

# NONCOGNITIVISM WITHOUT EXPRESSIVISM

## Abstract

According to expressivists, moral language does not try to describe the world; rather, it expresses desire-like states of mind. According to noncognitivists, moral beliefs are not representational mental states; rather, they have a desire-like direction of fit. 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: we should distinguish between moral semantics and moral psychology. Once we do, we open up space for a new metaethical position, which combines a noncognitivist moral psychology with a descriptivist semantics for moral language. After developing this position, I advertise its virtues. The noncognitivist aspect of the theory can be implemented by making a modest adjustment to a standard semantics for belief reports. This adjustment allows us to capture all of the explanatory benefits standardly associated with expressivism. At the same time, the descriptivist semantics allows us to avoid the Achilles' heel of expressivism: the Frege-Geach Problem. We get all of the advantages of expressivism without the headaches.

Word Count: 10,071 (12,345 with footnotes)

## 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 major 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.<sup>1</sup>

To hear many metaethicists tell it, these aren't two distinct faultlines; it's just the same one twice described. It's usually thought that expressivism goes hand-in-hand with noncognitivism and that descriptivism goes hand-in-hand with cognitivism. And these views do make for natural pairings. If to believe that stealing is wrong is to be in some desire-like state, then it's natural to think that in asserting, 'Stealing is wrong' one expresses this very state. If instead moral belief aims to accurately represent the world, then it's natural to think that moral assertions also aim to accurately represent the world.

Natural as these pairings may be, this paper aims to pull them apart. Specifically, I argue for *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, truth conditional semantics for moral language.<sup>2</sup>

At first blush, this may seem like an ungainly combination, with nothing to recommend it other than a perverse desire to stake out unoccupied regions of logical space. However, I'll argue that descriptivist noncognitivism has major advantages. 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, the descriptivist semantics avoids the main hurdle facing expressivism: the Frege-Geach Problem. We get the joys of expressivism without the tears.

Here's the game plan. Since descriptivist noncognitivism is an unfamiliar position, I start with a detailed implementation (§2). I show that, far from being *ad hoc*, the view can be implemented by combining an off-the-shelf descriptivist semantics for moral expressions with a simple modification of a standard semantics for belief reports. §§3-4 advertise the

<sup>1</sup>The early 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 2019.

<sup>2</sup>In the current literature, only Horgan and Timmons 2006 have explored the idea that one's moral psychology could come apart from one's theory of moral language. They defend a 'cognitivist expressivism', which combines a cognitivist moral psychology with an expressivist semantics. However, a number of authors have questioned whether the view they develop really earns the 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 my purposes, I won't take a stand on this issue. Regardless of what one thinks of cognitivist expressivism, I'll argue that its flipside—descriptivist noncognitivism—not only earns the title; it can be motivated on principled grounds.

advantages of the resulting position. §3 highlights the benefits that come from its descriptivist element; §4 trumpets the benefits that come from its noncognitivist component. §5 considers objections to my view and discusses some choice points in its implementation. I conclude (§6) that descriptivist noncognitivism is not just a theoretical novelty, but a serious contender in the metaethical landscape.

## 2 Developing Descriptivist Noncognitivism

In order to implement my view, I start by reviewing a standard descriptivist semantics for moral language. I then show that by combining this semantics with a new twist on an old 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.

The main reason for this choice of focus is that there is a general consensus about how to analyze such modal constructions, at least in broad outlines. The standard approach is due to Kratzer 1981, 1991, 2012. On Kratzer’s analysis, the extension of any modal depends on two parameters, both of which are determined by the context of utterance. 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. In the case of moral uses of modals (and deontic modals more generally), it is typically assumed that the modal base is circumstantial (Kratzer 1991; Hacquard 2011): for any world  $w$ ,  $f(w)$  delivers a set of propositions describing various circumstances that obtain at  $w$ . The second parameter is an ordering source  $g$ , a function from a world to a set of propositions that induces a ranking over worlds. A necessity modal (represented here as ‘ $\Box$ ’) is analyzed as a universal quantifier over the  $g$ -best of the  $f$ -accessible worlds. Formally:

**Contextualist Ought**  $\llbracket \Box \phi \rrbracket^{f,g,w} = 1$  iff  $\forall v \in \text{BEST}_{g(w),f(w)} : \llbracket \phi \rrbracket^{f,g,v} = 1$ ,

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

For example, (1) says that in all of the morally best of the circumstantially similar worlds, people give to charity.

This semantics does not take a stand on what makes one world morally better than another. It is compatible with realist views, according to which the ordering source reflects the objective, mind-independent moral facts. It is also compatible with subjectivist views, according to which the ordering source reflects the conative attitudes of the speaker, as well as various group contextualist views, according to which the ordering source reflects the conative attitudes shared by the conversational participants. For present purposes, we need not commit ourselves to any particular interpretation of the ordering source. In due course we will examine some constraints that bear on this choice.

Whichever way we go on this metaethical question, this semantics should be regarded as descriptivist, not expressivist. In order to make this point, it will be helpful to sharpen our understanding of the descriptivism/expressivism distinction. Admittedly, this is a tricky task, since different authors use these labels in different—and sometimes incompatible—ways. So no way of drawing this distinction will win universal agreement. That said, I take it that the core idea behind descriptivism is roughly this: 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.

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  $\phi$  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.<sup>3</sup> And a psychologistic semantics is typically regarded as an *alternative* to a truth conditional semantics.

<sup>3</sup>See e.g., Rosen 1998: 387-388; Wedgwood 2007: 5; Schroeder 2008a, 2010; Charlow 2015.

Admittedly, not all expressivists agree on this last point. For example, both [Yalcin 2011, 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. Consequently, they deny that the contents of moral utterances can be modeled with sets of worlds.

We are now in a position to see why Contextualist *Ought* is a descriptivist semantics. The crucial point is that while Contextualist *Ought* postulates non-worldly parameters—an ordering source and a modal base—the values of these parameters are determined by the context of utterance. And so the semantic value of any moral sentence, relative to such a sequence of parameters, will be a set of worlds. To make this a bit more concrete, suppose that (1) is uttered by a particular agent, Ana, in a particular context. According to Contextualist *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. Consequently, Ana’s utterance of (1) has representational truth conditions. It is true at a world  $w$  if  $w$  meets this condition, false otherwise.

## 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:

- (2) 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. Let’s take this step by step.

The leading analysis of belief reports in the semantics literature, due to [Hintikka 1962](#), analyzes *believes* as 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 Contextualist *Ought*, this semantics leads to a form of cognitivism. On the resulting combination, (2) says that Ana believes that all of the morally best of the accessible worlds are worlds where people give to charity. Now, what it means for her to have this belief depends on how one interprets the moral ordering source. But however we resolve this interpretative question, (2) will ascribe Ana a mental state with a mind-to-world direction of fit. And 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, 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  $\alpha$ 's intrinsic desire function ( $\delta_\alpha$ ), a function from a world  $w$  to a set of propositions representing  $\alpha$ '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:

$$\delta_{Ana}(w) = \{\text{wellbeing is promoted}\}.$$
<sup>4</sup>

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  $\alpha$ 's intrinsic desires (at  $w$ ) than  $u$  does.<sup>5</sup> Given this, we could propose that *believes* shifts the ordering source to the believer's intrinsic desire function:

<sup>4</sup>Insofar as we are after an analysis of moral beliefs (as opposed to normative beliefs more generally), we may 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.

<sup>5</sup>How 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  $u$ .

**Shifty Believes**  $\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,w} = 1$  iff  $\forall v \in \text{Dox}_\alpha^w : \llbracket \phi \rrbracket^{f,\delta_\alpha,v} = 1$ .

This 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 (2) using Contextualist *Ought* and Shifty *Believes*:

$$\llbracket (2) \rrbracket^{f,g,w} = 1 \text{ iff } \forall v \in \text{Dox}_{\text{Ana}}^w : \forall u \in \text{BEST}_{\delta_{\text{Ana}}(v),f(v)} : \text{people give to charity at } u.$$

In words: (2) 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 (2) 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 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 ( $\delta_\alpha^w$ ) be a constant function from an arbitrary world  $u$  to a set of propositions representing  $\alpha$ '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:

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

Noncognitivists can propose that belief reports shift the ordering source to the agent's world-indexed desire function:

**Noncognitivist Believes**  $\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,w} = 1$  iff  $\forall v \in \text{Dox}_\alpha^w : \llbracket \phi \rrbracket^{f,\delta_\alpha^w,v} = 1$ .

On this view, (2) 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.

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 well-being. 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:

- (3) Ana believes she ought to save the child.
- (4) Ana believes she ought to let the child drown.

Shifty *Believes* predicts that (3) is false and (4) 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 (3) is true and (4) 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 Shifty *Believes* and Noncognitivist *Believes* make different predictions about what Ana believes in this scenario. They also make different predictions about what constitutes a good reason for holding a moral belief. Consider: when Ana's psychoanalyst tells her that she only cares about accruing material goods, does this provide Ana with a reason to believe that she ought to let the child drown? Shifty *Believes* says 'Yes'. After all, this testimony is evidence in favor of the proposition: *your (Ana's) desires would be best satisfied by letting the child drown*. According to Shifty *Believes*, for Ana to believe she ought to let the child drown is just for her to have this belief about her own desires. But this prediction seems incorrect. Perhaps the psychoanalyst's testimony gives her a reason to believe that *she believes* that she ought to let the child drown. But, intuitively, it provides no reason (or, at least, no good reason) to believe that she ought to do so. Noncognitivist *Believes* avoids this consequence.



According to Noncognitivist *Believes*, the reasons for having a moral belief are not reasons to believe one has a certain desire. Rather, they are reasons to have this desire in the first place.<sup>6</sup>

### 2.3 Precedents and extensions: epistemic modals

To my knowledge, no one has defended Noncognitivist *Believes*. However, the basic idea that attitude verbs shift the parameters relative to which modals are evaluated has already gained considerable currency in the literature on epistemic modals.

To motivate this, consider a simple contextualist semantics for epistemic *might* ( $\diamond$ ), on which *might* existentially quantifies over the worlds in an epistemic modal base (a function from a world to a contextually relevant body of information):

**Contextualist *Might***  $\llbracket \diamond \phi \rrbracket^{f,g,w} = 1$  iff  $\exists v \in \bigcap f(w) : \llbracket \phi \rrbracket^{f,g,v} = 1$ .

As a number of writers have observed, combining Contextualist *Might* with Hintikka *Believes* makes implausible predictions about belief reports embedding epistemic modals:

- (5) Ana believes that it might be snowing.

On the resulting combination, (5) says that Ana believes that there is a world compatible with the contextually determined information where it is snowing. But there are two related difficulties with this analysis. The first is that (5) could be true in a situation where Ana has no views about which possibilities are compatible with the speaker's information, or the conversational context more generally. The second problem is that, as Yalcin 2011 notes, this seems to misdiagnose the subject matter of modal beliefs. Ana's belief seems to be about whether it is snowing, not about whether the proposition that it's snowing is compatible with a certain body of information.

<sup>6</sup>Noncognitivist *Believes* also differs from hybrid views of moral belief, of the sort defended by e.g., Ridge 2006b, 2007, 2014; Boisvert 2008; Hay 2013; Laskowski 2019; Perl 2020. According to Noncognitivist *Believes*, the truth of (2) does not depend in any way on the content of Ana's descriptive beliefs. For example, we do not require that Ana has any beliefs about her intrinsic desires, or about which worlds are accessible from her doxastic alternatives. More generally, our view predicts that anyone who has any descriptive beliefs at all, together with an intrinsic desire function, will have various moral beliefs.

This distinguishes our view from hybrid theories. On the hybrid picture, moral beliefs do not require just any old descriptive belief, paired with some desire-like state. Rather, they require descriptive beliefs of a very specific sort. For example, consider the advisor-based view defended in Ridge 2006b, 2007. On this view, to believe that people ought to give to charity is to both: (i) have a descriptive belief that any suitable advisor would approve of people giving to charity, (ii) approve of such an advisor. Clearly, our view makes no such requirements. Ana can have a moral belief without having any opinions about what any advisor would recommend.

As various theorists have pointed out, a natural solution is to let *believes* shift the modal base. Here is one way of developing this idea, due to [Ninan 2018](#).<sup>7</sup> Let an agent  $\alpha$ 's belief function at  $w$  ( $s_\alpha^w$ ) be a constant function from an arbitrary world to  $\alpha$ 's doxastic alternatives at  $w$ . Ninan proposes that *believes* shifts the modal base to the believer's belief function:

**Modal Base Shift**  $\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,w} = 1$  iff  $\forall v \in \text{Dox}_\alpha^w : \llbracket \phi \rrbracket^{s_\alpha^w, g, v} = 1$ .

On the resulting approach, (5) does not require that Ana has any beliefs about which possibilities are consistent with any body of information—be it her own information or the speaker's. All that is required is for her doxastic alternatives to include at least one world where it's snowing.

For my purposes, nothing hinges on beliefs embedding epistemic modals. So I will ignore any effect that *believes* has on the modal base going forward, and stick with a view on which *believes* only shifts the ordering source. The important point is that the sort of maneuver I am defending here is not without precedent, since it closely parallels moves that have been defended on independent grounds.<sup>8</sup>

## 2.4 Taking Stock

At first blush, descriptivist noncognitivism will strike many as a bizarre position, unmotivated at best, incoherent at worst. In this section, I've sought to dispel this impression. The view is perfectly coherent, and it can be given a straightforward compositional implementation. The general strategy was to combine a familiar descriptivist semantics for moral discourse (Contextualist *Ought*) with an ordering source-shifting semantics for belief reports (Noncognitivist *Believes*). This semantics for belief reports is relatively conservative, in that it makes a small modification to the leading treatment of attitude verbs in the literature. Moreover, analogous modifications have been independently proposed for handling embedded epistemic modals.

<sup>7</sup>Ninan focuses on supposition ascriptions, but the extension to belief is straightforward. For other proposals on which *believes* can shift the body of information that an embedded epistemic modal quantifies over, see [Stephenson 2007](#); [Yalcin 2007](#); [Hacquard 2010](#); [Silk 2017](#).

<sup>8</sup>Of course, one could hold the view that *believes* shifts both the modal base and the ordering source, i.e.:

$$\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,w} = 1 \text{ iff } \forall v \in \text{Dox}_\alpha^w : \llbracket \phi \rrbracket^{s_\alpha^w, \delta_\alpha^w, v} = 1.$$

However, this predicts that belief reports embedding moral uses of modals always quantify over the believer's doxastic alternatives. And this seems doubtful when it comes to reports like the following:

- (i) Ana believes the US ought not have invaded Iraq.

This could be true even if all of Ana's doxastic alternatives are worlds where the US invaded Iraq. So I am tentatively inclined to conclude that *believes* does not always shift the modal base to the believer's belief function. It only does so—and perhaps optionally—when the embedded modal is epistemic.

This implementation is offered as a ‘proof of concept’: it is one way to develop descriptivist noncognitivism, not the only way. For example, it is also possible to implement the view in a relativist setting. And this relativist implementation may carry certain advantages when it comes to responding to potential objections. However, so as not to hold up the action, I will defer this relativist implementation to the Appendix.

Having shown how to implement descriptivist noncognitivism, I now turn to consider the main advantages of the view. §3 argues that the descriptivist element of the view avoids the chief objection to expressivism: the Frege-Geach Problem. §4 shows that the noncognitivist element preserves all of the explanatory virtues associated with expressivism.

### 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.<sup>9</sup> 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:

- (6) It’s not the case that people ought to give to charity.
- (7) Charity alleviates suffering and people ought to give.
- (8) (Either) charity does not alleviate suffering or people ought to give.
- (9) If charity alleviates suffering, then 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 (6)-(9). 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:

---

<sup>9</sup>The *locus classicus* of the problem is [Geach 1964](#). For an overview, see [Schroeder 2008b](#).

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.

To their credit, expressivists have not left this challenge unanswered. The last forty years have seen no shortage of attempts to solve both facets of the Frege-Geach Problem.<sup>10</sup> While these approaches are often ingenuous, empirical and explanatory difficulties persist, and it remains very much an open question whether any of these approaches provides a fully adequate solution. By contrast, descriptivist noncognitivism avoids the Frege-Geach Problem entirely. After all, Contextualist *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.<sup>11</sup>

<sup>10</sup>See Blackburn 1984, 1988; Gibbard 1990, 2003; Horgan and Timmons 2006; Ridge 2006b, 2014; Schroeder 2008a; Silk 2014; Charlow 2015; Starr 2016; Willer 2017, among many others.

<sup>11</sup>A detailed comparison with particular expressivist views is outside the scope of this paper. However, it may be helpful to make a few remarks about how our framework differs from Gibbard's expressivism, which is arguably the closest expressivist cousin to the view developed here. Gibbard's key insight is that just as we can apply Boolean operations to sets of worlds, we can also apply them to sets of  $w, \pi$  pairs, where  $\pi$  is a set of norms (Gibbard 1990) or a hyperplan (Gibbard 2003). This allows Gibbard to adopt a Boolean analysis of complex constructions such as (6)-(9). Similarly, Gibbard can follow possible worlds semanticists in analyzing truth in terms of set membership, where the sets are now composed of  $w, \pi$  pairs.

While this is an elegant framework, a number of philosophers have worried that it is explanatorily inadequate. Any adequate semantics for normative language will explain why (1) and (6) are inconsistent. Now, Gibbard has a 'cheap' explanation of this: (1) and (6) express sets of  $w, \pi$  pairs with an empty intersection. But, the complaint runs, this cannot be the full story. After all, Gibbard is still committed to a psychologistic semantics: he aims to explain the semantic properties of normative sentences in terms of properties of the states of mind that these sentences express. So it is incumbent on Gibbard to tell us what states of mind (1) and (6) express, and why these states of mind are inconsistent. Gibbard 2003 tries to discharge this burden by appealing to a primitive notion of disagreement. Roughly, (1) expresses the state of planning to give to charity, and (6) expresses a state of mind that disagrees with planning to give to charity. However, this has left a number of commentators dissatisfied, since it seems tantamount to stipulating the existence of mental states with the desired properties. (For elaboration of this objection, see Dreier 2006, 2009: §5; Schroeder 2008a: chp.3, 2008c: §3.2; 2010: 131-133; 2015; Willer 2017: 177-179. For responses, see Silk 2014; Baker and Woods 2015; Carballo 2020.)

Descriptivist noncognitivists avoid this difficulty. First, they can simply say that (1) and (6) are inconsistent (provided they are uttered against the background of the same contextual parameters) because these assertions make inconsistent claims about what the world is like. Second, descriptivist noncognitivists are not encumbered by the explanatory commitments of a psychologistic semantics. They make no claim that moral sentences express states of mind at all, let alone that their semantic properties are to be explained in terms of the states of mind they express.

### 3.2 Revenge of the Frege-Geach Problem?

Some may think that this is too quick. Certain versions of the Frege-Geach Problem involve moral belief reports. Perhaps, some might suggest, these versions of the problem reveal a structural shortcoming in any noncognitivist theory of moral belief, regardless of whether that theory is paired with an expressivist or a descriptivist semantics.

Consider the negation problem for belief (Unwin 1999; Dreier 2006; Schroeder 2008a; Carballo 2020). This problem arises from the fact that there are three different places where one can insert negation into a moral belief report:

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

(10)–(12) 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

Some might worry that this version of the Frege-Geach challenge afflicts all noncognitivists, descriptivists or not.

However, if we work through what our semantics (Contextualist *Ought* + Noncognitivist *Believes*) says about (10)–(12), we see that it nicely captures the differences between these three states. In doing so, it yields a new solution to the negation problem.

To state the solution, it will help to introduce some terminology. Say that  $\alpha$  *decidedly desires*  $p$  at a world  $w$  if and only if, for all of  $\alpha$ 's doxastic alternatives  $v$ , all of the most desired of the  $v$ -accessible worlds (by the lights of  $\delta_\alpha^w$ ) are worlds where  $p$  holds. And say that  $\alpha$  *decidedly does not desire*  $p$  at  $w$  if and only if, for all  $\alpha$ 's doxastic alternatives  $v$ , **not** all of the most desired of the  $v$ -accessible worlds (by the lights of  $\delta_\alpha^w$ ) are worlds where  $p$  holds. If

we compute the truth conditions of (10)-(12) using Contextualist *Ought* and Noncognitivist *Believes*, we get the equivalences in Table 2.

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

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 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.)

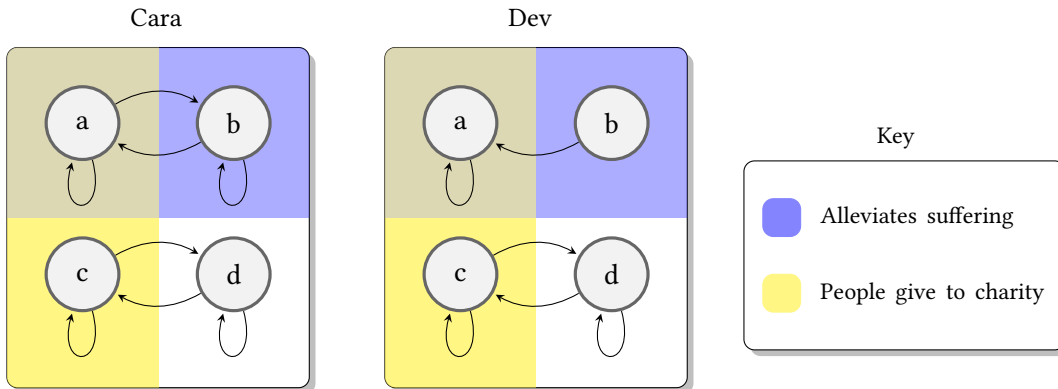


Figure 1: Two states of mind. (Arrows indicate the most desirable of the accessible worlds.)

This solution to the negation problem avoids challenges that arise for other solutions

that have been proposed in the literature. For example, one alternative solution maintains that (11) 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 (Schroeder 2008a,c). 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 (10)–(12): (10) says Benny desires praising for not giving to charity; (11) says Cara desires not praising for giving; (12) says Dev does not desire praising for giving. While this is in many respects an elegant solution, it makes substantive assumptions about the structure of moral judgment—assumptions whose plausibility one might well question. In order for (11) 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.<sup>12</sup>

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:

- (13) 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: (13) is true iff 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.<sup>13</sup>

<sup>12</sup>Another approach to the negation problem is suggested by Gibbard 2003. On one way of developing Gibbard's approach, (10) says that Benny plans to not give to charity, (11) says that Cara is in a mental state that disagrees with planning to give, and (12) says that Dev is in a mental state that does not disagree with planning to not give. This approach raises concerns—developed in Dreier 2006 and broached in the previous footnote—about whether we are entitled to help ourselves to an unreduced notion of disagreement in characterizing the relevant mental states. By contrast, the account developed here characterizes these belief states directly, without appeal to disagreement. (For discussion of how descriptivist noncognitivists should make sense of disagreement, see §4.4.)

<sup>13</sup>What about beliefs involving conditionals? There are a variety of possible worlds semantics for conditionals that could be combined with descriptivist noncognitivism. Given our reliance on a Kratzer's semantics for modals, one natural option would be to adopt Kratzer's restrictor analysis of conditionals; see Kratzer 1986, 2012.

### 3.3 Semantic Simplicity

I've argued that descriptivist noncognitivism bypasses the Frege-Geach Problem. In my eyes, this is a significant advantage. After all, the Frege-Geach Problem has been a thorn in the expressivist's side for nearly sixty years. And while there have been valiant attempts to remove the thorn, it remains unclear whether any of these extraction efforts has fully succeeded.

Moreover, even if some version of expressivism can handle the Frege-Geach Problem (a big if!), parsimony considerations may still favor descriptivist noncognitivism. The development of possible worlds semantics stands out as one of the great success stories in philosophy of language and linguistics over the last hundred years. The basic framework is simple, elegant, and explanatorily fruitful. By contrast, most of the expressivist frameworks proposed to date involve additional complexities—complexities that have no discernible empirical motivation.

Expressivists may take issue with this last claim. Even if we cannot find any linguistic data that favors metaethical expressivism over descriptivism, expressivism promises to solve long-standing metaethical puzzles. If solving these puzzles requires some additional semantic complexity, so be it!

However, once we examine these explanatory benefits in more depth, we will find that they are really benefits of a noncognitivist psychology rather than an expressivist semantics. Descriptivist noncognitivism allows us to enjoy these benefits without abandoning the possible worlds framework that has served us so well.

## 4 Noncognitivist Benefits

Despite the objections that have been levied against it, expressivism holds enduring appeal. This is due to the fact that expressivism promises to deliver four major metaethical benefits:

- *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.

#### 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 descriptivist noncognitivism 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. To see this, recall that Contextualist *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.

In response, it will be helpful to recall a point noted in §2.1. Contextualist *Ought* is compatible with any number of metaethical accounts of the moral ordering source, including naturalistic accounts. For example, it is compatible with a naturalist version of realism, according to which the moral ordering source is provided by some body of natural facts. And it is also compatible with various subjectivist and contextualist proposals, according to which the moral ordering source is provided by the conative attitudes of the speaker, or the conversational interlocutors, or some larger community.

Expressivists may retort that such naturalist realist or subjectivist views are 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 phe-

nomenon, 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:

(14) Charitable giving promotes wellbeing.

Moore invites us to compare the questions:

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

(16) Granted that we ought to give to charity, but ought we give to charity?

According to Moore, (15) is an ‘open’ question, whereas (16) is not. Moore took this to show that (15) and (16) differ in meaning, and hence that (1) and (14) 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 (14) differ in meaning because (14) 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.<sup>14</sup>

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).<sup>15</sup>

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  $\phi$ ing is  $N$  without believing that they (or anyone else) morally ought to  $\phi$ .

<sup>14</sup>See Ayer 1936: chp. 6; Hare 1952: 92-93. See Horgan and Timmons 1992 for a variant of this argument.

<sup>15</sup>See Harman 1977; Smith 1994 for related worries.

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  $\phi$ ing is  $N$  without desiring  $\phi$ ing, and hence without believing people ought to  $\phi$ . 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.<sup>16</sup>

### 4.3 Motivational Internalism

Many metaethicists have been attracted to the ‘internalist’ idea that moral judgments are inherently action-guiding.<sup>17</sup> 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.<sup>18</sup> 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  $\phi$ , then they will be at least somewhat disposed to try to  $\phi$ .

<sup>16</sup>While 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 believes that  $\phi$ ing is the only way to bring about  $p$ , then they cannot coherently wonder whether one ought to  $\phi$ . 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.4, 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.

<sup>17</sup>This idea traces back to Falk 1947. For contemporary defenses, see Dreier 1990; Smith 1994; Korsgaard 1996; Blackburn 1998; Gibbard 2003; Wedgwood 2007; Egan 2012, among many others.

<sup>18</sup>For discussion, see Stocker 1979; Smith 1994; Mele 1996; Svavarsdóttir 1999.

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.<sup>19</sup> By contrast, descriptivists are often thought to be at an explanatory disadvantage. After all, the worry runs, if claiming that one morally ought to  $\phi$  is just to claim that  $\phi$ ing has some descriptive property  $F$ , why should we expect this claim to carry any motivational oomph?

But once we are careful to distinguish between expressivism and noncognitivism, we see that expressivism does not—taken by itself—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 functional role of desire is to motivate action: desiring  $\phi$  disposes an agent to try to  $\phi$ . So if Ana believes everyone ought to give to charity, it follows that Ana will be disposed to try to make everyone give. And this would involve trying to give 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 1944](#):

Suppose that two people have decided to dine together. One suggests a restaurant where there is music; another expresses his disinclination to hear music

<sup>19</sup>Virtually all of the major defenses of expressivism appeal to some version of Motivational Internalism. See e.g., [Stevenson 1944](#); [Blackburn 1998](#); [Gibbard 1990, 2003](#). See also [Darwall et al. 1992](#) for the view that the best version of the Open Question Argument should appeal to the motivational profile of moral beliefs.

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 for two people to ‘disagree in attitude’? While the phenomenon is surely familiar, it has proved challenging to provide a theory of disagreement in attitude that avoids either overgenerating or undergenerating disagreements (cf. Dreier 2006). However, I think that our solution to the Negation Problem provides just what is needed. Recall our terminology from §3.2:  $\alpha$  decidedly desires  $p$  at  $w$  iff for all of  $\alpha$ ’s doxastic alternatives  $v$ , all of the  $v$ -accessible worlds ranked highest by  $\delta_\alpha^w$  are worlds where  $p$  holds. And  $\alpha$  decidedly does not desire  $p$  at  $w$  iff for all  $\alpha$ ’s doxastic alternatives  $v$ , not all of the  $v$ -accessible worlds ranked highest by  $\delta_\alpha^w$  are worlds where  $p$  holds. We can propose:

**Disagreement in Attitude**  $\alpha$  and  $\beta$  disagree in their desires regarding  $p$  iff  $\alpha$  decidedly desires  $p$  and  $\beta$  decidedly does not desire  $p$ , or *vice versa*.

This nicely captures our intuitions about the disagreements—or lack thereof—that obtain 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 Kuala Lumpur is north of Singapore, and Dev has no opinion about their locations, it seems wrong to describe them as disagreeing in their descriptive beliefs.

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:

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

1.  $\llbracket \alpha$  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  $\alpha$  and  $\beta$  disagree over whether people ought to give to charity iff  $\alpha$  believes people ought to give to charity, and  $\beta$  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  $\alpha$  decidedly desires people to give to charity, and  $\beta$  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.

The upshot: combining Contextualist *Ought* with Noncognitivist *Believes* yields a new theory of moral disagreement—one that has considerable explanatory value. Once again, all of the explanatory value comes from the noncognitivist theory of moral belief. Nothing in the story presupposed expressivism.<sup>20</sup>

#### 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.

<sup>20</sup>An interesting question is whether expressivists could take on board the analysis of disagreement developed here. I think the answer will depend on the details of one's expressivist theory. The account of disagreement developed here—much like our earlier solution to the negation problem—only works because decidedly desiring involves two levels of quantification: we quantify over all of the agent's doxastic alternatives  $v$ , and also over all of the most desirable of the  $v$ -accessible worlds. These two layers of quantification follow immediately from Noncognitivist *Believes*, when combined with Contextualist *Ought*. By contrast, no extant expressivist semantics for belief reports has quite this structure. That said, I see no reason in principle why one could not devise an expressivist semantics for *ought* and *believes* that mimics these predictions. Such an expressivist account would be interesting in its own right, and would have significant advantages over expressivist views that simply help themselves to an unreduced notion of disagreement (see fns. 10 & 11). However, from the perspective of this paper, it is unclear why one would pursue this expressivist variant of our theory. If I'm right, descriptivist noncognitivism retains all of the advantages associated with expressivism while hewing to the most successful semantic framework around. Why overhaul our semantic framework only to mimic its predictions?

## 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 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  $\lceil \text{Ought } p \rceil$  is to get their audience to believe *ought*  $p$ .

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  $\phi$ , in a context with a sequence of contextually determined parameters  $c$ , is to get their audience to believe  $\llbracket \phi \rrbracket^c$ .<sup>21</sup>

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

<sup>21</sup>See [Stalnaker 1978](#) for an influential early statement of this view.

between moral talk and moral thought, we have trouble explaining how speakers use the former to influence the latter.<sup>22</sup>

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  $\lceil \text{Ought } p \rceil$ , they will thereby come to believe *ought*  $p$ . The second response involves rejecting Assertoric Update in favor of a pragmatic theory that plays better with our framework. Let us take each in turn.

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 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.<sup>23</sup> 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:

- (17) 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

<sup>22</sup>This problem is not specific to moral language. A similar worry arises for the treatment of epistemic modals discussed in §2.4, which combines a contextualist semantics for epistemic modals (Contextualist *Might*) with a shifty semantics for belief reports (Modal Base Shift).

<sup>23</sup>Mandelkern 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 1979b](#); [Björnsson and Finlay 2010](#); [Plunkett and Sundell 2013](#); [Silk 2015](#); [Khoo and Knobe 2018](#).



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.

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 her ultimate aim is 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 for descriptivist noncognitivists to secure thought-talk connections. According to this response, someone who asserts  $\lceil \text{Ought } p \rceil$  in a context *c* aims to get their audience to believe  $\llbracket \text{Ought } p \rrbracket^c$ —just as Assertoric Update has it. However, achieving this aim will normally—and foreseeably—have the consequence that their audience will also believe ought *p*. In order to secure this result, we relied on Prospective Contextualism. However, there are other possible route to the same conclusion. We could also make much the same argument in a relativist setting; see Appendix A.3 for discussion.

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.<sup>24</sup> 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

<sup>24</sup>For more on this distinction, see Lewis 1980; Yalcin 2007; Ninan 2010; Rabern 2012.

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  $\phi$  is to get one's audience to believe  $\llbracket\phi\rrbracket$ .

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.

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 iff 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:

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

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 (18) 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 (18) guarantees the falsity of the first.<sup>25</sup>

An alternative response is to scrutinize whether (18) 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:

- (19) 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, (19) provides evidence that it's possible for two speakers to disagree despite neither of them speaking falsely.

In fact, recent 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.<sup>26</sup> 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 (18) countenances.

<sup>25</sup>Here too, Prospective Contextualism is not the only way to go. We can achieve a similar result in a relativist setting. See Appendix A.4 for discussion.

<sup>26</sup>The 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.

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 (18). This is an 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.

## 6 Conclusion

This paper introduced and defended a new metaethical position: 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. It can be given a precise compositional implementation in a possible worlds semantics. And it avoids the main problems facing expressivism while reaping all of the benefits. In light of these advantages, I urge that this view deserves a seat at the metaethical table.

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 introducing a view like descriptivist noncognitivism is that it helps us resist this natural—but ultimately mistaken—line of argument.

## Appendix: Relativist Noncognitivism

In order to integrate descriptivism with noncognitivism, I relied on a contextualist semantics. But a similar integration can be achieved in a relativist setting. In this appendix, I show how this can be done. I then explore the consequences of this relativist alternative for the issues discussed in §5.

### A.1 A relativist semantics

According to moral relativists, moral utterances take on different truth-values relative to different assessors. For example, Ana's utterance of (1) might be true for Ana, but false for Benny.<sup>27</sup> To implement this idea, we need only make a couple of adjustments to our con-

---

<sup>27</sup>For a general overview and defense of relativism, see [MacFarlane 2014](#). For a defense of relativism about moral claims, see e.g., [Egan 2012](#).

textualist semantics. First, we enrich our circumstances of evaluation: they are no longer worlds, but rather centered worlds  $\langle w, x \rangle$ —that is, ordered pairs of a world and an assessor. Second, we take the ordering source to be a function from centered worlds to sets of propositions:

**Relativist Ought**  $\llbracket \Box \phi \rrbracket^{f,g,\langle w,x \rangle} = 1$  iff  $\forall v \in \text{BEST}_{g(w,x),f(w)} : \llbracket \phi \rrbracket^{f,g,\langle v,x \rangle} = 1$ .

Like contextualism, relativism is officially compatible with different conceptions of the moral ordering source. However, in order to generate concrete predictions, it will be helpful to get a particular view of the ordering source on the table:

**Relativist Ordering Source** For moral uses of deontic modals, the ordering source is the intrinsic desire function of the assessor.

On this view, Ana’s utterance of (1) is true for her (that is, true at the centered world centered on her) since her intrinsic desires give high marks to charitable giving. But her utterance is false for Benny, since his desires assign low marks thereto.

Is this semantics descriptivist? Not according to our earlier definition of descriptivism (§2.1). After all, the relativist holds that the assertoric content of a moral sentence is not a set of worlds, but rather a set of centered worlds. However, I think that a plausible case can be made for thinking that these centered contents are still representational. It’s just that what they represent is one’s place in the world. For example, the assertoric content of (1) will be a set of centered worlds representing self-locating information about one’s desires—specifically, that all of the nearby worlds that come closest to conforming to one’s intrinsic desires are worlds where people give to charity.<sup>28</sup>

Of course, if we regard relativism as a type of descriptivism, then our earlier definition will have to be tweaked. But the tweak is straightforward enough: a semantics is descriptivist if, for any moral sentence  $\phi$  and any sequence of contextual parameters  $c$ , it says that  $\llbracket \phi \rrbracket^c$  is either a set of worlds or a set of centered worlds.<sup>29</sup>

<sup>28</sup>Here I follow Lewis 1979a in thinking of sets of centered worlds as representing self-locating or *de se* information. See also Stephenson 2007; Egan 2007, 2010, 2012.

<sup>29</sup>Some might wonder whether this revision is too permissive. After all, the relativist’s formal apparatus is quite similar to that of Gibbard-style expressivists. Both hold that assertoric contents are sets of ordered pairs of a world and some non-worldly entity. For the relativist, that entity is an assessor; for Gibbard, it is a norm or a hyperplan. So if we classify relativism as a form of descriptivism, is there any principled reason not to classify Gibbardian expressivism similarly? I think there is. The crucial difference is not in the formalism itself, but in how the formalism is interpreted. I’ve argued that the relativist’s assertoric contents are still representational. Gibbardian expressivists develop their approach as an alternative to a representational semantics. For Gibbard, an utterance of (1) *expresses* one’s desire-like states of mind. And Gibbard, like virtually all expressivists, maintains that there is an important difference between expressing a state of mind  $m$  and representing oneself as being in  $m$  (Gibbard 1986: 473; see also Schroeder 2010: chp.4).

## A.2 Belief reports for relativists

Next, consider how to analyze moral belief reports in a relativist setting. We will continue to operate with the Hintikkan assumption that *believes* universally quantifies over the believer's doxastic alternatives, now taken to be centered worlds. In order to reconcile relativism and noncognitivism, we can deploy a variant of our earlier reconciliation strategy (§2.2). As before, we let *believes* shift the ordering source in the index to the believer's world-indexed intrinsic desire function:

**Relativist Believes**  $\llbracket \alpha \text{ believes } \phi \rrbracket^{f,g,\langle w,x \rangle} = 1$  iff  $\forall \langle v, y \rangle \in \text{Dox}_\alpha^w : \llbracket \phi \rrbracket^{f,\delta_\alpha^w,\langle v,y \rangle} = 1$ ,  
 where  $\text{Dox}_\alpha^w = \{ \langle v, y \rangle \mid \text{Being } y \text{ at } v \text{ is compatible with } \alpha \text{'s beliefs at } w \}$ .

This semantics for belief reports (when combined with Relativist *Ought*) is equivalent to Noncognitivist *Believes* (when combined with Contextualist *Ought*). Both predict that (2) is true iff Ana decidedly desires that people give to charity.

So descriptivist noncognitivism can be implemented in either a contextualist or a relativist setting. Moreover, the choice between these implementations has implications for the objections facing descriptivist noncognitivism. To see this, let's revisit the objections from §5 from a relativist perspective.

## A.3 Thought-talk connections for relativists

In §5.1 we confronted the worry that descriptivist noncognitivism has trouble underwriting the connections between moral thought and moral talk. In response, we saw that one option would be to combine Prospective Contextualism with the idea that we normally have introspective access to our desires (Normal Access) to derive the result that normally anyone who updates on the assertoric content of a moral assertion will be in a desire-like state of mind.

A similar result can be achieved in a relativist framework. According to the relativist, the assertoric content of Ana's utterance of (1) is a set of centered worlds. Assuming Relativist Ordering Source, this will be the set of centered worlds such that the center's intrinsic desires give high marks to the accessible worlds where people give to charity. Now suppose a hearer *H* comes to believe this centered content. Then *H* will thereby come to self-attribute the property of being in this desire-like state of mind. By Normal Access, it follows that if conditions are normal, *H* is in this desire-like state of mind. So once again we get the result that anyone who updates on the assertoric content of (1) will, in normal circumstances, believe that people ought to give to charity.

#### A.4 Truth-value assessments for relativists

The next objection was that there is a tension between our truth conditions for moral assertions and our account of moral disagreement—a tension revealed by the oddity of:

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

As we saw in §5.2, Prospective Contextualism provides one way of explaining the apparent oddity of (18). Relativism provides another. If Benny believes that it's not the case that people ought to give to charity, then Ana's utterance of (1) is guaranteed to be false at his context of assessment.

Relativism and Prospective Contextualism thus agree in a number of their predictions. But they are not equivalent. Prospective Contextualism predicts that our moral utterances are false whenever we happen to be conversing with an interlocutor who does not share our desires—the dispute between Ana and Benny being a case in point. Some might find this consequence implausible. Relativism avoids this result. For the relativist, Ana's assertion is still true for her, despite Benny's contrary desires. And it is also true for *us*, provided we share Ana's desires. Another empirical difference concerns cases where eavesdroppers have desires that conflict the speaker's. (Imagine that Benny was not a part of Ana's conversation, but overhead her utterance. Is (18) any more felicitous?) Admittedly, the empirical data concerning eavesdroppers are rather murky, and a matter of ongoing controversy.<sup>30</sup> The present point is simply that our intuitions about such cases—whatever they may be—will help adjudicate between the two frameworks.

## References

- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- Derek Baker and Jack Woods. How Expressivists Can and Should Explain Inconsistency. *Ethics*, 125:391–424, 2015.
- Bob Beddor. Noncognitivism and Epistemic Evaluations. *Philosophers' Imprint*, 19(10):1–27, 2019.
- Bob Beddor and Andy Egan. Might do Better: Flexible Relativism and the QUD. *Semantics and Pragmatics*, 11, 2018.

<sup>30</sup>For empirical studies on eavesdropping intuitions involving epistemic modals, see [Knobe and Yalcin 2014](#); [Khoo 2015](#); [Beddor and Egan 2018](#); [Phillips and Mandelkern 2020](#).

- Gunnar Björnsson and Stephen Finlay. Metaethical Contextualism Defended. *Ethics*, 121: 7–36, 2010.
- Gunnar Björnsson and Tristram McPherson. Moral Attitudes for Non-Cognitivists: Solving the Specification Problem. *Mind*, 123(489):1–38, 2014.
- Simon Blackburn. *Spreading the Word*. Clarendon Press, Oxford, 1984.
- 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.
- Alejandro Pérez Carballo. Negation, Expressivism, and Intentionality. *The Philosophical Quarterly*, 70(279):246–267, 2020.
- 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.
- James Dreier. Relativism (and Expressivism) and the Problem of Disagreement. *Philosophical Perspectives*, 23(1):79–110, 2009.
- Andy Egan. Epistemic Modals, Relativism, and Assertion. *Philosophical Studies*, 133(1):1–22, 2007.
- Andy Egan. Disputing about Taste. In Feldman and Warfield, editors, *Disagreement*. Oxford University Press, Oxford, 2010.
- Andy Egan. Relativist Dispositional Theories of Value. *Southern Journal of Philosophy*, 50 (4):557–582, 2012.
- 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. An Expressivistic Theory of Normative Discourse. *Ethics*, 96(3):472–485, 1986.
- 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.
- Valentine Hacquard. Modality. In *Semantics: An International Handbook of Natural Language Meaning*, volume HSK 33.2, pages 1484–1515. Mouton de Gruyter, Berlin, 2011.
- 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.
- 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.
- Joshua Knobe and Seth Yalcin. Epistemic Modals and Context: Experimental Data. *Semantics and Pragmatics*, 7(10):1–21, 2014.
- 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. Conditionals. *Chicago Linguistics Society*, 22(2):1–15, 1986.
- 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. Attitudes *De Dicto* and *De Se*. *Philosophical Review*, 88:513–543, 1979a.
- David Lewis. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1):339–359, 1979b.
- David Lewis. Index, Context, and Content. In Kanger and Ohman, editors, *Philosophy and Grammar*, pages 79–100. Reidel, Amsterdam, 1980.
- John MacFarlane. *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford, 2014.
- 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.
- Caleb Perl. A User’s Guide to Hybrid Tools. *Mind*, 129(513):129–158, 2020.
- Jonathan Phillips and Matthew Mandelkern. Eavesdropping: What is it Good For? manuscript, 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. Ecumenical Expressivism: The Best of Both Worlds? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 2. Oxford University Press, 2007.
- 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. Attitudes and Epistemics. In *Expressing Our Attitudes*, volume 2. Oxford University Press, New York, 2015.
- 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.
- William Starr. Dynamic Expressivism About Deontic Modality. In Charlow and Chrisman, editors, *Deontic Modality*, pages 355–394. Oxford University Press, Oxford, 2016.
- 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.
- Nicholas Unwin. Quasi-Realism, Negation, and the Frege-Geach Problem. *The Philosophical Quarterly*, 49(196):337–352, 1999.
- Ralph Wedgwood. *The Nature of Normativity*. Oxford University Press, Oxford, 2007.
- Malte Willer. Advice for Noncognitivists. *Pacific Philosophical Quarterly*, 98(S1):174–207, 2017.
- Seth Yalcin. Epistemic Modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Nonfactualism about Epistemic Modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Seth Yalcin. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012.