

NONCOGNITIVISM WITHOUT EXPRESSIVISM

Bob Beddor

10 · 8 · 20

Outline

- 1 Two debates
- 2 A Toy Implementation
- 3 Descriptivist Benefits
- 4 Noncognitivist Benefits
- 5 Conclusion

Two distinctions

Descriptivism vs. Expressivism

- Descriptivists hold that moral discourse aims to describe the world.
- Expressivists hold that moral discourse does not aim to describe the world; rather, it serves to express the speaker's desire-like attitudes.

Cognitivism vs. Noncognitivism

- Cognitivists hold that moral belief has a mind-to-world direction of fit: it aims to represent the world.
- Noncognitivists hold that moral belief has world-to-mind direction of fit: to believe people ought to give to charity is just to desire that people give to charity (or approve of it, etc.).

What is the relation between these two debates?

Two debates

Standard Answer

Descriptivism \iff cognitivism; Expressivism \iff noncognitivism

My Aim: Pull these pairings apart; distinguish moral semantics from moral psychology.

The plan

Specifically, I want to explore the prospects of:

Descriptivist Noncognitivism

Descriptivist semantics + Noncognitivist psychology

Retains all of the explanatory advantages associated with expressivism

- Captures the connection between moral judgment & motivation
- Accounts for 'open question' phenomena while maintaining naturalistic respectability

At the same time, avoids the main challenge facing expressivism: the Frege-Geach Problem

Outline

- 1 Two debates
- 2 A Toy Implementation**
- 3 Descriptivist Benefits
- 4 Noncognitivist Benefits
- 5 Conclusion

First Ingredient: Descriptivist Semantics

For concreteness, I'll focus on moral uses of deontic modals:

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

Contextualist *Ought*

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

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

- Kratzer [1981, 1991, 2012]

For example:

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

says that all of the morally best of the accessible worlds are worlds where ppl give to charity.

First Ingredient: Descriptivist Semantics

Key idea behind descriptivism: moral discourse purports to represent the world in the same way that non-moral discourse does.

One way of cashing this out:

Descriptivism (First Pass)

A semantics is descriptivist if it assigns every utterance of a moral sentence representational truth conditions.

First Ingredient: Descriptivist Semantics

One way of making this more precise:

Descriptivism (Second Pass)

A semantics $\llbracket \cdot \rrbracket$ is descriptivist if, for every moral sentence φ and every sequence of contextually determined parameters c , $\llbracket \varphi \rrbracket^c$ is a set of worlds.

Contextualist *Ought* is descriptivist: in any context c , the semantic value of a moral use of a modal, relative to c , is a set of worlds.

One consequence of this is that the truth or falsity of any given moral assertion will be settled by the world alone.

Second Ingredient: Noncognitivist Belief Reports

It's usually thought that a descriptivist semantics leads to cognitivism.

I think this is too quick: it all depends on our semantics for belief reports, e.g.:

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

Standard semantics for belief reports:

Hintikka *Believes*

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

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

Observation: Contextualist *Ought* + Hintikka *Believes* leads to cognitivism.

Standard semantics for belief reports:

Hintikka *Believes*

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

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

According to Contextualist *Ought* + Hintikka *Believes*:

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

= true at w iff for every world v consistent with Ana's beliefs at w : all of the morally best of the v -accessible worlds are worlds where ppl give to charity.

This ascribes to Ana a representational mental state.

My proposal: In addition to quantifying over the believer's doxastic alternatives, *believes* also shifts the ordering source in the index to one provided by the believer's desire-like states.

Shifty Hintikka Semantics

Let a world-indexed desire function (δ_α^w) be a constant function from an arbitrary world u to a set of propositions representing α 's intrinsic desires at w .

Toy example: suppose 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:

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

Note that a world-indexed desire function has the same structure as an ordering source: both are functions from worlds to sets of propositions.

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

Noncognitivist *Believes*

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

(2) Ana believes that people morally ought to give to charity

= true at w iff 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 are worlds where ppl give to charity.

This gives us a genuinely noncognitivist view: moral beliefs depend on whatever desires you actually hold (not what you believe your desires to be).

Outline

- 1 Two debates
- 2 A Toy Implementation
- 3 Descriptivist Benefits**
- 4 Noncognitivist Benefits
- 5 Conclusion

Frege-Geach Problem

- **Embedding Problem:** Providing a principled and predictive semantics that explains how moral vocabulary embeds in logically complex constructions.
- **Semantic Relations Problem:** Providing a plausible account of semantic notions such as truth, validity, and consistency that apply to normative sentences.

Descriptivist noncognitivism avoids both of these problems, since it embraces a possible worlds semantics.

- Can use the standard possible worlds analyses of logically complex constructions.
- Can use the standard possible worlds definitions of truth, validity, and consistency.

Frege-Geach Problem

In short, the Frege-Geach Problem simply does not arise for descriptivist noncognitivism.

Worry: Don't some versions of the Frege-Geach Problem arise for any noncognitivist theory of moral belief reports?

Revenge of the Frege-Geach Problem?

Negation Problem

- Unwin 1999; Dreier 2006; Schroeder 2008; Pérez-Carballo 2020

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

The Negation Problem

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

Solving the Negation Problem

However, if we work through our semantics (Contextualist *Ought* + Noncognitivist *Believes*), we see that it straightforwardly captures these differences, yielding a new solution to the Negation Problem.

Some terminology

- α *decidedly desires* p at a world w iff for all of α 's doxastic alternatives v , all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds.
- α *decidedly does not desire* p at w iff for all α 's doxastic alternatives v , **not** all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds.

Solving the Negation Problem

Applying our semantics:

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

A Toy Model

- Cara and Dev are both unsure whether charity alleviates suffering; at two of their doxastic alternatives (a & b) it does; at two others (c & d) it does not.
- At a & c people give to charity; at b & d they do not.
- a & b are both accessible from each other, likewise with c & d.
- For Cara, a & b are equally desirable, as are c & d. For Dev, a is preferable to b, but c & d are equally desirable.




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

Revenge of the Frege-Geach Problem?

Taking stock: applying our semantics (Contextualist *Ought* + Noncognitivist *Believes*) yields a straightforward solution to the Negation Problem.

This should bolster confidence that the Frege-Geach Problem does not automatically undermine any noncognitivist theory of moral belief.

Has the Frege-Geach Problem Already Been Solved?

A different way of objecting to my argument is to point out that a number of expressivists have tried to solve the Frege-Geach Problem.

- E.g. Blackburn [1984; 1988]; Gibbard [1990; 2003]; Horgan & Timmons [2006]; Ridge [2006]; Schroeder [2008]; Silk [2014]; Charlow [2015]; Starr [2016]; Willer [2017].

If any of these solutions succeeds, then doesn't this undermine the argument for descriptivism?

Many (all?) of the proposed solutions to the Frege-Geach Problem face difficulties; at present there is no consensus on whether any succeeds.

Even if some sophisticated form of expressivism can handle the Frege-Geach Problem (a big if!), it's not clear whether there is any empirical motivation for the extra semantic machinery that the expressivist employs.

Many (all?) of the proposed solutions to the Frege-Geach Problem face difficulties; at present there is no consensus on whether any succeeds.

Even if some sophisticated form of expressivism can handle the Frege-Geach Problem, there is no discernible *empirical* motivation for the complicated semantic machinery that the expressivist employs.

By contrast, descriptivist noncognitivism has the virtue of semantic simplicity and conservatism: we employ a classical Kraterian semantics for modals, without any extra bells & whistles.

Outline

- 1 Two debates
- 2 A Toy Implementation
- 3 Descriptivist Benefits
- 4 Noncognitivist Benefits**
- 5 Conclusion

Why like expressivism?

- Naturalistic respectability
- The ‘Open Question’ phenomenon
- Motivational internalism
- Explaining normative disagreement

I’m going to argue that these are really advantages of noncognitivism, not expressivism *per se*.

Naturalistic Respectability

On the noncognitivist view, for Ana to believe people (morally) ought to give to charity is for her to decidedly desire charitable giving.

So normative belief is analyzed as a conative relation towards a descriptive proposition (*people give to charity*).

The Open Question Phenomenon

Moore [1903] infamously outlined a recipe for refuting any naturalistic paraphrase of (1) (*People (morally) ought to give to charity*), e.g.

(7) Charitable giving promotes wellbeing.

The Open Question Phenomenon

Moore invites us to compare the questions:

- (8a) Granted that charitable giving promotes wellbeing, but ought we give to charity?
- (8b) Granted that we ought to give to charity, but ought we give to charity?

According to Moore, (8a) is an 'open' question, whereas (8b) is not.

Moore took this to show that (8a) and (8b) must differ in meaning.

Common Complaints

- Assumes that meanings are always transparent to the speaker.
- Seems to prove to much, since it threatens to “bring the whole enterprise of conceptual analysis to a standstill” (Darwall et al. 1992)

Open Question Argument Reframed

We can avoid these complaints by reframing the Open Question Phenomenon in terms of *belief*, rather than meaning.

Doxastic Openness

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

Noncognitivism explains Doxastic Openness:
one can believe that ϕ ing is N without decidedly desiring ϕ ing, and hence without believing people ought to ϕ .

Open Question Argument Reframed

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.

Motivational Internalism

Basic idea: moral beliefs are intrinsically motivational.

One way of cashing this out:

Motivational Internalism

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

Motivational Internalism

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

Expressivists often claim to explain why Motivational Internalism holds.

- e.g., Stevenson [1944]; Blackburn [1998]; Gibbard [1990, 2003]

Motivational Internalism

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

But note that motivational internalism is a claim about moral *belief*, not moral language.

So what really explains motivational internalism is noncognitivism, not expressivism.

Explaining Disagreement

Final (putative) advantage of expressivism: explains how two people can agree on all the descriptive facts but still disagree over what people ought to do.

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.

Disagreement in Attitude

This idea traces back to Stevenson [1944: 3]:

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

Descriptivist noncognitivists are also free to appeal to 'disagreement in attitude' to explain normative disagreements.

OK, but what does it mean for two people to disagree in attitude?

Proves difficult to provide a substantive theory that avoids overgenerating or undergenerating disagreement (cf. Dreier [2006])

Disagreement in Attitude

However, I think that our solution to the Negation Problem provides just what is needed:

Disagreement in Attitude

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

Terminological Reminder

- α *decidedly desires* p at a world w iff for all of α 's doxastic alternatives v , all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds.
- α *decidedly does not desire* p at w iff for all α 's doxastic alternatives v , **not** all of the most desired of the v -accessible worlds (by the lights of δ_α^w) are worlds where p holds.

Disagreement in Attitude

Bonus: This account dispels a residual worry—namely, that anyone who appeals to some notion of ‘disagreement in attitude’ will be saddled with a disjunctive account of disagreement.

Our semantics for belief reports allows us to provide a general semantics for disagreement ascriptions:

Disagreement Ascriptions

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

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

or *vice versa*.

- Cf. Beddor [2019]

Disagreement in Attitude

Disagreement Ascriptions

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

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

or *vice versa*.

As a special case: Ana and Benny disagree over whether people ought to give to charity iff both:

- 1 Ana believes that people ought to give charity,
- 2 Benny believes that it's not the case that people ought to give to charity,

or *vice versa*.

Disagreement in Attitude

Equivalently: Ana and Benny disagree over whether people ought to give to charity iff both:

- 1 Ana decidedly desires that people give to charity,
- 2 Benny decidedly does not desire that people give to charity,

or *vice versa*.

Explaining Disagreement

The upshot: Contextualist *Ought* + Noncognitivist *Believes* yields a new theory of moral disagreement.

Once again, all of the work is being done by the noncognitivist theory of moral belief. Nothing in the story presupposed expressivism.

Outline

- 1 Two debates
- 2 A Toy Implementation
- 3 Descriptivist Benefits
- 4 Noncognitivist Benefits
- 5 Conclusion**

Wrapping Up

In this talk, I've tried to make the case for a new metaethical position:

Descriptivist Noncognitivism

Descriptivist semantics + Noncognitivist psychology

Virtues:

- Can be given a precise compositional semantics in line with the standard analysis of deontic modals
- Avoids the main problems for expressivism, while reaping all of the explanatory benefits

More general lesson: Need to be careful to distinguish moral semantics from moral psychology

THE END

Bonus Slides: Objections

Worry: this account does not underwrite the pragmatics of moral assertion.

Moral Pragmatics

Normally, a speaker's goal in asserting $\lceil \text{Ought } p \rceil$ is to get their audience to believe ought p .

Assertoric Update

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

Observation: Contextualist *Ought* & Noncognitivist *Believes* & Assertoric Update \nrightarrow Moral Pragmatics.

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.

- Cf. Mandelkern [2020]

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.

Given Assertoric Update, the goal of asserting $\lceil \text{Ought } p \rceil$ in a context c is to get your audience to believe $\llbracket \text{Ought } p \rrbracket^c$. Given Prospective Contextualism + Contextualist *Ought*, this means that your audience will believe that they desire p . Given Normal Access, it follows that normally, this will obtain only if they desire p .

Some of the work that has been traditionally assigned to assertoric content should be reassigned to semantic content.

Semantic Update

Normally, a speaker's goal in asserting a sentence φ , in a context with a sequence of contextually determined parameters c , is to get their audience to believe $[[\varphi]]$.

Observation: Contextualist *Ought* + Noncognitivist *Believes* + Semantic Update \Rightarrow Moral Pragmatics.