egan. Relativist Dispositional Theories of Value. *Southern Journal of Philosophy*, 50(4):557–582, 2012.

Hartry Field. Epistemology without metaphysics. *Philosophical Studies*, 143(2):249–290, 2009.

Nils Franzén. Evaluative discourse and affective states of mind. *Mind*, 129(516):1095–1126, 2019.

Allan Gibbard. *Wise Choices, Apt Feelings*. Harvard University Press, Cambridge, MA, 1990.

Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, MA, 2003.

Daniel Greco. A puzzle about epistemic akrasia. *Philosophical Studies*, 167(2):201–219, 2014.

Daniel Greco. Epistemological open questions. *Australasian Journal of Philosophy*, 93(3):509–523, 2015.

Daniel Greco. Justification and Excuses in Epistemology. *Noûs*, forthcoming.

Daniel Greco and Brian Hedden. Uniqueness and metaepistemology. *Journal of Philosophy*, 113(8):365–395, 2016.

Valentine Hacquard. On the Event Relativity of Modal Auxiliaries. *Natural Language Semantics*, 18(1):79–114, 2010.

John Hawthorne and Arturs Logins. Graded epistemic justification. *Philosophical Studies*, 178:1845–1858, 2021.

John Hawthorne and Amia Srinivasan. Disagreement without transparency: Some bleak thoughts. In Christensen and Lackey, editors, *Epistemology of Disagreement*, pages 9–30. Oxford University Press, Oxford, 2013.

Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.

Laurence Horn. *A Natural History of Negation*. University of Chicago Press, Chicago, 1989.

Klemens Kappel and Emil F.L. Moeller. Epistemic expressivism and the argument from motivation. *Synthese*, 191:1529–1547, 2014.

Antti Kauppinen. What Makes a Sentiment Moral? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1. Oxford University Press, New York, 2010.

Antti Kauppinen. Epistemic norms and epistemic accountability. *Philosophers' Imprint*, 18, 2018.

Christoph Kelp and Mona Simion. Criticism and Blame in Action and Assertion. *Journal of Philosophy*, 114 (2):76–93, 2017.

Sebastian Köhler. Do Expressivists Have an Attitude Problem? *Ethics*, 123(3):479–507, 2013.

Christine Korsgaard. *Creating the Kingdom of Ends*. Cambridge University Press, Cambridge, 1996.

Angelika Kratzer. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1:337–355, 1977.

Angelika Kratzer. The Notional Category of Modality. In Eikmeyer and Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. W. de Gruyter, Berlin, 1981.

Angelika Kratzer. Modality. In von Stechow and Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*. W. de Gruyter, Berlin, 1991.

Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, Oxford, 2012.

Arie Kruglanski. *Lay Epistemics and Human Knowledge.* Springer Science & Business Media, 1989.

Arie Kruglanski, Donna M. Webster, and Adena Klem. Motivated resistance and openness to persuasion in the presence or absence of prior information. *Journal of Personality and Social Psychology*, 65(5):861–876, 1993.

Daniel Lassiter. *Must*, Knowledge, and (In)directness. *Natural Language Semantics*, 24:117–163, 2016.

Stephanie Leary. Defending internalists from acquired sociopaths. *Philosophical Psychology*, 7:878–895, 2017.

Clayton Littlejohn. *Justification and the Truth Connection.* Cambridge University Press, Cambridge, 2012.

B.J.C Madison. On justifications and excuses. *Synthese*, 195(10):4551–4562, 2018.

Paul McNamara. Must I do what I ought (or will the least I can do do)? In Brown and Carmo, editors, *Deontic Logic, Agency, and Normative Systems.* Springer-Verlag, Berlin, 1996.

Alfred Mele. Internalist Moral Cognitivism and Listlessness. *Ethics*, 106:727–753, 1996.

J.S. Mill. *Utilitarianism.* 1861.

Alexander Miller. *Introduction to Contemporary Metaethics.* Polity, Cambridge, 2003.

G.E. Moore. *Principia Ethics.* Cambridge University Press, Cambridge, 1903.

Jennifer Nagel. Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86:279–294, 2008.

Dilip Ninan. Relational Semantics and Domain Semantics for Epistemic Modals. *Journal of Philosophical Logic*, 47(1):1–16, 2018.

Kate Nolfi. Why only evidential considerations can justify belief. In McHugh, Way, and Whiting, editors, *Normativity: Epistemic and Practical.* Oxford University Press, Oxford, 2018.

Adam Pichovy. What do we want from a theory of epistemic blame? *Australasian Journal of Philosophy*, forthcoming.

Alvin Plantinga. *Warrant: The Current Debate.* Oxford University Press, Oxford, 1993.

Paul Portner and Aynat Rubinstein. Extreme and non-extreme deontic modals. In Charlow and Chrisman, editors, *Deontic Modality.* Oxford University Press, Oxford, 2016.

Sarah Zoe Raskoff. Getting expressivism out of the woods. *Ergo*, 5(36):947–969, 2018.

Lindsay Rettler. In defense of doxastic blame. *Synthese*, 195(5):2205–2226, 2018.

Michael Ridge. Expressivism and epistemology: Epistemology for ecumenical expressivists. *Aristotelian Society Supplementary Volume*, 81(1):83–108, 2007.

Michael Ridge. How to be an epistemic expressivist. In McHugh, Way, and Whiting, editors, *Metaepistemology.* Oxford University Press, Oxford, 2018.

A.L. Roskies. Are ethical judgments intrinsically motivational? lessons from acquired sociopathy. *Philosophical Psychology*, 16:51–66, 2003.

T.M. Scanlon. *What We Owe to Each Other.* Harvard University Press, Cambridge, MA, 1998.

T.M. Scanlon. *Moral Dimensions: Permissibility, Meaning, Blame.* Belknap Press, Cambridge, MA, 2008.

Karl Schafer. Doxastic planning and epistemic internalism. *Synthese*, 191:2571–2591, 2014.

Karl Schafer. Epistemic planning, epistemic internalism, and luminosity. In McHugh, Way, and Whiting, editors, *Metaepistemology.* Oxford University Press, Oxford, 2018.

Mark Schroeder. *Being For: Evaluating the Semantic Program of Expressivism.* Oxford University Press, Oxford, 2008.

Mark Schroeder. *Noncognitivism in Ethics.* Routledge, New York, 2010.

Andrew Sepielli. Normative Uncertainty for Non-Cognitivists. *Philosophical Studies*, 160:191–207, 2012.

Nishi Shah. A new argument for evidentialism. *The Philosophical Quarterly*, 56(225):481–498, 2006.

Alex Silk. How to Be an Ethical Expressivist. *Philosophy and Phenomenological Research*, 90(3), 2014.

Alex Silk. How to Embed an Epistemic Modal: Attitude Problems and Other Defects of Character. *Philosophical Studies*, 174:1773–1799, 2017.

Neil Sinhababu. The Desire-Belief Account of Intention Explains Everything. *Noûs*, 47(4):680–696, 2013.

Neil Sinhababu. *Humean Nature.* Oxford University Press, New York, 2017.

Aaron Sloman. 'ought' and 'better'. *Mind*, 79:385–394, 1970.

Michael Smith. *The Moral Problem.* Blackwell, Oxford, 1994.

Tamina Stephenson. Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.

Matthias Steup. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue.* Oxford University Press, Oxford, 2001.

C.L. Stevenson. *Ethics and Language.* Yale University Press, New Haven, 1944.

C.L. Stevenson. *Facts and Values.* Greenwood Press, Westport, CT, 1963.

Michael Stocker. Desiring the Bad: An Essay in Moral Psychology. *The Journal of Philosophy*, 76:738–753, 1979.

P.F. Strawson. Freedom and Resentment. In *Freedom and Resentment and other Essays*, pages 1–25. Methuen, London, 1974.

Jonathan Sutton. *Without Justification.* MIT Press, Cambridge, MA, 2007.

Sigrun Svavarsdóttir. Moral Cognitivism and Motivation. *The Philosophical Review*, 108:161–219, 1999.

Teemu Toppinen. Expressivism and moore's paradox: A reply to woods. *Journal of Ethics and Social Philosophy*, 8(1):1–5, 2014.

Mark van Roojen. Moral Rationalism and Rationalist Amoralism. *Ethics*, 120:495–525, 2010.

Kai von Fintel and Sabine Iatridou. How to say 'ought' in foreign: the composition of weak necessity modals. In Guéoron and Lecarme, editors, *Time and Modality (Studies in Natural Language and Linguistic Theory 75)*. Springer, 2008.

Ralph Wedgwood. *The Nature of Normativity.* Oxford University Press, Oxford, 2007.

Jonathan Weisberg. Belief in psyontology. *Philosophers' Imprint*, 20(11):1–27, 2020.

Timothy Williamson. Justifications, Excuses, and Sceptical Scenarios. In Dorsch and Dutant, editors, *The New Evil Demon Problem.* Oxford University Press, Oxford, forthcoming.

Daniel Wodak. Expressivism and the varieties of normativity. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 12, pages 265–293. Oxford University Press, Oxford, 2017.

Jack Woods. Expressivism and Moore's Paradox. *Philosophers' Imprint*, 14:1–12, 2014.

Seth Yalcin. Epistemic Modals. *Mind*, 116(464):983–1026, 2007.

Seth Yalcin. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012.