

Noncognitivism Without Expressivism

Bob Beddor

October 9, 2022

Forthcoming in *Philosophy and Phenomenological Research*

Abstract

According to expressivists, normative language expresses desire-like states of mind. According to noncognitivists, normative beliefs have a desire-like functional role. What is the relation between these two doctrines? It is widely assumed that expressivism commits you to noncognitivism, and *vice versa*. This paper argues against this assumption. I advance a view that combines a noncognitivist psychology with a descriptivist semantics for normative language. While this might seem like an ungainly hybrid, I argue that it has important advantages over more familiar metaethical positions. The noncognitivist aspect of the theory captures all of the explanatory benefits standardly associated with expressivism. At the same time, the descriptivist element allows us to avoid all of the semantic headaches for expressivism.

1 Stage Setting

One major faultline in metaethics concerns the meaning of moral language. According to descriptivists, someone who asserts ‘Stealing is wrong’ is doing much the same thing as someone who asserts ‘Global temperatures are rising’: both aim to accurately describe the world. Expressivists deny this. According to expressivists, moral language does not aim to describe the world, but rather to express the speaker’s desire-like attitudes.

Another metaethical faultline concerns the nature of moral belief. According to cognitivists, someone who believes that stealing is wrong is in the same sort of psychological state as someone who believes that global temperatures are rising: both are in states that

aim to accurately represent the world. Noncognitivists deny this. According to noncognitivists, to believe that stealing is wrong is to hold some desire-like attitude towards stealing, for example, to prefer that no one steals, or to disapprove of stealing.¹

To hear many metaethicists tell it, these are not two distinct faultlines; it is just the same one described twice. It's usually thought that expressivism goes hand-in-hand with noncognitivism and that descriptivism goes hand-in-hand with cognitivism. Natural as these pairings may be, this paper makes the case for teasing them apart. I advance a novel form of *descriptivist noncognitivism*. This view adopts a noncognitivist moral psychology, according to which having a moral belief is just a matter of being in some conative state. But it combines this with a descriptivist semantics for moral language. I argue that this view gives us the best of both worlds. The noncognitivist element captures all of the selling points associated with expressivism. It explains why moral judgments give rise to the 'Open Question' phenomenon; it accounts for the close connection between moral belief and motivation; and it explains how people who agree on all of the descriptive facts can still have a moral disagreement. At the same time, we retain a simple possible worlds semantics for moral language, without any expressivist bells and whistles. As a result, we sidestep the Frege-Geach Problem altogether. We get the joys of expressivism without the tears.

The idea that we can solve metaethical problems by distinguishing between moral semantics and moral psychology is fairly simple. But it faces an obvious objection. Surely, the objection runs, there is a very close connection between moral language and moral belief. For one thing, we ascribe moral beliefs using moral vocabulary (e.g., 'Mary believes that stealing is wrong'). For another thing, we use moral language to influence the moral beliefs of our interlocutors. Normally, when I assert 'Stealing is wrong', I am trying to get my audience to believe that stealing is wrong. It is difficult to see how we could account for these connections if we divorce moral semantics from moral psychology.²

¹Early noncognitivists denied the existence of moral beliefs (e.g., [Ayer 1936](#): 108; [Russell 1935](#): 231–232). However, more recent noncognitivists have sought to 'save the appearances' of moral cognitivism, allowing that we can speak truly when we ascribe moral beliefs (e.g., [Blackburn 1993, 1998](#); [Gibbard 2003, 2008](#); [Yalcin 2012](#)). For further discussion of how to reconcile noncognitivism with the existence of moral beliefs, see [Ridge 2006a](#); [Schroeder 2010](#): chp.5; [Björnsson and McPherson 2014](#); [Köhler 2017](#); [Beddor 2019b](#).

²To my knowledge, the only author who has defended a version of descriptivist noncognitivism is [Kalderon 2004](#). Kalderon's account is particularly vulnerable to this objection. Kalderon does not provide a semantics for moral discourse or for moral belief reports, with the result that his account leaves a number of pressing questions about the connection between moral language and moral belief unanswered. I compare my approach with Kalderon's in §2.3. Another interesting point of comparison is [Horgan and Timmons 2006](#), who defend a 'cognitivist expressivism', which combines a cognitivist moral psychology with an expressivist semantics. Not everyone has been convinced that Horgan and Timmons' view deserves its title. For example, Schroeder objects that the only cognitivist element of their view is that they develop a semantics for *believes* according to which it can embed moral vocabulary—an idea that is perfectly consistent with noncognitivism ([2009](#): 259). For the purposes of this paper, I won't take a stand on this issue. Regardless of what one thinks of cognitivist expres-

This paper develops a distinctive version of descriptivist noncognitivism that explains these connections. Here’s the basic recipe. The first ingredient is a familiar contextualist semantics for moral discourse, according to which moral sentences are evaluated using a contextually determined parameter that generates an ordering over worlds. The second ingredient is a new proposal about the semantics of moral belief reports: moral belief reports shift the ordering to be provided by the believer’s desire-like states of mind. The resulting view, I argue, qualifies as a form of descriptivist noncognitivism. Moreover, it’s a form that captures the systematic connections between moral semantics and moral psychology.

2 Developing Descriptivist Noncognitivism

This section offers a ‘proof of concept’ of descriptivist noncognitivism. I start by reviewing a familiar descriptivist semantics for moral language. I then show that by combining this semantics with a new twist on a familiar theory of belief reports, we get descriptivist noncognitivism as a result.

2.1 Descriptivist Truth Conditions

For the sake of specificity, I focus on moral evaluations made using modals (*ought*, *should*, *may*). For example:

- (1) People (morally) ought to give to charity.

In order to implement my proof of concept, I will use a well-known semantics for deontic modals, due to [Kratzer 1981, 1991, 2012](#). (I rely on Kratzer’s semantics mainly for familiarity and convenience; as we will see shortly, the basic idea behind my approach can be generalized to a variety of other semantics.) 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. Necessity modals universally quantify over the g -best of the f -accessible worlds:

Kratzerian Ought $[[Ought \phi]]^{f,g,w} = 1$ iff $\forall v \in \text{BEST}_{g(w),f(w)} : [[\phi]]^{f,g,v} = 1$,

where $\text{BEST}_{g(w),f(w)}$ is the set of worlds in $\bigcap f(w)$ ranked highest by $g(w)$.

sivism, I’ll argue that its flipside—descriptivist noncognitivism—not only earns the title; it can be motivated on principled grounds.

A major advantage of Kratzerian *Ought* is that it offers a uniform semantics for all uses of modals. Modal language is extremely flexible: we use the same modal expression (e.g., *ought*) to make moral evaluations (1), to dispense prudential advice (2), to discuss what is required by some contextually salient regulations (3), and so on.

- (2) You ought to take the A train.
- (3) Hospital visitors ought to leave by 8pm.

Kratzerian *Ought* captures this flexibility without positing rampant ambiguity. On the resulting semantics, all modals share a quantificational semantic core. The differences between them correspond to different modal bases or ordering sources. For example, (1) is evaluated using a moral ordering source, which ranks the accessible worlds along a moral dimension. (2) is most naturally evaluated using a teleological ordering source, which ranks the accessible worlds based on the extent to which they comply with the addressee's goals.

As formulated, Kratzerian *Ought* does not tell us where the ordering source or the modal base come from. But one natural answer—explicitly endorsed by [Kratzer 1981, 2012](#)—is that they are fixed by the context of utterance. That is:

Parameter Contextualism Assertions of *Ought* ϕ are evaluated using a modal base and ordering source supplied by the context of utterance.

For the purposes of developing our proof of concept, it will be helpful to assume Parameter Contextualism. We shall consider an alternative in due course (§3.3).

Given Parameter Contextualism, Kratzerian *Ought* has good claim to be considered a descriptivist semantics.³ In order to see why, let's start by sharpening our understanding of the descriptivist/expressivist distinction. Admittedly, this is a tricky task, since different authors use these labels in different ways. So no way of drawing this distinction will win universal agreement. That said, I take it that the core idea behind descriptivism is that moral sentences purport to represent the world in the same way that non-moral sentences do. This suggests the following sufficient condition for a semantics to be descriptivist:

Descriptivism (First Pass) A semantics is descriptivist if it assigns every utterance of a moral sentence representational truth conditions.

³Arguably, Parameter Contextualism is a metasemantic view. So it may be more accurate to view the conjunction of Kratzerian *Ought* and Parameter Contextualism as a combination of a semantic thesis with a metasemantic thesis. For my purposes, little hinges on this choice of labeling.

What does it mean for a sentence to have representational truth conditions? A natural thought is that a sentence has representational truth conditions provided it can be modeled with a classical proposition—i.e., a set of worlds. For example, your standard possible worlds semantics will model the meaning of ‘Grass is green’ with the set of worlds where grass is green. This semantics can be viewed as encoding representational truth conditions. An utterance of ‘Grass is green’ represents the world as being one where grass is green. This representation is true at some world w if grass is green at w , false otherwise.

This suggests a more perspicuous way of unpacking our formulation of descriptivism:

Descriptivism (Second Pass) A semantics $[[\cdot]]$ is descriptivist if, for every moral sentence ϕ and every sequence of contextually determined parameters c , $[[\phi]]^c$ is a set of worlds.

Expressivists reject descriptivism, thus construed. Most expressivists embrace some form of a ‘psychologistic semantics’, according to which the meanings of moral utterances are desire-like states of mind, or some formal object that models such a state.⁴ And a psychologistic semantics is typically regarded as an *alternative* to a truth conditional semantics. Admittedly, not all expressivists agree on this last point. For example, both [Yalcin 2012](#) and [Silk 2014](#), building on ideas in [Gibbard 2003](#), have suggested that expressivists can help themselves to truth conditional resources. Still, even these expressivists agree that moral sentences do not have *representational* truth conditions. For them, an utterance of a moral sentence cannot be evaluated as true or false relative to a world alone, but only relative to an ordered pair of a world and some entity that models a desire-like state of mind.⁵

We are now in a position to see why Kratzerian *Ought*, when paired with Parameter Contextualism, is descriptivist. The crucial point is that while Kratzerian *Ought* postulates non-worldly parameters—an ordering source and a modal base—Parameter Contextualism says that these parameters are determined by the context of utterance. So the semantic value of any moral sentence, relative to such a sequence of parameters, will be a set of worlds. To make this more concrete, suppose that (1) is uttered by a particular agent, Ana, in a particular context. According to Kratzerian *Ought*, the content of Ana’s utterance in that context—its assertoric content—is a set of worlds. Specifically, it’s the set of worlds w that meets the following condition: at all of the best of the w -accessible worlds, people give to charity.

⁴See e.g., [Rosen 1998](#): 387-388; [Schroeder 2008a, 2010](#); [Charlow 2015](#).

⁵Of course, there remains a question of whether this difference is significant. I return to this in §3.

2.2 Noncognitivist Belief Reports

It's usually thought that a descriptivist semantics leads to a cognitivist theory of moral belief. However, a descriptivist semantics does not, taken on its own, say anything about moral belief. To get an account of moral belief, we would need to extend our semantics to encompass moral belief reports, such as:

- (4) Ana believes that people (morally) ought to give to charity.

When we look at how to analyze these belief reports, we encounter various choice points. Some routes lead to cognitivism; others to noncognitivism.

As a starting point, I will operate with a familiar semantics for belief reports, due to Hintikka 1962. (As with our assumption of Kratzerian analysis of modals, nothing essentially hinges on this choice of framework; the basic strategy proposed here can be adapted to other analyses of belief reports.) According to Hintikka's semantics, *believes* is a universal quantifier over the believer's doxastic alternatives—that is, the set of worlds compatible with what the agent believes. Formally:

Hintikka Believes $[[\alpha \text{ believes } \phi]]^{f,g,w} = 1$ iff $\forall v \in \text{Dox}_\alpha^w : [[\phi]]^{f,g,v} = 1$,

where $\text{Dox}_\alpha^w = \{v \mid v \text{ is compatible with what } \alpha \text{ believes at } w\}$.

Combined with Kratzerian *Ought*, this semantics leads to a form of cognitivism. On the resulting combination, (4) says that Ana believes that all of the morally best of the accessible worlds are worlds where people give to charity. On this analysis, (4) ascribes to Ana a mental state with a mind-to-world direction of fit. This makes it a cognitivist view *par excellence*.

However, we should not be so quick to regard a Hintikkan semantics as the final word on belief reports. Here is a variant that I find particularly promising, and that will serve as the guiding idea in what follows: in addition to quantifying over the believer's doxastic alternatives, moral belief reports shift the ordering source in the index to one provided by the believer's desire-like states.

Fleshing this out requires care; as we'll see, different versions of this thought will yield different results. As a first pass (to be revised shortly), let us start by introducing an agent α 's intrinsic desire function (δ_α), a function from a world w to a set of propositions representing α 's intrinsic desires at w . As a toy example, suppose that at w the only thing that Ana cares about for its own sake is the promotion of wellbeing. Then:

$$\delta_{Ana}(w) = \{\text{wellbeing is promoted}\}.$$
⁶

⁶Insofar as we are after an analysis of moral beliefs (as opposed to normative beliefs more generally), we may

Intrinsic desire functions have the same structure as ordering sources: both are functions from worlds to sets of propositions. Consequently, we can enlist intrinsic desire functions to do the work of an ordering source, specifically, inducing a ranking over worlds. A world v is better than a world u , by the lights of $\delta_\alpha(w)$, provided v comes closer to satisfying α 's intrinsic desires (at w) than u does.⁷ As a first pass, then, we might propose:

Shifty Believes Moral belief reports shift the ordering source to the believer's intrinsic desire function.

This proposal takes us closer to noncognitivism, but not quite there. Rather, what we get is a form of cognitivist subjectivism. To see this, let us compute the truth conditions for (4) using Shifty *Believes* and Kratzerian *Ought*:

$$\begin{aligned} \llbracket \text{Ana believes people (morally) ought to give to charity} \rrbracket^{f,g,w} = 1 \text{ iff} \\ \forall v \in \text{Dox}_{\text{Ana}}^w : \llbracket \text{People (morally) ought to give to charity} \rrbracket^{f,\delta_{\text{Ana}},v} = 1 \text{ iff} \\ \forall v \in \text{Dox}_{\text{Ana}}^w : \forall u \in \text{BEST}_{\delta_{\text{Ana}}(v),f(v)} : \text{people give to charity at } u. \end{aligned}$$

In words: (4) is true at some world w if and only if, for every world v consistent with Ana's beliefs at w , all the v -accessible worlds that come closest to satisfying Ana's intrinsic desires at v are worlds where people give to charity. This is a bit of a mouthful, but the basic idea is not particularly complicated. To make it more intuitive, assume the accessibility relation 'sees' all and only the nearby worlds. Then (4) says that Ana believes that, of all the nearby worlds, her intrinsic desires give highest marks to those where people give to charity.

This is cognitivist subjectivism, not noncognitivism. According to cognitivist subjectivism, moral beliefs are representational mental states. It's just that what they represent is something about the agent's desires. By contrast, the whole idea behind noncognitivism

wish to impose restrictions on the intrinsic desires at issue. After all, someone might have an intrinsic desire for knowledge without thinking that the attainment of knowledge is *morally* required. This is an instance of what Miller 2003 calls the 'moral attitude problem': the problem of distinguishing the conative attitude that constitutes a genuinely moral judgment from other desires. This problem arises for many metaethical theories, including most forms of subjectivism, expressivism, and sentimentalism, and by now a number of solutions have been proposed. Perhaps the moral attitude is closely tied to the reactive attitudes, such as shame, blame, guilt, and resentment (Gibbard 1990; Kauppinen 2010; Björnsson and McPherson 2014; cf. Schroeder 2008a). Or perhaps it is marked by an impartiality condition; perhaps the moral attitude is a desire that we would be disposed to hold even from behind a veil of ignorance (Harsanyi 1977; cf. Kauppinen 2010). For our purposes, we need not take a stand on how to solve this problem. Going forward, the reader is encouraged to interpret my talk of 'intrinsic desires' in whatever way best accords with their preferred solution.

⁷How exactly does an intrinsic desire function induce a ranking over worlds? The short answer: however an ordering source does. A simple option, in line with the treatment in Kratzer 1981, is to say that v is at least as good as u , by the lights of $\delta_\alpha(w)$, just in case every proposition in $\delta_\alpha(w)$ which is true at v is also true at u .

is that moral beliefs are non-representational, motivational states. For the noncognitivists, moral beliefs are not *about* desires. Rather, they *are* desires.

However, a noncognitivist view is just around the corner. We just need to anchor the intrinsic desire function to a particular world. Let a world-indexed intrinsic desire function (δ_α^w) be a constant function from an arbitrary world u to a set of propositions representing α 's intrinsic desires at w (where w may or may not be the same as u). To illustrate, suppose again that the only thing that Ana intrinsically desires at w is the promotion of wellbeing. Then for any world u —even those where Ana has different desires—we get:

$$\delta_{Ana}^w(u) = \{\text{wellbeing is promoted}\}.$$

Noncognitivists can propose:

Noncognitivist *Believes* Moral belief reports shift the ordering source to the agent's world-indexed desire function.

To illustrate, let us recompute the truth conditions for (4) using Noncognitivist *Believes*:

$$\begin{aligned} \llbracket \text{Ana believes people (morally) ought to give to charity} \rrbracket^{f,g,w} = 1 \text{ iff} \\ \forall v \in \text{Dox}_{Ana}^w : \llbracket \text{People (morally) ought to give to charity} \rrbracket^{f,\delta_{Ana}^w,v} = 1 \text{ iff} \\ \forall v \in \text{Dox}_{Ana}^w : \forall u \in \text{BEST}_{\delta_{Ana}^w(v),f(v)} : \text{people give to charity at } u. \end{aligned}$$

In words: (4) is true at w if and only if, for every world v consistent with Ana's beliefs at w , all the v -accessible worlds that come closest to satisfying Ana's intrinsic desires at w (not v !) are worlds where people give to charity. So if the only thing Ana intrinsically desires at w is the promotion of wellbeing, this means that (4) is true if and only if, for every doxastic alternative v , all the worlds nearby to v where wellbeing is promoted are giving worlds.

While the difference between this proposal and Shifty *Believes* may seem subtle, it has important implications. According to Shifty *Believes*, your moral beliefs depend on your beliefs about your desires. According to Noncognitivist *Believes*, your moral beliefs depend on your actual desires. To illustrate the difference, consider the following scenario:

Desire Delusion As before, Ana's only intrinsic desire is to promote wellbeing. But Ana's psychoanalyst claims otherwise: during session, he declares that she only cares about amassing material goods. Ana finds this assessment disturbing, but accepts it nonetheless. Leaving her therapist's office, she spots a child drowning in a pond. As so often happens, she can save the child easily, but at the cost of ruining her clothes.

Now, consider the following possible belief attributions:

- (5) Ana believes she ought to save the child.
- (6) Ana believes she ought to let the child drown.

Shifty *Believes* predicts that (5) is false and (6) is true. After all, Ana believes her psychoanalyst's claim that she only cares about accruing material goods. And at all worlds consistent with her beliefs—and at all worlds nearby to those worlds—saving the child is not conducive to that aim. By contrast, Noncognitivist *Believes* predicts that (5) is true and (6) is false. After all, Ana intrinsically desires promoting wellbeing. And at all worlds consistent with her beliefs—and at all worlds nearby to those worlds—saving the child is the only way of achieving this aim. Only the latter prediction is compatible with noncognitivism. One of the core noncognitivist ideas—to be explored in more detail in §4—is that moral beliefs motivate action in a desire-like way. If this is right, we would expect Ana's moral belief to dispose her to perform whatever action her (actual) desires motivate her to perform. By stipulation, this action is saving the child.

So by combining a contextualist semantics for moral discourse with an ordering-source shifting semantics for moral beliefs, we get a version of descriptivist noncognitivism. Let me pause here to make a couple remarks on the resulting implementation. My proof of concept relied on Kratzer's analysis of modals and Hintikka's analysis of belief reports, both of which have been contested.⁸ However, the basic idea behind my approach is not wedded to the details of either analysis. My approach only requires two assumptions: first, the semantics for moral discourse contains some contextually supplied parameter that induces an ordering or ranking; second, this parameter can be shifted by belief reports to be supplied by desire-like states. These assumptions are fairly minimal, and do not incur any of the most controversial commitments of Kratzer's semantics for modals or Hintikka's semantics for belief reports.

Next, I have formulated Noncognitivist *Believes* as the view that *moral* belief reports shift the ordering source in the index to the believer's desire-like states. This leaves open the possibility that when a non-moral deontic modal is embedded in a belief report, the ordering source does not shift. Consider:

- (7) Ana believes hospital visitors ought to leave by 8pm.

⁸To flag two of the most controversial commitments of a Hintikka semantics: it predicts that beliefs are closed under logical consequence; it also imposes no constraints on which concepts or expressions can appear in belief reports.

While (7) can be read as ascribing a moral belief to Ana, it could also be read as ascribing a belief that the hospital regulations require visitors to leave by 8pm. The latter reading does not impute any desires to Ana; it could be true even if she doesn't care whether visitors comply with the regulations. The framework developed here is flexible enough to allow for both readings. The moral reading is captured by shifting the ordering source to Ana's world-indexed desire function; the non-moral reading is captured using an unshifted ordering source provided by the hospital regulations.

Some might wonder why moral belief reports shift the ordering source, whereas other belief reports do not. Here, two points are worth noting. First, other researchers have independently advanced analogous proposals about the behavior of epistemic modals in belief reports. Consider:

(8) Ana believes that Kendra could be stuck in traffic.

As several authors have noted, it is natural to evaluate the modal in (8) as quantifying over the possibilities compatible with Ana's information (rather than, say, the information of the speaker). In order to capture this reading, these authors propose that *believes* shifts the modal base to the believer's doxastic alternatives.⁹ Importantly, this shift is only obligatory when the modal is given an epistemic reading; if *could* is read as conveying physical or metaphysical possibility (for example), then the modal base remains unshifted. This shows that my proposal is not without precedent or parallel: others have argued, on independent grounds, both that belief reports shift the parameters we use to evaluate modals, and that these shifts are restricted to certain readings of the modal.

Second, the idea that moral belief reports behave differently than descriptive belief reports gains plausibility when we consider the motivations for noncognitivism. Many metaethicists have been attracted to noncognitivism precisely because moral beliefs appear to have a distinctive functional role. As noted above, moral beliefs seem to be intimately tied to motivation, in a way that descriptive beliefs are not. If the functional profile of moral beliefs differs from that of descriptive beliefs, we should not be surprised if moral belief reports behave differently from descriptive belief reports.¹⁰

⁹For proposals in this vein, see [Yalcin 2007](#); [Stephenson 2007](#); [Hacquard 2010](#); [Silk 2017](#); [Ninan 2018](#). (Ninan's semantics is a particularly close analogue of the semantics developed here.)

¹⁰Do only moral belief reports shift the ordering source? This depends on how far we want to take the noncognitivist agenda. Some noncognitivists hold that all normative beliefs—not just moral beliefs—involve desire-like states of mind. Those attracted to this view could propose that every genuinely normative belief ascription shifts the ordering source in the index to some desire-like state of mind, but different species of normative belief are associated with different desire-like states. For further development of this idea, see [Beddor 2021](#).

2.3 Comparison with Kalderon

Having laid out my proof of concept, let me briefly compare the resulting view to the only extant defense of descriptivist noncognitivism, due to [Kalderon 2004](#). While both Kalderon and I agree that it is important to distinguish moral semantics from moral psychology, there are some crucial differences between our approaches.

First, Kalderon does not provide a semantics for moral discourse. According to Kalderon, moral sentences express moral propositions, but he does not tell us what their propositional content is. By contrast, I have showed descriptivist noncognitivism can be implemented using an off-the-shelf semantics for deontic modals.

Second, a key feature of Kalderon's view is his commitment to a version of moral fictionalism. According to Kalderon, when Ana utters (1), she is not genuinely asserting it. Rather, she is performing a distinct speech act of 'quasi-assertion' (2004: 112). This is a sort of pretend-assertion, as when an author purports to make assertions about their fictional creations. Consequently Ana is not committed to the literal truth of (1), just as e.g., Agatha Christie is not committed to the literal truth of her quasi-assertions about Hercule Poirot. This gives rise to familiar worry for moral fictionalism: is it really plausible that people vehemently proclaiming their moral positions in the midst of a heated debate are merely pretending to make assertions? By contrast, my implementation of descriptivist noncognitivism avoids this worry. Ana's utterance of (1) is a genuine assertion, and she is committed to its literal truth.

A third major difference concerns the treatment of moral belief. According to Kalderon, just as we do not really assert moral propositions, we do not really believe moral propositions. But why not? This begins to look rather mysterious. After all, we ascribe moral beliefs to agents by embedding moral sentences under the verb *believes*. If, as Kalderon maintains, the semantic value of an unembedded moral sentence is a moral proposition, why doesn't the belief report ascribe a belief in this very same proposition? By contrast, my version of descriptivist noncognitivism avoids this explanatory challenge. On the view developed here, moral belief reports ascribe genuine beliefs, but having a moral belief just amounts to being in a certain-desire like state of mind. Importantly, this is not merely stipulated; rather, it follows from a general semantics for moral belief reports (Noncognitivist *Believes*), together with our descriptivist semantics for moral language.¹¹

Having developed my implementation of descriptivist noncognitivism and highlighted its advantages over Kalderon's approach, I'll now discuss how the resulting view gets us the best of both worlds.

¹¹Thanks to a referee for helpful discussion here.

3 Descriptivist Benefits

3.1 Possible Worlds Semantics and the Frege-Geach Problem

If asked to list the main objections to expressivism, one would spring to the top of every metaethicist's list: the Frege-Geach Problem.¹² This problem has two parts. The first part is what we might call the 'embedding problem': expressivists owe us a plausible and principled explanation for how moral expressions embed in complex constructions, for example:

- (9) It's not the case that people ought to give to charity.
- (10) Charity alleviates suffering and people ought to give.
- (11) Charity does not alleviate suffering or people ought to give.

It is worth pausing to review why these constructions cause trouble for expressivists. On a possible worlds semantics, logically complex constructions are analyzed in terms of Boolean operations. Conjunction is analyzed in terms of intersection, disjunction in terms of union, negation in terms of complementation, and so on. This provides a simple, principled, and predictive way of analyzing (9)-(11). By contrast, expressivists deny that the meanings of moral utterances can be modeled with sets of worlds. The challenge for expressivists is thus to develop some equally plausible, and equally principled, analysis of complex constructions embedding moral vocabulary.

The second facet of the Frege-Geach Problem is what we might call the 'semantic relations problem': expressivists owe us an account of semantic notions such as truth, consistency, and validity. Here too, there are textbook possible worlds analyses of these notions: a proposition p is true at a world w if and only if $w \in p$; a set of propositions is consistent if and only if there is some possible world where they are all true; a sequence of premises entails a conclusion if and only if for any possible world where all the premises are true, the conclusion is also true. But these analyses are unavailable to expressivists, since they deny that the contents of moral utterances can be modeled with sets of worlds.

Descriptivist noncognitivism avoids both parts of the Frege-Geach Problem. After all, Kratzerian *Ought* is a possible worlds semantics. Consequently, we can help ourselves to the standard possible worlds accounts of both embeddings and semantic relations. In short: *the Frege-Geach Problem simply does not arise for descriptivist noncognitivism.*

Some may think this is too quick. The rest of this section addresses two potential misgivings.

¹²The *locus classicus* of the problem is Geach 1964. This problem has spawned a vast literature; for a sampling of contributions, see Blackburn 1988, 1993; Gibbard 1990, 2003; van Roojen 1996; Schroeder 2008a,b; Silk 2014.

3.2 Revenge of the Frege-Geach Problem?

First, some might worry that the Frege-Geach Problem reveals a structural shortcoming in any noncognitivist theory of moral belief, regardless of whether that theory is paired with a descriptivist semantics. To flesh out this concern, consider the ‘Negation Problem’ (Unwin 1999; Dreier 2006; Schroeder 2008a; Pérez Carballo 2020). This problem arises from the fact that there are three different places where one can insert negation into a moral belief report:

- (12) Benny believes we ought not to give to charity. ($B \Box \neg$)
- (13) Cara believes it’s not the case we ought give to charity. ($B \neg \Box$)
- (14) Dev does not believe that we ought to give to charity. ($\neg B \Box$)

(12)–(14) all ascribe different mental states. To see this, suppose that benighted Benny disapproves of any action that would remedy socio-economic inequalities; as a result, he hopes that everyone ceases their philanthropic efforts. His state of mind is different from that of carefree Cara, who is indifferent as to whether people give to charity. And both differ from divided Dev, who doesn’t have a settled opinion on whether charitable giving is good.

The problem for noncognitivists is that our conative attitudes do not seem to have enough structure to model these differences (see Table 1).

Benny	believing ought not give	desiring not giving
Cara	believing not ought give	??
Dev	not believing ought give	not desiring giving

Table 1: The Negation Problem

However, if we work through what our semantics (Kratzerian *Ought* + Noncognitivist *Believes*) says about (12)–(14), we see that it nicely captures the differences between these three states.

It will help to introduce some terminology. Say that α *decidedly desires* p at a world w if and only if, for all of α ’s doxastic alternatives v , all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds. And say that α *decidedly does not desire* p at w if and only if, for all α ’s doxastic alternatives v , **not** all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds. If we compute the truth conditions of (12)–(14) using Kratzerian *Ought* and Noncognitivist *Believes*, we get the equivalences in Table 2.

These states are distinct. For example, here’s a toy model to distinguish Cara from Dev. Suppose Cara and Dev are both unsure whether charity alleviates suffering. At two of their

Benny	believing ought not give	decidedly desiring not giving
Cara	believing not ought give	decidedly not desiring giving
Dev	not believing ought give	not decidedly desiring giving

Table 2: The Negation Problem Solved

doxastic alternatives (a and b) it does; at two others (c and d) it does not. Suppose that a and b are both accessible from each other, and likewise with c and d (i.e., $f(a) = f(b) = \{a, b\}$; $f(c) = f(d) = \{c, d\}$). And suppose that at a and c people give to charity, and that at b and d they do not. Now, Cara is indifferent as to whether people give to charity regardless of whether doing so alleviates suffering. From the point of view of her desires, a and b are equally desirable, as are c and d . Thus she decidedly does not desire giving: no matter which of her doxastic alternatives obtains, at least some of the most desired of the accessible worlds are worlds where people refrain from giving. By contrast, Dev thinks that if philanthropy alleviates suffering, it is obligatory; otherwise, it is optional. From the point of view of his desires, a is preferable to b , but c and d are equally desirable. So he does not decidedly desire giving. But he also does not decidedly not desire giving. (See Figure 1.)

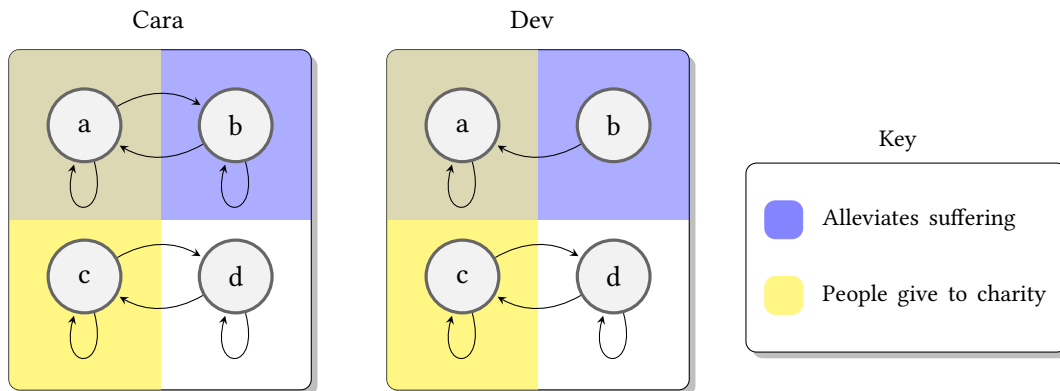


Figure 1: Two states of mind. (Arrows indicate the most desirable of the accessible worlds.)

This solution to the Negation Problem nicely avoids problems that arise for other solutions that have been proposed in the literature. For example, one alternative solution maintains that (13) ascribes a *sui generis* conative attitude such as *tolerating not giving to charity* (Blackburn 1988; Horgan and Timmons 2006). This proposal faces a number of daunting questions. How do we derive the ascription of this *sui generis* attitude from our semantics

for belief reports, together with our semantics for negation? And what explains why this *sui generis* attitude stands in various coherence relations with other moral beliefs? For example, why is *tolerating not giving to charity* inconsistent with *believing people ought to give*?

Another approach identifies moral beliefs with conative attitudes directed towards various reactive attitudes.¹³ On a simple version of this view, to believe people ought to give to charity is to desire praising for giving to charity. This extra complexity distinguishes (12)-(14): (12) says Benny desires praising for not giving to charity; (13) says Cara desires not praising for giving; (14) says Dev does not desire praising for giving. In many respects this is an elegant solution. However, it makes substantive assumptions about the structure of moral judgment—assumptions whose plausibility one might well question. In order for (13) to be true, must Cara desire that people refrain from praising charitable giving? Couldn't her indifference extend to how people bestow their praise? By contrast, the solution to the Negation Problem offered here avoids these concerns.

While I have focused on the interactions between belief and negation, my semantics extends smoothly to cover the interactions with other logical operators. Take disjunction:

- (15) Dev believes that (either) charity does not alleviate suffering or that people ought to give.

Helping ourselves to a Boolean semantics for disjunction, we get the following truth conditions: (15) is true if and only if every world v consistent with Dev's descriptive beliefs is such that either charity does not alleviate suffering at v , or every v -accessible world ranked highest by Dev's intrinsic desires is one where people give to charity. For a model where this ascription comes out true, refer back to Figure 1.¹⁴

3.3 Can Expressivists Help Themselves to the Same Resources?

A second potential misgiving concerns whether descriptivism is really necessary for solving the Frege-Geach Problem. Could we enjoy the same benefits in expressivist setting? A full comparison with expressivist solutions to the Frege-Geach Problem falls outside the scope of this paper. That said, it will be helpful to compare descriptivist noncognitivism to one expressivist framework that might seem particularly well-suited for delivering the same explanatory benefits: relational expressivism.¹⁵

¹³See Schroeder 2008a,c. See also Sepielli 2012 and Beddor 2020 for discussion of whether this approach can be used to model normative uncertainty.

¹⁴This is an example of a belief in a 'mixed disjunction.' See Schroeder 2015a for discussion of the difficulties that mixed disjunctions pose for certain expressivist theories.

¹⁵I borrow the label 'relational expressivism' from Schroeder 2013; see also Toppinen 2013 for development and defense (Toppinen calls this view 'higher-order state expressivism'.) See Ridge 2014 for a closely related

Relational expressivism consists in two main tenets. The first is a thesis about moral psychology. Whereas traditional noncognitivists maintain that moral beliefs are just desire-like states of mind, relational expressivists maintain that ‘moral beliefs consist in a certain relation holding between one’s ordinary descriptive belief state and some kind of desire-like attitudinal state’ (Schroeder 2013: 289).¹⁶ While this characterization is very abstract, Schroeder argues that this thesis is in principle able to deliver a plausible and productive account of complex moral beliefs, provided we assume that the relevant relation R has the property of being *belief-monotonic*: whenever R holds between a desire-like attitude and a descriptive belief state, R also holds between the same desire-like attitude and any strictly more opinionated descriptive belief state.

Our semantics for belief reports can be viewed as one way of implementing this abstract psychological thesis. According to Noncognitivist *Believes*, you believe people (morally) ought to give to charity if and only if you decidedly desire charitable giving. To be in this state is for a certain relation to hold between your descriptive beliefs and your desire-like attitudes. Specifically, it’s for the most desired of all the worlds accessible from your doxastic alternatives to be worlds where people give. Moreover, this relation is belief-monotonic. (After all, if for every world v in $Dox(w)$ the most desired of the v -accessible worlds are giving worlds, the same will hold for any subset of $Dox(w)$.) So Noncognitivist *Believes* offers a detailed implementation of the relational expressivist’s psychological thesis, one that can be used to deliver precise predictions about the conditions under which someone holds a complex belief with moral content. And these predictions in turn provide support for Schroeder’s contention that this psychological thesis can overcome the Frege-Geach Problem at the level of normative thought.

The second tenet of relational expressivism concerns moral language. Relational expressivists claim that moral sentences *express* relational states of mind. Here is where our view gets off the boat. The whole point of descriptivist noncognitivism is that you can get the psychology associated with expressivism without incurring the semantic commitments.

This is a difference, to be sure. But why does it matter? Why not simply say that (1) expresses the state of decidedly desiring giving to charity, and that (9) expresses the state of not decidedly desiring giving to charity, etc.? If we have solved the Frege-Geach Problem at the level of belief, what is to stop us from leveraging this into a general solution to the Frege-Geach Problem?

view. (Thanks to a referee for encouraging me to discuss the connection with relational expressivism.)

¹⁶As Schroeder observes, relational expressivism can be viewed as a generalization of hybrid expressivism, of the sort defended by e.g., Ridge 2006b; Boisvert 2008; Hay 2013; Laskowski 2019; Perl 2020. For further discussion of the relation between the two, see Toppinen 2013.

While this is a natural thought, there are important questions about the semantic details. In particular, what is our relational expressivist lexical entry for *ought*? One option would be to follow traditional expressivists in rejecting truth conditional semantics. For example, one familiar implementation of traditional expressivism takes the meaning of a sentence to be whatever state of mind it expresses.¹⁷ Now, since we have seen that our noncognitivist psychology captures the full range of logically complex thoughts (§3.2), the resulting view is well-positioned to overcome the first facet of the Frege-Geach Problem, the embedding problem. But recall that expressivists also face the semantic relations problem, the problem of explaining semantic notions such as truth, validity, and consistency. Earlier, we saw that descriptivists have a simple and elegant way of interdefining these notions, courtesy of possible worlds semantics. For expressivists who reject truth conditional semantics, it is far from straightforward how to recover these explanations. Now, this is not to say that there is no way of doing so. The point is just that expressivists have their work cut out for them; merely endorsing relational expressivism does not automatically solve these issues.¹⁸

In light of this hurdle, perhaps relational expressivists should try a different tack. Taking a cue from Yalcin 2012 and Silk 2014, they could try appealing to truth conditional resources. Here is one way this might go. Recall that our framework in §2 counted as descriptivist because it embraced Parameter Contextualism (the view that the ordering source and modal base are fixed by the context of utterance). Relational expressivists could retain our truth conditional semantics for modals (Kraterzerian *Ought*) while rejecting Parameter Contextualism. In its place, they could advance an expressivist account of the ordering source:

Ordering Expressivism Assertions of *Ought* ϕ are evaluated using an ordering source that is not supplied by the context of utterance, but which instead models a conative attitude expressed by the speaker.¹⁹

Since the resulting semantics is still truth conditional, it would seem well-positioned to handle the semantic relations problem.²⁰

¹⁷Both Schroeder 2013 and Toppinen 2013 suggest a semantic implementation of relational expressivism along these lines.

¹⁸Schroeder 2013 acknowledges this difficulty. His strategy is to set aside *validity*, and focus on two properties that are closely correlated with validity, what he calls the ‘inconsistency property’ and the ‘inference-licensing property’ (294-295). But even if relational expressivists can make sense of properties that are *correlated* with validity, we still will not have given account of validity, which is what we were after. Moreover, it is not hard to discern the source of the problem. As Schroeder notes, validity is closely related to truth (2013: 294). And relational expressivists who reject truth conditional semantics face difficult questions about how best to assign truth-values to moral sentences.

¹⁹I borrow the label, ‘Ordering Expressivism’ from Silk 2014. One could also opt for an expressivist interpretation of where the modal base comes from; see Yalcin 2007, 2012.

²⁰On the resulting view, the content of a moral sentence, relative to a context of utterance, cannot be modeled

In my eyes, something along these lines is the most promising version of expressivism. Still, I think it would be premature to declare a tie. One worry for Ordering Expressivism is that it is highly doubtful that *all* modals are evaluated relative to ordering sources that express the speaker's desire-like states. Recall that one virtue of Kratzerian *Ought* is that it provides a uniform semantics for all flavors of modals. Consider again:

- (2) You ought to take the 4 train.
- (3) Hospital visitors ought to leave by 8pm.

(2) has a natural interpretation on which the ordering source is provided by the addressee's goals, regardless of whether these goals are shared by the speaker. Similarly, (3) has a natural interpretation on which the ordering source is provided by the hospital regulations, regardless of whether those regulations accord with the speaker's desires.

These examples show that Ordering Expressivism does not hold in full generality. In examples like (2) and (3), it is overwhelmingly natural to suggest that the ordering source is supplied by the context of utterance—just as Parameter Contextualism maintains. Of course, one could restrict Ordering Expressivism to distinctly moral (or, more generally, normative) uses of modals. But then ordering expressivists will be saddled with a disjunctive semantics (or a disjunctive metaseantics): for some uses of modals, the parameters are contextually supplied; for others, they are supplied directly by a desire-like state of mind. But, as far as I can tell, there is no semantic evidence to motivate this disjunctive approach. Absent such evidence, considerations of unity and parsimony speak in favor of the descriptivist option: Parameter Contextualism applies to all uses of modals, moral and otherwise.

Ordering Expressivists may insist that even though there is no semantic evidence for a disjunctive approach, there are philosophical considerations in its favor—namely, all of the explanatory benefits that expressivism delivers. In the next section, I'll examine these putative advantages in more depth. On closer consideration, we find that these are really advantages of a noncognitivist psychology rather than an expressivist semantics. If this is right, then we will have failed to uncover any positive reason for adopting a disjunctive approach, or an expressivist view more generally.

This point can be framed in more general terms. A peculiar feature of moral expressivism is that it is a semantic thesis, yet it is rarely motivated by semantic data. On the contrary: ever since the Frege-Geach Problem was introduced, moral expressivists have been playing catch up with possible worlds semantics; the name of the game has been to come

by a set of worlds. But it can be modeled by a set of world, moral ordering source pairs. In this regard, Ordering Expressivism resembles the expressivist framework developed by [Gibbard 1990, 2003](#).

up with expressivistly kosher explanations of linguistic data that already have perfectly good descriptivist explanations. This is well-illustrated by Ordering Expressivism, which solves the semantic relations problem by mimicking possible worlds semantics as closely as possible while still retaining its expressivist credentials. Even if this can be done, there remains a question as to why we should bother. If it turns out that we can retain all of the benefits of expressivism thanks to a conservative modification of our treatment of belief reports, why go the expressivist route?

4 Noncognitivist Benefits

Let us now turn to consider the major metaethical benefits that expressivism promises to deliver. A standard list includes:

- *Naturalistic respectability*: Expressivists claim to give an account of normative thought and talk in terms of conative attitudes which can be understood without recourse to any ‘spooky’ non-natural properties.
- *The ‘Open Question’ phenomenon*: Even if an agent knows all of the descriptive facts about a course of action, the question of whether they ought to perform this action might remain ‘open’ for them. Expressivists purport to explain this phenomenon.
- *Motivational internalism*: Expressivists seek to explain the intimate connection between moral judgment and motivation.
- *Explaining disagreement*: Expressivists claim to explain how two people can agree on all of the descriptive facts while still disagreeing over moral matters.

In this section, I consider these putative advantages one by one. I’ll argue that each turns out to be an advantage of noncognitivism rather than expressivism. Consequently, they provide no reason for favoring expressivism over descriptivist noncognitivism.

4.1 Naturalistic Respectability

Since its earliest days, expressivism has been advertised as a naturalistically kosher theory of moral thought and talk. This naturalistic ambition was front and center in the work of the early emotivists (e.g., [Ayer 1936](#)), who sought to understand moral discourse without appealing to any mysterious, unverifiable entities. And it is also prominent in the work of contemporary quasi-realists such as [Blackburn 1998](#) and [Gibbard 1990, 2003, 2008](#), who

aim to situate moral judgments in a naturalistic worldview. According to the expressivist picture, ‘ought’-y judgments are analyzed in terms of desire-like states of mind, which are themselves directed towards non-normative contents.

Important as this advantage may be, it is shared by descriptivist noncognitivism. According to the noncognitivist part of the theory, for Ana to believe that people ought to give to charity is for her to decidedly desire that people give to charity. So our framework also explains moral thought in naturalistically respectable terms.

Some might object that this only tells half of the story. According to this objection, we should aim to give a naturalistically respectable account of *both* moral thought and moral talk. Even if descriptivist noncognitivism provides the former, it doesn’t provide the latter. After all, Kratzerian *Ought* analyzes unembedded moral sentences using a moral ordering source. But it does not tell us how to understand this moral ordering source in naturalistically respectable terms.

However, Kratzerian *Ought* is perfectly compatible with any number of metaethical accounts of the moral ordering source, including naturalistic accounts. For reasons that will become clear in §5, I think the most promising way of developing descriptivist noncognitivism will explain the moral ordering source in terms of the conative attitudes of the conversational interlocutors. If this is right, then there is no need to invoke *sui generis* moral properties in our semantics.²¹

Expressivists may retort that any such view is implausible on independent grounds. After all, expressivists often maintain that such views are unable to capture the other data points that motivate expressivism: the Open Question phenomenon, motivational internalism, or disagreement data. In order to assess this worry, let us investigate these further data points.

4.2 The Open Question Phenomenon

Moore 1903 infamously outlined a recipe for refuting any naturalistic analysis of a normative claim. Take a naturalistic paraphrase of (1), for example:

(16) Charitable giving promotes wellbeing.

Moore invites us to compare the questions:

²¹As a reviewer helpfully emphasizes, this approach incurs an additional level of context-sensitivity beyond that which is hardwired into the Kratzerian semantics. For Kratzer, context determines whether a given modal is used to make a moral evaluation, but nothing in this approach requires that there is further contextual shiftiness within moral uses. By contrast, an important feature of the sort of views to be explored here is that there is no single propositional content that is contributed by every moral use of a modal sentence. Instead, in different contexts (1) will contribute different contents, each of which is naturalistically respectable in its own right.

- (17) Granted that charitable giving promotes wellbeing, but ought we give to charity?
 (18) Granted that we ought to give to charity, but ought we give to charity?

According to Moore, (17) is an ‘open’ question, whereas (18) is not. Moore took this to show that (17) and (18) differ in meaning, and hence that (1) and (16) also differ in meaning. Whereas Moore used this style of argument to support his non-naturalist metaethics, many expressivists have co-opted it for their own ends. According to the expressivist diagnosis, (1) and (16) differ in meaning because (16) is purely descriptive, whereas (1) expresses approval of charitable giving. Expressivists thus purport to account for the ‘openness’ of moral judgments while maintaining naturalistic respectability.²²

History has not been kind to the Open Question Argument. One common complaint is that the argument assumes that the meaning of an expression is always transparent to the speaker. A related objection is that the argument proves too much, since it seems to rule out the possibility of any non-trivial semantic analysis. As Darwall et al. put it, the argument threatens to “bring the whole enterprise of conceptual analysis to a standstill” (1992: 115).²³

These are fair criticisms of standard formulations of the Open Question Argument. But we can devise an alternative version that sidesteps these problems. The variant I have in mind does not make any claims about *meaning* at all. Rather, it focuses entirely on *belief*.

One of Moore’s key observations was that a coherent agent can believe that charitable giving promotes wellbeing without believing that people ought to give to charity. More generally:

Doxastic Openness For any natural property N , a coherent agent can believe that ϕ ing is N without believing that they (or anyone else) morally ought to ϕ .

This is an important insight, and any adequate moral psychology should account for it. Noncognitivism offers a promising explanation. According to noncognitivists, one can believe that ϕ ing is N without desiring ϕ ing, and hence without believing people ought to ϕ . This explanation avoids any dubious commitments to the transparency of meaning or the impossibility of conceptual analysis.

So by recasting the Open Question Argument in terms of belief, we avoid the implausible commitments of traditional formulations. But once we reformulate the argument in this fashion, we find that its real beneficiary is noncognitivism, not expressivism.²⁴

²²See Ayer 1936: chp. 6; Hare 1952: 92-93. See Horgan and Timmons 1992 for a variant of this argument.

²³See Harman 1977; Smith 1994 for related worries.

²⁴While noncognitivism accounts for Doxastic Openness, it does predict that certain normative questions will be doxastically closed for certain agents. In particular, it predicts that if an agent intrinsically desires p and

4.3 Motivational Internalism

Many metaethicists have been attracted to the ‘internalist’ idea that moral judgments are inherently action-guiding.²⁵ To motivate this idea, suppose Ana tells us that she is firmly convinced that everyone ought to donate to charity. But, she hastens to add, she feels no inclination whatsoever to donate. Most of us would react to this speech with surprise; we might even wonder whether Ana really does hold the moral belief that she professes.

Now, we should be careful not to overstate the connection between moral judgment and motivation. Akrasia is a familiar phenomenon: people regularly fail to act in accordance with their moral ideals, due to apathy or competing preferences.²⁶ So any halfway plausible formulation of the internalist thesis will need to be suitably qualified, for example:

Motivational Internalism Necessarily, if an agent believes that they morally ought to ϕ , then they will be at least somewhat disposed to try to ϕ .

A formulation along these lines makes room for akrasia. Someone might be disposed to try to perform some action while failing to do so, since this disposition might be masked by any number of factors.

Suppose we accept Motivational Internalism. This thesis calls for an explanation. What endows moral beliefs with their motivational force? Expressivists claim to provide an answer.²⁷ However, once we are careful to distinguish between expressivism and noncognitivism, we see that expressivism does not really explain Motivational Internalism. Motivational Internalism is a thesis about the connection between *belief* and motivation. What really explains this connection is noncognitivism. And it is easy to see how the explanation goes. According to noncognitivism, moral beliefs are desire-like states. Part of the

believes that ϕ ing is the only way to bring about p , then they cannot coherently wonder whether one ought to ϕ . Some may find this implausible. Suppose Edmond’s only intrinsic desire is to wreak vengeance on anyone who has wronged him, and he is certain that publicly humiliating his enemies is the only way of bringing this about. Still, it seems, he might be unsure whether he ought to publicly humiliate his enemies.

In response, we should remind ourselves that every noncognitivist needs to solve the ‘Moral Attitude Problem’: they need some way of distinguishing the conative attitudes that constitute moral beliefs from other desires. Different theorists try to solve this problem in different ways. As we noted in fn. 6, some have proposed that moral attitudes are distinguished by their connections with the reactive attitudes. On this view, if Edmond’s desire for vengeance constitutes a moral attitude, he must be disposed to blame those who stand in the way of his desire, and feel remorse if he himself frustrates its fulfillment. Others have proposed that moral attitudes must be impartial, in the sense that they would be held even from behind a veil of ignorance. Once we stipulate that Edmond’s vengeful desire checks all of these boxes, it becomes more plausible to maintain that he really is committed to believing that he ought to humiliate his enemies.

²⁵This idea traces back to Falk 1947. For contemporary defenses, see Dreier 1990; Smith 1994; Korsgaard 1996; Blackburn 1998; Gibbard 2003; Egan 2012, among many others.

²⁶For discussion, see Stocker 1979; Smith 1994; Mele 1996; Svavarsdóttir 1999; van Roojen 2010.

²⁷Virtually all of the major defenses of expressivism appeal to some version of Motivational Internalism. See e.g., Stevenson 1944; Blackburn 1998; Gibbard 1990, 2003; Charlow 2015.

functional role of desire is to motivate action: desiring ϕ disposes an agent to try to ϕ . So if Ana believes everyone ought to give to charity, it follows that Ana will be disposed to try to make everyone give (including herself). This explanation is silent on whether moral language gets an expressivist or a descriptivist semantics.

4.4 Explaining Disagreement

Let us turn to the last apparent advantage of expressivism on our list: its ability to explain moral disagreement. Plausibly, people can disagree over some moral matter without disagreeing in their descriptive beliefs. Ana and Benny might agree on all of the descriptive facts concerning charitable giving—for example, they might agree that giving $\$x$ to such-and-such causes would produce y utiles in such-and-such people. But they might still disagree over whether people ought to give, so long as Ana wants to promote wellbeing and Benny does not.

According to expressivists, the lesson to be learned here is that two people can ‘disagree in attitude’—that is, they can disagree in virtue of having clashing desire-like states. This idea traces back to Stevenson:

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... (Stevenson 1944: 3)

What does it mean to ‘disagree in attitude’? While Stevenson describes a familiar phenomenon, it has proved challenging to provide a theory of disagreement in attitude that avoids either overgenerating or undergenerating disagreements (cf. Dreier 2006). However, our solution to the Negation Problem provides a promising path forward. Recall our terminology from §3.2: α decidedly desires p at w if and only if for all of α ’s doxastic alternatives v , all of the v -accessible worlds ranked highest by δ_α^w are worlds where p holds. And α decidedly does not desire p at w if and only if for all α ’s doxastic alternatives v , not all of the v -accessible worlds ranked highest by δ_α^w are worlds where p holds. We can propose:

Disagreement in Attitude α and β disagree in their desires regarding p iff α decidedly desires p and β decidedly does not desire p , or *vice versa*.

An advantage of this account is that it captures our intuitions about the disagreements—or lack thereof—between our earlier characters, Ana, Benny, Cara, and Dev. It correctly

predicts that Ana and Benny disagree in their desires regarding charitable giving: she decidedly desires it; he decidedly does not. It also correctly predicts that Ana disagrees with Cara. (Recall that Cara is indifferent to whether people give to charity no matter what the world is like. Hence she, like Benny, decidedly does not desire charitable giving.) At the same time, our account predicts that Ana does not disagree with divided Dev, who is unsure whether people ought to give. This also seems like the right result. Compare: if Ana believes that Singapore is north of Jakarta, and Dev has no opinion about the matter, it seems wrong to describe them as disagreeing.

As a bonus, this account dispels a residual worry that some may have about the whole strategy of appealing to disagreement in attitude. The worry is that this strategy saddles us with a disjunctive account of disagreement. After all, the reasoning goes, we will wind up with one story about disagreements in descriptive belief, and some separate story about disagreements in attitude. But we won't have explained why both these phenomena deserve the name 'disagreement.'

Our semantics for belief reports allows us to provide a general semantics for disagreement ascriptions—one that encompasses both moral and descriptive disagreements. The starting idea is this: all disagreements arise from conflicting beliefs. This is, I think, a very intuitive view. (I suspect the appeal of this view is part of what makes some philosophers skeptical of the very idea of disagreement in attitude.) The next step is to semantically ascend, converting this intuitive thought into a semantics for disagreement ascriptions:

Disagreement Ascriptions $\llbracket \alpha \text{ and } \beta \text{ disagree over whether } \phi \rrbracket^{f,g,w} = 1$ iff both:

1. $\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,w} = 1$,
2. $\llbracket \beta \text{ believes } \neg \phi \rrbracket^{f,g,w} = 1$

or *vice versa*.

Applied to moral disagreements, this account predicts that α and β disagree over whether people ought to give to charity if and only if α believes people ought to give to charity, and β believes that it's not the case that people ought to give to charity (or *vice versa*). Given our semantics, this will occur just in case α decidedly desires people to give to charity, and β decidedly does not desire that people give to charity (or *vice versa*). So disagreement in attitude is really just a special instance of the more general phenomenon of doxastic disagreement.

One virtue of this account is that it explains why our intuitions about disagreement line up so closely with our intuitive belief ascriptions. For example, our judgment that Ana disagrees with Benny and Cara over whether people ought to give tracks the fact that Ana

believes that people ought to give, whereas Benny and Cara believe that it's not that people ought to give. Disagreement Ascriptions explains why these judgments align.²⁸

The upshot: combining Kratzerian *Ought* with Noncognitivist *Believes* yields a promising new approach to disagreement. Of course, this is hardly the final word on the matter; a full assessment of this approach deserves a paper of its own. For present purposes, the main point is this: if anything like this theory is on the right track, then all of the explanatory work comes from the noncognitivist treatment of moral belief. Nothing in the story presupposed expressivism.²⁹

4.5 Taking Stock

This section considered four arguments that have been used to motivate expressivism. Each of these arguments has an important grain of truth, but the import of each has been systematically misunderstood. Properly unpacked, these are all arguments for a noncognitivist psychology rather than an expressivist semantics.

Having laid out descriptivist noncognitivism and advertised its benefits, I now turn to consider some potential objections.

5 Questions and Objections

5.1 Thought-talk connections

I've argued that we need to distinguish between moral semantics and moral psychology. But surely, one might protest, there is some close connection between the two. This gives rise to a natural worry for descriptivist noncognitivism: it draws too great a gulf between the meanings of our moral assertions and the nature of moral belief.

One way of fleshing out this worry is to focus on the pragmatics of moral communication. Consider: What is Ana trying to accomplish when she asserts (1)? Here's a very

²⁸For further development and defense of this approach to disagreement, see [Beddor 2019a](#). There, I argue that this account of disagreement is well-positioned to avoid problems facing other treatments of disagreement.

²⁹An interesting question is whether expressivists could take on board the analysis of disagreement developed here. For example, we saw in §3.3 that relational expressivists might consider combining Noncognitivist *Believes* with Ordering Expressivism. Those who go this route might also adopt Disagreement Ascriptions. Such an expressivist account would be interesting in its own right, and would have advantages over many extant expressivist treatments of disagreement. (See [Toppinen 2013: 272-280](#) and [Schroeder 2015b: 23](#) for discussion of the problems that disagreement poses for relational expressivists.) However, from the perspective of this paper, it is unclear why one would pursue this expressivist variant of our theory. A central theme in this paper is that there is no need to go expressivist after all: descriptivist noncognitivism gets all of the advantages associated with expressivism while still allowing us to retain a traditional possible worlds semantics.

plausible answer: She intends to get her audience to believe that people ought to give to charity. More generally:

Moral Pragmatics Normally, a speaker's goal in asserting *Ought* ϕ is to get their audience to believe ought ϕ .

For descriptivist cognitivists, there is no mystery as to why Moral Pragmatics holds. It follows from a textbook pragmatics, according to which the goal of making an assertion is to get your interlocutors to believe its assertoric content:

Assertoric Update Normally, a speaker's goal in asserting a sentence ψ , in a context with a sequence of contextually determined parameters c , is to get their audience to believe $\llbracket\psi\rrbracket^c$.³⁰

But for descriptivist noncognitivists, Moral Pragmatics is not an instance of Assertoric Update. On our framework, if Ana asserts (1) in some context c , the assertoric content of her utterance ($\llbracket(1)\rrbracket^c$) is a possible worlds proposition. Now, exactly what proposition this is will depend on how we understand the moral ordering source. For the sake of illustration, suppose we adopt a simple subjectivist conception of the ordering source. Then $\llbracket(1)\rrbracket^c$ is a proposition about Ana's desires. Specifically, it's the proposition that, in all of the accessible worlds ranked highest by Ana's desires, people give to charity. Now for the trouble: by the noncognitivist's lights, just because Ana's hearer H comes to believe this proposition about Ana does not mean that H will believe people ought to give to charity. (H might grant that Ana desires for people to give to charity, without sharing her desire.) So it seems we cannot use Assertoric Update to underwrite Moral Pragmatics. Some may think this portends a more general trouble for descriptivist noncognitivism: once we draw a divide between moral talk and moral thought, we have trouble explaining how speakers use the former to influence the latter.

While I agree that this is a serious concern, I think we should not admit defeat too readily. Let me sketch two possible responses. The first response is to keep Assertoric Update, but supplement it with an argument that normally when a hearer updates on the assertoric content of *Ought* ϕ they will thereby come to believe ought ϕ . The second response involves rejecting Assertoric Update in favor of a pragmatic theory that plays better with our framework.

To introduce the first response, note that our statement of the problem two paragraphs above assumed a speaker subjectivist conception of the ordering source. But these days

³⁰See Stalnaker 1978 for an influential early statement of this view.

speaker subjectivism is widely regarded as a non-starter, and a variety of sophisticated contextualist proposals have been offered as replacements. For example, we might adopt a version of group contextualism, and hold that the ordering source reflects the shared desires of the conversational participants. Better, we could adapt an idea from [Mandelkern 2020](#) and hold that the ordering source reflects the shared desires of the conversational interlocutors *after the assertion has been made*. As Mandkelkern notes, this qualification is in a better position to capture the dynamic force of modal assertions: in making a modal assertion, we do not simply aim to describe the conversational state as it currently stands; often, we aim to change this state.³¹ This yields a version of what Mandelkern calls ‘Prospective Contextualism’:

Prospective Contextualism The moral ordering source is the set of propositions describing the intrinsic desires shared by all of the conversational participants after the moral assertion is made.

On this view, the assertoric content of (1) is roughly equivalent to the following:

- (19) All of the current conversational interlocutors, after hearing this assertion, will be in a desire-like state that gives highest marks to the accessible worlds where people give to charity.

Suppose we combine Prospective Contextualism with Assertoric Update. Now consider some interlocutor *H* who comes to believe the assertoric content of (1). Since *H* is one of the conversational participants—and knows as much—*H* is committed to believing that *H*’s own intrinsic desires give highest marks to the accessible worlds where people give to charity. Now, add the further assumption that people typically have introspective access to their own desires:

Normal Access Normally, an agent believes they are in a particular desire-like state of mind only if they are in that desire-like state of mind.

Clearly, it would be much too strong to claim that we always have introspective access to our desires. We’ve already seen one case where such introspective access fails (Desire Delusion); such cases are crucial for distinguishing descriptivist noncognitivism from cognitivist subjectivism. But the idea that we normally have such access is hard to dispute.

³¹Mandelkern develops a version of this semantics for epistemic modals. Here I extend his ideas to moral uses of deontic modals. A number of other contextualists have defended the idea that in making moral assertions, we often try to change the conversational parameters; see, for example, [Lewis 1979](#); [Plunkett and Sundell 2013](#); [Silk 2015](#); [Khoo and Knobe 2018](#).

So if H is in normal circumstances, then H will, post-assertion, decidedly desire for people to give to charity. By Noncognitivist *Believes*, H will believe that people ought to give to charity. Moreover, all of this is foreseeable by our speaker, Ana. That is, assuming Ana is aware of the assertoric content of (1), she can know that, normally, anyone who comes to believe this content will believe that people ought to give to charity. So if she aims to get her interlocutors to hold this moral belief, a reasonable way of trying to accomplish this goal is by asserting (1).

This provides one way to secure thought-talk connections. According to this response, someone who asserts *Ought* ϕ in a context c aims to get their audience to believe $\llbracket \text{Ought } \phi \rrbracket^c$ —just as Assertoric Update has it. However, achieving this aim will normally—and foreseeably—have the consequence that their audience will also believe *ought* ϕ .³²

The second response to the challenge is more direct. Note that the trouble only arises because our view requires us to distinguish between assertoric content and compositional semantic value.³³ The assertoric content of a moral sentence, relative to a context, is a set of worlds. The compositional semantic value of a moral sentence is a function from modal bases and ordering sources to a set of worlds. Once we distinguish between assertoric content and compositional semantic value, we can ask ourselves whether some of the theoretical work that has been traditionally allocated to assertoric content should be re-assigned to semantic value. In particular, we might venture the hypothesis that the primary pragmatic function of asserting a sentence is to get one’s interlocutors to believe its semantic content, rather than its assertoric content:

Semantic Update Normally, a speaker’s goal in asserting a sentence ψ is to get one’s audience to believe $\llbracket \psi \rrbracket$.

³²This response can also be used to underwrite various thought-talk connections between a speaker’s assertions and their *own* beliefs. As a reviewer notes, it sounds incoherent to claim:

- (i) ?? People (morally) ought to give to charity, but I don’t believe they (morally) ought to do so.

To explain the oddity of this conjunction, we can appeal to a speaker inclusion constraint: the speaker’s intrinsic desires always contribute to determining the moral ordering source. This constraint follows immediately from Prospective Contextualism: according to Prospective Contextualism, the moral ordering source pools the intrinsic desires shared by all the conversational interlocutors post-assertion, and the speaker will always be one such interlocutor. (See Egan et al. 2005 for discussion of an analogous constraint on epistemic modals.) To see how the explanation goes, suppose a speaker S utters (i). Assume S is sincere, and hence believes the assertoric content of (i). By the speaker inclusion constraint, S is committed to believing that S ’s desire-like attitudes give top marks to the accessible worlds where people give to charity. By Normal Access, it follows that if S is in normal circumstances, then S ’s desire-like attitudes give top marks to accessible worlds where people give to charity, and hence that S believes that people (morally) ought to give to charity (by Noncognitivist *Believes*). But this contradicts the second conjunct of (i). So we have an explanation of why conjunctions like (i) seem bizarre: in normal circumstances, anyone who sincerely uttered them would be committed to a contradiction.

³³For more on this distinction, see Lewis 1980; Yalcin 2007; Ninan 2010; Rabern 2012.

This replacement validates Moral Pragmatics. When Ana asserts (1), she is trying to get her audience to believe its semantic content ($\llbracket(1)\rrbracket$). By Noncognitivist *Believes*, this will happen if and only if her audience believes that people ought to give to charity.³⁴

Some might worry that if we revise our pragmatics in this way, there will be no work left for assertoric content to do. However, this would be too quick; there are a number of other roles that assertoric contents are well-suited to fill. One particularly important role is to serve as the bearers of truth and falsity. Assertoric contents can play this role because they have representational truth conditions: relative to any context, the truth or falsity of Ana's assertion depends entirely on what the world is like.

Where does this leave us? Our objector rightly insisted that there is a close connection between moral assertion and moral belief—a connection seen most clearly in the pragmatics of moral communication. But we also saw that descriptivist noncognitivists have two promising strategies for securing this connection. I will defer a detailed cost-benefit analysis of these strategies to another occasion. The important point is that the availability of these strategies should bolster confidence that our framework can preserve the important connections between moral thought and moral talk.

5.2 Truth-value assessments

Another concern is that our representational truth conditions stand in tension with our account of moral disagreement. To see the objection in its starkest form, let us again assume a speaker subjectivist conception of the moral ordering source. Then Ana's assertion of (1) is true if and only if her intrinsic desires give high marks to the accessible worlds where giving occurs. Now, we've seen that our framework still allows us to explain why Benny disagrees with Ana, in virtue of having conflicting desires (§4.4). But a residual worry remains. Suppose Benny grasps the truth conditions for Ana's utterance. Then he should be willing to say things like:

- (20) ? What Ana says is perfectly true. Still, I disagree with her over whether people ought to give to charity.

³⁴This proposal can also be leveraged to provide an alternative explanation of the oddity of (i) (*People (morally) ought to give to charity, but I don't believe that they morally ought to do so*). Suppose we adjust our understanding of sincerity. Rather than saying that someone sincerely asserts a sentence only if they believe its assertoric content, we could say that someone sincerely asserts a sentence only if they believe its semantic content. From this it follows that anyone who sincerely uttered the first conjunct of (i) would believe its semantic content ($\llbracket\textit{People (morally) ought to give to charity}\rrbracket$), and hence, by Noncognitivist *Believes*, that they believe people (morally) ought to give. But this would contradict the second conjunct of (i). Consequently (i) possesses one of the cardinal defects of the Moore-paradoxical sentence, '*p*, but I don't believe *p*': it cannot be both true and sincerely believed.

Some might find this prediction implausible. If Benny believes that people ought not give to charity, then presumably he will deem Ana's utterance of (1) to be false.

Here too, I want to sketch two possible responses. As with the previous objection, one option is to switch to a more sophisticated conception of the moral ordering source. For example, Prospective Contextualism offers a clean explanation of why (20) has an air of infelicity. If, post-assertion, Benny still does not desire charitable giving, then Ana's bid to change the conversational context has failed. According to Prospective Contextualism, this renders her assertion false. So the second conjunct of (20) guarantees the falsity of the first.

An alternative response is to scrutinize whether (20) is really infelicitous. A number of philosophers and linguists have argued on independent grounds that it's possible to disagree with a speaker while judging that their utterance is literally true. Consider disagreements with implicatures:

- (21) a. *A*: John has two sisters.
b. *B*: ✓I disagree/#That's false.

If *B* thinks John has three sisters, then *B* can felicitously register disagreement with *A*. But *B* cannot felicitously dispute the truth of *A*'s assertion (Horn 1989; Sundell 2011; Khoo 2015). Of course, the analogy only goes so far: on our view, Benny is disagreeing with Ana's belief, not with an implicature. Still, (21) provides evidence that it's possible for two speakers to disagree despite neither of them speaking falsely.

In fact, experimental work by Khoo and Knobe 2018 suggests that moral disagreements are precisely such a case. Khoo and Knobe recruited participants through Amazon Mechanical Turk and presented them with vignettes in which two speakers make different claims about the moral status of a particular action (e.g., 'What so-and-so did was wrong'; 'No, what so-and-so did was not wrong.'). Khoo and Knobe found high rates of agreement with the claim that the two speakers disagree, but significantly lower rates of agreement with the claim that at least one of their judgments must be incorrect.³⁵ These data provide independent reason to think that our truth-value assessments of moral utterances can come apart from our intuitions about disagreement—indeed, they can come apart in precisely the way that (20) countenances.

So we have two strategies for responding to the objection at hand. Which response should we prefer? Ultimately this depends on just how bad speakers find (20). This is an

³⁵The rates of agreement with the latter claims were lower when it was stipulated that the two speakers came from different cultures, and lowest when it was stipulated that the speakers came from different planets. But in each of these conditions participants strongly assented to the claim that the two speakers disagree.

empirical issue, on which we can afford to remain noncommittal. However this issue plays out, truth-value assessments should cause little concern for descriptivist noncognitivists.³⁶

6 Conclusion

This paper offered a defense of descriptivist noncognitivism. This view combines a descriptivist semantics for moral language with a noncognitivist moral psychology. Despite the apparent oddity of this combination, there is much to be said on its behalf. In contrast with previous explorations of this idea, I showed that the view can be given a precise compositional implementation in a possible worlds semantics. Moreover, it reaps all of the benefits standardly associated with expressivism, thanks to its noncognitivist psychology. Finally, the implementation developed here has ample resources for securing the connections between moral belief and communication.

This paper also carries a more general lesson, which is that we should be careful to distinguish between semantics and psychology. If I'm right, many metaethicists have wrongly concluded that the problems for a semantic framework (expressivism) are problems for a psychological thesis (noncognitivism), and similarly that the advantages of the latter thesis redound to the credit of the former. One advantage of exploring a view like descriptivist noncognitivism is that it helps us resist this natural—but ultimately mistaken—line of argument.³⁷

³⁶A more revisionary option is also worth mentioning. We could recast our semantics for *ought* in a relativist setting: we could take our circumstances of evaluation to be centered worlds (ordered pairs of a world and assessor). We could then propose that the ordering source used to evaluate an unembedded moral use of a modal is supplied by the desire-like attitudes of the assessor. The semantics for belief reports would still work much as before. As before, we could adopt a Hintikka semantics, on which *believes* universally quantifies over the believer's doxastic alternatives (now modeled as centered worlds). As before, moral belief reports would shift the ordering source to be supplied by the believer's world-indexed desire function. This 'relativist noncognitivism' still has good claim to be a descriptivist view. While the assertoric content of a moral sentence will not be a set of worlds, it will be a set of centered worlds. And a strong case can be made for thinking that these centered contents are still representational: what they represent is self-locating information about one's own desires. This framework provides another explanation of the apparent oddity of (20). If Benny desires that people avoid giving, Ana's utterance is false at his context of assessment. Assuming that agents typically evaluate utterances for truth using their own context of assessment, we would expect Benny to deem her utterance false. Whether one finds this relativist solution attractive will depend on one's more general views about the viability of the relativist program.

³⁷For helpful comments on this material, I am grateful to Dan Drucker, Simon Goldstein, Dan Greco, Dan Hoek, Stephanie Leary, Ben Lennertz, Matt Mandelkern, the NUS Philosophy Reading Group, two anonymous referees for this journal, and audiences at Stanford University, the University of Vienna, the NY Philosophy of Language Workshop, and the 7th Formal Ethics Conference.

References

- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- Bob Beddor. Subjective Disagreement. *Noûs*, 53(4):819–851, 2019a.
- Bob Beddor. Noncognitivism and epistemic evaluations. *Philosophers' Imprint*, 19(10):1–27, 2019b.
- Bob Beddor. Fallibility for Expressivists. *Australasian Journal of Philosophy*, 98(4):767–777, 2020.
- Bob Beddor. Moral and epistemic evaluations: A unified treatment. *Philosophical Perspectives*, 35(1):23–49, 2021.
- Gunnar Björnsson and Tristram McPherson. Moral Attitudes for Non-Cognitivists: Solving the Specification Problem. *Mind*, 123(489):1–38, 2014.
- Simon Blackburn. Attitudes and Contents. *Ethics*, 98(3):501–517, 1988.
- Simon Blackburn. *Essays in Quasi-Realism*. Oxford University Press, Oxford, 1993.
- Simon Blackburn. *Ruling Passions*. Clarendon Press, Oxford, 1998.
- Daniel Boisvert. Expressive-Assertivism. *Pacific Philosophical Quarterly*, 89:169–203, 2008.
- Nate Charlow. Prospects for an Expressivist Theory of Meaning. *Philosophers' Imprint*, 15(23), 2015.
- Stephen Darwall, Allan Gibbard, and Peter Railton. Toward *Fin de Siècle* Ethics: Some Trends. *The Philosophical Review*, 101(1):115–189, 1992.
- James Dreier. Internalism and Speaker Relativism. *Ethics*, 101(1):6–26, 1990.
- James Dreier. Negation for Expressivists: A Collection of Problems with a Suggestion for Their Solution. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1. Oxford University Press, Oxford, 2006.
- Andy Egan. Relativist Dispositional Theories of Value. *Southern Journal of Philosophy*, 50(4):557–582, 2012.
- Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Preyer and Peter, editors, *Contextualism in Philosophy*. Oxford University Press, Oxford, 2005.
- W.D. Falk. 'Ought' and Motivation. *Proceedings of the Aristotelian Society*, 48:492–510, 1947.
- Peter Geach. Assertion. *Philosophical Review*, 74:449–465, 1964.
- Allan Gibbard. *Wise Choices, Apt Feelings*. Harvard University Press, Cambridge, MA, 1990.
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, MA, 2003.
- Allan Gibbard. *Reconciling Our Aims*. Oxford University Press, Oxford, 2008.
- Valentine Hacquard. On the Event Relativity of Modal Auxiliaries. *Natural Language Semantics*, 18(1):79–114, 2010.
- R.M. Hare. *The Language of Morals*. Oxford University Press, Oxford, 1952.

- Gilbert Harman. *The Nature of Morality*. Oxford University Press, New York, 1977.
- John Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge, 1977.
- Ryan Hay. Hybrid Expressivism and the Analogy Between Pejoratives and Moral Language. *European Journal of Philosophy*, 21(3):450–474, 2013.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- Terry Horgan and Mark Timmons. Troubles for New Wave Moral Semantics: The ‘Open Question Argument’ Revived. *Philosophical Papers*, 21(3):153–175, 1992.
- Terry Horgan and Mark Timmons. Cognitivist Expressivism. In Horgan and Timmons, editors, *Metaethics After Moore*. Oxford University Press, New York, 2006.
- Laurence Horn. *A Natural History of Negation*. University of Chicago Press, Chicago, 1989.
- Mark Eli Kalderon. Open Questions and the Manifest Image. *Philosophy and Phenomenological Research*, 68(2):251–289, 2004.
- Antti Kauppinen. What Makes a Sentiment Moral? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1. Oxford University Press, New York, 2010.
- Justin Khoo. Modal Disagreements. *Inquiry*, 5(1):1–24, 2015.
- Justin Khoo and Joshua Knobe. Moral Disagreement and Moral Semantics. *Noûs*, 52(1):109–143, 2018.
- Sebastian Köhler. Expressivism, Belief, and All that. *Journal of Philosophy*, 2017.
- Christine Korsgaard. *Creating the Kingdom of Ends*. Cambridge University Press, Cambridge, 1996.
- 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.
- Nicholas Laskowski. The Sense of Incredibility in Ethics. *Philosophical Studies*, 176(1):93–115, 2019.
- David Lewis. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1):339–359, 1979.
- David Lewis. Index, Context, and Content. In Kanger and Ohman, editors, *Philosophy and Grammar*, pages 79–100. Reidel, Amsterdam, 1980.
- Matthew Mandelkern. How to Do Things With Modals. *Mind and Language*, 35(1):115–138, 2020.

- Alfred Mele. Internalist Moral Cognitivism and Listlessness. *Ethics*, 106:727–753, 1996.
- Alexander Miller. *Introduction to Contemporary Metaethics*. Polity, Cambridge, 2003.
- G.E. Moore. *Principia Ethics*. Cambridge University Press, Cambridge, 1903.
- Dilip Ninan. Semantics and the Objects of Assertion. *Linguistics and Philosophy*, 33(5): 355–380, 2010.
- Dilip Ninan. Relational Semantics and Domain Semantics for Epistemic Modals. *Journal of Philosophical Logic*, 47(1):1–16, 2018.
- Alejandro Pérez Carballo. Negation, Expressivism, and Intentionality. *The Philosophical Quarterly*, 70(279):246–267, 2020.
- Caleb Perl. A User’s Guide to Hybrid Tools. *Mind*, 129(513):129–158, 2020.
- David Plunkett and Timothy Sundell. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers’ Imprint*, 13:1–37, 2013.
- Brian Rabern. Against the Identification of Assertoric Content with Compositional Value. *Synthese*, 189:75–96, 2012.
- Michael Ridge. Saving the Ethical Appearances. *Mind*, 115(459):633–650, 2006a.
- Michael Ridge. Ecumenical Expressivism: Finessing Frege. *Ethics*, 116:302–336, 2006b.
- Michael Ridge. *Impassioned Belief*. Oxford University Press, New York, 2014.
- Gideon Rosen. Blackburn’s *Essays in Quasi-Realism*. *Noûs*, 32(3):386–405, 1998.
- Bertrand Russell. *Religion and Science*. Oxford University Press, New York, 1935.
- Mark Schroeder. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford University Press, Oxford, 2008a.
- Mark Schroeder. What is the Frege-Geach Problem? *Philosophy Compass*, 3(4):703–720, 2008b.
- Mark Schroeder. How Expressivists Can and Should Solve Their Problem with Negation. *Noûs*, 42(4):573–599, 2008c.
- Mark Schroeder. Hybrid Expressivism: Virtues and Vices. *Ethics*, 119:257–309, 2009.
- Mark Schroeder. *Noncognitivism in Ethics*. Routledge, New York, 2010.
- Mark Schroeder. Tempered Expressivism. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press, New York, 2013.
- Mark Schroeder. Attitudes and Epistemics. In *Expressing Our Attitudes*, volume 2. Oxford University Press, New York, 2015a.
- Mark Schroeder. *Expressing our Attitudes: Explanation and Expression in Ethics*, volume 2. Oxford University Press, New York, 2015b.
- Andrew Sepielli. Normative Uncertainty for Non-Cognitivists. *Philosophical Studies*, 160: 191–207, 2012.

- Alex Silk. How to Be an Ethical Expressivist. *Philosophy and Phenomenological Research*, 90 (3), 2014.
- Alex Silk. Normative Language in Context. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 12. Oxford University Press, New York, 2015.
- Alex Silk. How to Embed an Epistemic Modal: Attitude Problems and Other Defects of Character. *Philosophical Studies*, 174:1773–1799, 2017.
- Michael Smith. *The Moral Problem*. Blackwell, Oxford, 1994.
- Robert Stalnaker. Assertion. In *Syntax and Semantics*, pages 315–332. New York Academic Press, New York, 1978.
- Tamina Stephenson. Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.
- C.L. Stevenson. *Ethics and Language*. Yale University Press, New Haven, 1944.
- Michael Stocker. Desiring the Bad: An Essay in Moral Psychology. *The Journal of Philosophy*, 76:738–753, 1979.
- Timothy Sundell. Disagreement about Taste. *Philosophical Studies*, 39(3):359–396, 2011.
- Sigrun Svavarsdóttir. Moral Cognitivism and Motivation. *The Philosophical Review*, 108: 161–219, 1999.
- Teemu Toppinen. Believing in Expressivism. In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press, New York, 2013.
- Nicholas Unwin. Quasi-Realism, Negation, and the Frege-Geach Problem. *The Philosophical Quarterly*, 49(196):337–352, 1999.
- Mark van Roojen. Expressivism and Irrationality. *The Philosophical Review*, 105(3), 1996.
- Mark van Roojen. Moral Rationalism and Rationalist Amoralism. *Ethics*, 120:495–525, 2010.
- Seth Yalcin. Epistemic Modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012.