

finagling frege

In his recent paper, ‘Ecumenical Expressivism: Finessing Frege’, Michael Ridge claims to show how to solve the famous Frege-Geach problem for expressivism by allowing that moral sentences express both beliefs and desires, rather than simply desires. A number of other theorists have made similar or related suggestions recently, including David Copp [2001], Frank Jackson [1999], Richard Joyce [2006], Stephen Barker [2002], David Alm [2000], Dorit Bar-On and Matthew Chrisman [forthcoming], Daniel Boisvert [unpublished], and Jon Tresan [unpublished], and Allan Gibbard [2003], James Lenman [2003], and Stephen Finlay [2005] have all made suggestions that have broad affinities to these kinds of view. But Ridge claims that his view differs from other views which hold that moral sentences express both beliefs and desires, because it has a better claim to be a version of expressivism that can make use of cognitivist resources, rather than a version of cognitivism which claims to be able to make use of expressivist resources. He claims this makes his hybrid view the true heir to traditional expressivist views.

The central claim of Ridge’s paper is that his view allows us to ‘finesse’ the Frege-Geach problem, and the basic idea is simple. The Frege-Geach problem is the problem of how to explain why normative predicates behave in all of the same ways that descriptive predicates do. Ridge’s idea is that ordinary descriptive language works in this way because it expresses beliefs; so normative language can work in that way by expressing beliefs, too. He just thinks that doesn’t preclude normative language from also expressing desires.

In this brief paper I do two things. First I draw one small observation about Ridge’s hybrid solution to the Frege-Geach problem. And then I pose an independent and I think fatal objection to his account as it stands, particularly if we take seriously his claim to have explained anything about *logical* validity, or even about what looks like good reasoning. In closing, I’ll argue briefly that I think the feature of his account that my objection turns on is central to his claim to be defending a hybrid view that is still a form of expressivism – in his terms, to be defending a form of ecumenical expressivism, instead of a version of ecumenical cognitivism.

I what is the Frege-Geach problem?

The Frege-Geach problem for expressivism is to give an account of the meanings of complex sentences in terms of the meanings of their parts, which can explain why sentences with those complex structures have the kinds of systematic semantic properties that they do. For example, any sentence of the form ‘not P’ should have the property of being inconsistent with the sentence ‘P’. That inconsistency is due to the meaning of ‘not’. So an account of the meaning of ‘not’ needs to be able to explain it. But since expressivists think that if ‘P’ is a normative sentence, then the meanings of both ‘P’ and ‘not P’ must be given by the desire-like attitudes that they express, ordinary expressivists must appeal to these attitudes in order to explain why ‘P’ and ‘not P’ are inconsistent.

For example, suppose that ‘P’ expresses desire A, and that ‘not P’ expresses desire B. An adequate expressivist account of negation must explain why ‘P’ and ‘not P’ are inconsistent by appealing to properties of A and B. Despite the fact that this is arguably one of the easiest tasks facing expressivists, Nicholas Unwin [1999], [2001] has argued persuasively that no expressivist has yet successfully accomplished this.¹ For historical reasons, however, most discussions of the Frege-Geach problem have focused on conditionals. They have tried to explain why arguments like the following are valid, where P and Q are atomic normative sentences:

S1	P
S2	$P \supset Q$.
S3	Q

Here is why to think this problem is hard. If this argument is valid, then at a minimum, the following must be a logically inconsistent set: $\{P, P \supset Q, \sim Q\}$. But if Unwin is right that expressivists have trouble even satisfactorily explaining why $\{P, \sim P\}$ is a logically inconsistent set, their prospects for explaining why $\{P, P \supset Q, \sim Q\}$ is logically inconsistent look dim. So explaining why moral *modus ponens* arguments are valid is at least as hard as explaining why normative sentences are inconsistent with their negations, which no one has yet shown how to do.²

Yet there is more to explaining why the *modus ponens* argument is valid than explaining why $\{P, P \supset Q, \sim Q\}$ is a logically inconsistent set. If an argument is valid, then there is not only a problem with accepting its premises and denying its conclusion. There is also a problem with accepting its premises and not going on to accept its conclusion. This is one of the important features of valid arguments. Valid

¹ See also Dreier [2006], Schroeder [forthcoming a], and Schroeder [forthcoming b].

² I show how expressivists can do it in Schroeder [forthcoming a], [forthcoming b].

arguments are useful in philosophy not only because they create rational pressure for those who accept their premises to not deny their conclusions (unless they give up one of the premises), but because they create rational pressure for those who accept their premises to actually *accept* their conclusions (unless they give up one of the premises).³ So accounting for the properties of conditionals is clearly a harder problem for expressivists than accounting for the properties of negations. Ridge claims that his hybrid account can ‘finesse’ this problem, by providing an adequate account of why moral *modus ponens* arguments are logically valid.

2 the idea

Ridge’s idea is simple. Ordinary expressivists run into trouble, because their views have too few resources, and resources of the wrong kind. An ordinary expressivist account of the validity of our *modus ponens* argument would have the following materials to work with: to each sentence it would assign a desire, call them D1, D2, and D3, and it would assign another desire, D~3, to the negation of sentence S3. To explain why $\{P, P \supset Q, \sim Q\}$ is an inconsistent set, the ordinary expressivist would have to explain why there is something inconsistent (in the right sort of way⁴) about having desires D1, D2, and D~3. And to explain why the argument creates rational pressure to accept its conclusion, the ordinary expressivist would have to explain why desires D1 and D2 create rational pressure (of the right kind⁵) to have desire D3.

On the other hand, if you think that moral sentences express both beliefs and desires, then you would have more resources to solve this problem. To each sentence you would assign a belief as well as a desire. So S1 would get assigned to belief B1 and desire D1, S2 would get B2 and D2, S3 would get B3 and D3, and the negation of S3 would get B~3 and D~3. Now to explain why $\{P, P \supset Q, \sim Q\}$ is an inconsistent set, what you have to explain is why there is something inconsistent about being in all six of the following states: B1, D1, B2, D2, B~3, and D~3. But this is easy to explain, provided that we assign the right beliefs to our sentences. So long as B1 and B2 correspond to a valid argument for B3, they will be inconsistent in the same way as for an ordinary valid descriptive argument.

For example, suppose that B1 is the belief that someone would have by accepting ordinary descriptive sentence ‘P*’, that B2 is the belief that someone would have by accepting ordinary descriptive

³ Let me be clear that nothing rests, here, on confusing the important distinction between inference and implication, nor on any particular theory about in what this rational pressure consists. All that I am saying, is that there is something rationally problematic about someone who clearheadedly believes the premises of a logically valid argument, is specifically thinking about the premises and the conclusion at the same time, and yet remains completely undecided with respect to the conclusion. See section 3, below.

⁴ See van Roojen [1996].

⁵ See van Roojen [1996].

sentence 'P* \supset Q*', and that B3 is the belief that someone would have by accepting ordinary descriptive sentence 'Q*', and that B \sim 3 is the belief that someone would have by accepting ordinary descriptive sentence ' \sim Q*'. Then the inconsistency in {P,P \supset Q, \sim Q} is explained by the inconsistency in {P*,P* \supset Q*, \sim Q*}. But that is ordinary descriptive inconsistency, with which there is not supposed to be any problem. This is why the hybrid approach to the Frege-Geach problem can seem initially promising.

It leads to the following idea: that we can use the descriptive contents of moral sentences in order to explain why accepting the premises of the argument commits to accepting the conclusion, too. Since B1 and B2 commit to B3, the idea goes, someone who accepts the premises of our *modus ponens* argument is committed to accepting its conclusion.

3 the small observation

However, this is not right. For on the hybrid view, accepting the conclusion of our argument requires being in two states. It requires having belief B3, true. But it also requires having desire D3. So even if there is rational pressure to have belief B3 so long as you accept sentences S1 and S2, that does not translate into rational pressure to accept the argument's conclusion, unless we somehow have an argument that the agent is under rational pressure to have desire D3, or else have some reason to think that she must *already* have desire D3.

Ridge's article doesn't explicitly address this problem. Strictly speaking, he stipulates it away. His official solution to the problem is to offer the following definition of validity:

An argument is valid just in case any possible believer who accepts all of the premises but at one and the same time denies the conclusion would thereby be guaranteed to have inconsistent beliefs.⁶

So here is my small observation: to offer this definition is simply to *assume* that explaining the inconsistency of accepting all of {P,P \supset Q, \sim Q} is sufficient to explain the validity of the argument. But that is just to assume that there is no problem about explaining the rational pressure to accept the conclusion, not to explain where it comes from.

I think this small observation points at a major issue for hybrid views like Ridge's. By telling us that sentences can express two distinct mental states, they undermine excluded middle, at least in principle. If 'Q' and ' \sim Q' each express a desire as well as a belief, there can be no rational pressure to accept either

⁶ Ridge [2006, 326].

'Q' or ' \sim Q', even if there is rational pressure to accept either 'Q*' or ' \sim Q*', where 'Q*' is a descriptive sentence expressing the same belief as 'Q'. For one may simply fail to have the right *desire* to accept one of these sentences. For example, to focus on a paradigm for many hybrid views, if accepting 'Jon is a nigger' requires having a certain derogatory attitude and accepting 'Jon is not a nigger' requires having a certain derogatory attitude, then if you don't have the requisite derogatory attitude, you won't accept either one. But without excluded middle, we need more than the inconsistency of $\{P, P \supset Q, \sim Q\}$ to generate rational pressure for someone who accepts 'P' and ' $P \supset Q$ ' to accept 'Q'. It may turn out to be perfectly rational, for all this says, to accept both 'P' and ' $P \supset Q$ ', and simply neither accept nor deny 'Q'.

This is *bad*. Imagine having a conversation with someone who believes that stealing is wrong, believes that if stealing is wrong, then murder is wrong, and yet is still uncertain whether murder is wrong:

- You: Don't you believe that stealing is wrong?
 Her: Oh, of course stealing is wrong.
 You: And don't you believe that if stealing is wrong, then murder is wrong?
 Her: Oh, definitely. After all, murder is much worse than stealing!
 You: And you don't *deny* that murder is wrong?
 Her: No, I haven't made up my mind about it. I'm just not sure what to think.
 You: Do you at least have some *evidence* that murder is not wrong?
 Her: Oh, no – none at all. That's part of why I'm having trouble making up my mind.
 You: So shouldn't you conclude that it *is* wrong? After all, you believe that stealing is wrong, and that if stealing is wrong, then murder is wrong.
 Her: Hmmm... I just don't see how that commits me to thinking that murder is wrong. What do you think?

Your interlocutor in this situation is exhibiting a paradigm case of irrationality, and that is because the argument from 'stealing is wrong' and 'if stealing is wrong, then murder is wrong' to 'murder is wrong' is transparently valid. Any theory that succeeded of convicting your interlocutor of irrationality only if she goes on to form the belief that murder is *not* wrong, but allows that she may be totally and rationally coherent so long as she does not form a view either way would be deeply problematic.

Ridge doesn't address this problem at all, or tell us what feature of his view allows him to deal with it. But this is not to say that he doesn't have the resources for an answer, given other resources of his account – he does. (That is why I called this the 'small observation'.) For on his view, D1, D2, and D3 are all the very same desire, no matter what moral *modus ponens* argument we are considering. He believes that every moral sentence whatsoever expresses the very same desire. So Ridge thinks that anyone who accepts the premises of the argument *already has* the desire expressed by its conclusion. Consequently, he is not puzzled about where the commitment would come from to have desire D3.

This feature of Ridge's view is important and interesting, and it turns out to be possible to prove (given certain constraints set by the aspiration of hybrid theories to offer a solution to Frege-Geach style worries that appeals to machinery that could work even if classical, non-hybrid forms of expressivism could not) that the problem I raise in this section commits hybrid theorists to a dilemma. The dilemma is that they must either hold, along with Ridge, that all moral sentences express the same desire, or they must allow not only that moral sentences express both desires and beliefs, but that complex moral sentences can express indefinitely many different desires. I discuss the issues that this raises for hybrid theories, any why they face this dilemma, in another paper.⁷ For my purposes here, I simply want to make the small observation that Ridge is wrong: it is this distinct feature of Ridge's account that deals with the problem, not his account of validity, which merely assumes that the problem does not exist.

4 the problem for ridge

That, as I said, was simply my small observation. So on to the main problem. The problem for Ridge on which I want to focus, is that his account does not really assign the right beliefs to moral sentences in order for the argument among the descriptive contents of the sentences to itself be valid, anyway. I said above that if there are descriptive sentences 'P*' and 'Q*' such that B1 is the belief expressed by 'P*', B2 is the belief expressed by 'P* \supset Q*', and B3 is the belief expressed by 'Q*', then the hybrid theorist would at least be off to a good start. But Ridge's view does not have this form. On Ridge's view, B1, B2, and B3 have contents of the form, 'A', 'B \supset C', 'D'. So on his view the descriptive argument equivocates.

Ridge puts the essentials of his view this way:

Perhaps we should therefore understand moral predicates as expressing both a speaker's attitude in favor of actions in general insofar as they have a certain property (whatever property guides the speaker's approval of actions quite generally) *and* a belief which makes anaphoric reference to that property. The speaker may or may not have a very clear idea of what the relevant property is. The crucial idea here would be one of *anaphoric pronominal back-reference to the relevant property*. An utterance of the sentence, "There is moral reason not to eat meat" would on this account express a speaker's attitudes in favor of actions insofar as they have a certain property and the belief that refraining from eating meat has *that property*.⁸

The key feature of Ridge's view is not just that sentences express both beliefs and desires. It is that *which* belief a sentence expresses is a matter of which desire it expresses. The belief involves reference to

⁷ Schroeder [unpublished].

⁸ Ridge [2006, 313-314], italics in original.

something that is determined by which desire that very sentence expresses. This is part of why Ridge's view looks like expressivism with beliefs attached, rather than like cognitivism with desires attached. The belief expressed by a sentence depends on which desire it (that very sentence) expresses.

Ridge calls the way in which this reference is determined 'anaphora', but we can illustrate it with the following diagram:⁹

DI: Disapproval of actions with certain property	BI: that being friendly has ← <i>that</i> property
D2: Disapproval of actions with certain property	B2: that if being friendly has ← <i>that</i> property, then so does being friendly to strangers
D3: Disapproval of actions with certain property	B3: that being friendly to strangers has ← <i>that</i> property

Ridge thinks, of course, as I explained in section 3, that DI, D2, and D3 are all identical. So he thinks that the 'certain property' is in each case the same. So long as this is true, the argument from BI and B2 to B3 cannot fail to be truth-preserving. That is a good thing for Ridge. But it is not enough. For not all truth-preserving arguments are logically valid, and it is certainly not always irrational to accept the premises and deny the conclusion of a truth-preserving argument.

5 logical validity and anaphora

Only arguments which are truth-preserving in virtue of their form are logically valid. For example, the following argument is truth-preserving but not logically valid (within the relevant fiction, a qualification I'll henceforth ignore):

- PI Superman flies.
- P2 If Clark Kent flies, then I'm a walrus.
- C I'm a walrus.

⁹ Note that Ridge's view does not officially appeal to demonstratives, which I've just used in order to illustrate its flavor. Technically, since being an anaphor is a property of elements of syntactic structure, and since anaphors require antecedents which are also part of syntactic structure, but Ridge wants his to have an antecedent that is not part of the descriptive content of the sentence at all, it is far from clear how to extend the notion of anaphora in order to do what Ridge requires of it. The most natural way to make sense of his view is as the idea that normative predicates have hidden complexity and involve a hidden pronoun, whose content is the same as 'dthat(the property I hereby express approval of by uttering this sentence)'. See Schroeder [unpublished] for further discussion.

Lois Lane accepts the premises of this argument, but not the conclusion, and rationally so. In fact, she even denies the conclusion, and rationally so. The argument is truth-preserving but not logically valid. For to *understand* that it is truth-preserving, it is not enough to understand the meanings of the logical terms involved. We must also supply the appropriate interpretation of ‘Superman’ and of ‘Clark Kent’, on which they designate the same person. It is because Lois Lane does not realize that Clark Kent is Superman, that she is not in a position to recognize that the argument is truth-preserving, and that it is rational for her to accept the premises while denying the conclusion.

There are many interpretations of the sentences in this argument on which the premises are true but the conclusion is false – they are simply models in which ‘Superman’ and ‘Clark Kent’ pick out different things. But these are still appropriate interpretations as far as validity goes – the only constraints on the interpretation function are that it assign the same interpretation to each name and predicate, and those models meet this constraint. These constraints don’t come from nowhere; they are the constraints that someone can recognize in virtue of understanding the sentences.

There should be a similar constraint on an adequate account of logical validity for a language including sentential anaphora. Compare the following:

- PI* Superman – he flies.
- P2* But Clark Kent – if he flies, then I’m a walrus.
- C* I’m a walrus.

In this argument, we replace ‘Superman’ and ‘Clark Kent’ with pronouns that are anaphoric on different antecedents – ‘Superman’ for the first, and ‘Clark Kent’ for the second. The fact that this argument involves anaphora should not make a difference as to its validity. If the first argument was logically invalid, then this one should be, too. Its truth-preservingness given the *preferred* interpretation of ‘Superman’ and ‘Clark Kent’ is insufficient for logical validity; logical validity requires truth-preservingness in any model.

But the only constraint on interpretations of these sentences is that they should assign each anaphoric pronoun the same referent as its antecedent. After all, this is all that you need to understand about the anaphoric pronouns, in order to understand and be in a position to accept or deny these sentences. So ‘he’ in the first sentence should get assigned the same interpretation as ‘Superman’ and ‘he’ in the second sentence should get assigned the same interpretation as ‘Clark Kent’ in the second. This much ought to be guaranteed by an adequate logical understanding of anaphora. But the logic of anaphora should not turn out to guarantee that ‘Superman’ and ‘Clark Kent’ turn out to be co-referential as a matter of logic – that is a consequence of their *meaning*, not a product of logic. So like the first argument, this

argument should also turn out to be logically invalid, even though it *is* truth-preserving. It is an important consequence of this that Lois Lane can rationally understand and accept both premises of this argument without accepting its conclusion – indeed, while denying the conclusion. This is because she does not realize that Clark Kent is Superman. Anyone who does not realize this can rationally accept the premises of the argument and deny its conclusion.

The argument just considered involves *sentential* anaphora. Each pronoun is anaphoric on a different antecedent, located in its particular sentence. That makes it different from *discourse* anaphora, in which each pronoun is anaphoric on a shared antecedent. For example, compare the following argument:

- O You know that Superman guy?
- PI† He flies.
- P2† And if he flies, then I'm a walrus.
- C† So I'm a walrus.

If this argument involves discourse anaphora, then the pronouns in both PI† and P2† are anaphoric on the same antecedent. So someone who understands these sentences as they are intended will have to realize that they both refer to the same thing as their antecedent, 'Superman'. So an adequate logic of discourse anaphora should guarantee that any interpretation of these sentences needs to assign them both the same referent as their antecedent, in this case, 'Superman'. So an adequate logic for anaphora should make this argument logically valid, but the two previous arguments invalid, even though all three are truth-preserving. This is because the truth-preservingness of the first two arguments is guaranteed only by the preferred interpretation of the terms involved, and not by the logical terms. Consequently, Lois Lane cannot understand a case of discourse anaphora like this one¹⁰, and rationally accept the premises while denying the conclusion.

6 back to ridge

The problem with Ridge's account is that it appeals to something that looks like *sentential* anaphora. Ridge's assumption that all moral sentences express the same desire is interesting, and plays a crucial role in his view. But it is also substantive. It is not a view that is shared even by others who think that moral sentences express desires. For example, ordinary expressivists like Allan Gibbard ([1990], [2003]) think that so far from all moral sentences expressing the same desire, every moral sentence expresses a different

¹⁰ She can, of course, misunderstand the pronoun in the second sentence as deictically referring to something else – for example, to someone else in the room during the conversation – but that does not involve understanding the argument as a case of discourse anaphora.

desire. The identity of D1, D2, and D3 is therefore like the identity of Superman and Clark Kent. It may be true, and it is certainly an interesting hypothesis. Its truth is sufficient to guarantee that someone who has belief B1 and B2 but also belief B~3 has at least one false belief. But it is not sufficient to guarantee any *logical* inconsistency between the contents of B1, B2, and B~3, and it is not sufficient to guarantee any *rational* inconsistency in having beliefs B1 and B2 while also having belief B~3. So Ridge's view, I claim, has a problem with logical validity.

This is a very serious problem. Arguments that are truth-preserving but not in virtue of their form are arguments whose premises it is rational to accept while denying their conclusion, so long as one is not, as Lois Lane is not, privy to certain background information. Ridge's problem is severe, because the background information required, on his view, is his thesis that all moral sentences express the very same desire-like attitude. But even if Ridge is right that this thesis is *true*, it is not a thesis to which most of us are privy. As I've just pointed out, it is a thesis that classical expressivists explicitly deny, and it is explicitly denied by cognitivists like me, who deny that moral sentences express desire-like attitudes at all. But this means that if Gibbard or I accept the premises of a moral *modus ponens* argument and deny its conclusion, we are being exactly as irrational as Lois Lane is, for accepting P1 and P2 but denying C. That is to say, it is a consequence of Ridge's view that neither Gibbard nor I would be in the least irrational for doing such a thing. This is very, very *bad*. We saw in section 2 that the *easy* case for hybrid views is being able to explain the rational inconsistency of accepting the premises of a valid argument and denying its conclusion. But Ridge's view doesn't even get us that!

Moreover, it is important to emphasize just how general this problem is. Mark van Roojen [2005] has argued that views like Ridge's have a problem accounting for the validity of arguments considered over a stretch of time, because the attitude expressed can change over time. The problem that van Roojen raises is structurally similar to the one that I am raising. But mine is more general. What I've just argued is that Ridge's view commits a fallacy of equivocation even for arguments considered wholly at a single time.

How bad is this problem? Could Ridge solve it by changing the feature of his view according to which the belief expressed by a moral sentence makes anaphoric reference to the desire expressed by that very sentence? What if he took discourse anaphora as a model, and proposed that the belief expressed by any moral sentence makes reference to the single property that he approves of things for having?

I think there is a very simple reason why Ridge did not take this view, and it occupies a large part of his paper. Given that a number of other authors have recently proposed hybrid views, all of which can yield the same sort of solution to the Frege-Geach problem as his own, Ridge was understandably concerned in his paper to distinguish his own view from the others. The other hybrid views, Ridge claims,

are really cognitivist views in hybrid clothing – they are what he calls versions of ‘ecumenical cognitivism’. Whereas his own view, Ridge claims, is an expressivist view in hybrid clothing – it is ‘ecumenical expressivism’. Only his own view is the proper heir to the expressivisms of Blackburn, Gibbard, and Horgan and Timmons, along with the rest of the noncognitivist tradition.

Ridge’s grounds for this classification are that on his view alone does the belief expressed by a sentence depend on the desire expressed by that very sentence. If Ridge held instead that the belief expressed by a sentence is fixed by a single general fact about the attitudes of that speaker, then he would merely have an ordinary context-dependent semantics for the descriptive content of normative sentences, and he would no longer be able to distinguish his view from that of Stephen Barker [2002], who has *precisely* this view. So it is not an idiosyncratic feature of Ridge’s view which disables it from accounting for logical validity. It is the very central feature on the grounds of which he takes his view to be distinctive.

7 from inconsistency to logic

One way of trying to fix Ridge’s account might be to back off of his claim to have accounted for *logical* validity, and propose that his account should instead be understood as an account of *informal* validity.¹¹ The idea would then be to define a logically valid argument as one that is informally valid under any interpretation. But to take this route, Ridge would need to do a lot more work. He would need to give us an adequate general picture, for example, of what an interpretation of a predicate consists in. Because he doesn’t share cognitivists’ view that predicates have properties or property-determining senses for their semantic values, he can’t run with their view. And since he doesn’t share pure expressivists’ view that normative sentences express a single non-cognitive attitude, he can’t run with their conception of what an interpretation of a predicate consists in, either.¹²

Filling in a general account of what an interpretation of a predicate consists in is a highly nontrivial project for Ridge, after all. Keep in mind that an interpretation of a descriptive predicate will have to assign it some descriptive property, but that an interpretation of a normative predicate will have to assign it some attitude, and then only derivatively a descriptive property. What single characterization of what an interpretation does could be neutral between these two kinds of assignments? Ridge has done nothing to tell us.

¹¹ It is clear that as things stand, Ridge does take himself to have accounted for logical validity. In fact, in Ridge [2007, 65], he re-states his account of validity in this way (italics added): ‘An argument is *logically* valid just in case it is such that, necessarily, anyone who accepts the premisses and at one and the same time denies the conclusion is thereby guaranteed to have contradictory beliefs.’

¹² See Schroeder [forthcoming b] chapters 6-8.

I conclude that Ridge has far more work to do, before he should be claiming to have a serious proposal for the semantics of natural-language moral sentences, let alone to have solved the Frege-Geach Problem ‘on the cheap’.¹³

references

- Alm, David [2000]. ‘Moral Conditionals, Non-Cognitivism, and Meaning.’ *The Southern Journal of Philosophy* 38(3): 355-377.
- Barker, Stephen [2002]. ‘Is Value Content a Component of Conventional Implicature?’ *Analysis* 60(3): 268-279.
- Bar-On, Dorit, and Matthew Chrisman [forthcoming]. ‘Ethical Neo-Expressivism.’ Forthcoming in *Oxford Studies in Metaethics*, volume 4.
- Boisvert, Daniel [unpublished]. ‘Expressive-Assertivism and ‘The Embedding Problem.’’ Unpublished paper.
- Copp, David [2001]. ‘Realist-Expressivism: A Neglected Option for Moral Realism.’ *Social Philosophy and Policy* 18: 1-43.
- Dreier, James [2006]. ‘Negation for Expressivists: A Collection of Problems with a Suggestion for their Solution.’ *Oxford Studies in Metaethics*, volume I. Oxford: Oxford University Press, 217-233.
- Finlay, Stephen [2005]. ‘Value and Implicature.’ *Philosophers’ Imprint* 5(4), www.philosophersimprint.org/005004/.
- Gibbard, Allan [1990]. *Wise Choices, Apt Feelings*. Cambridge: Harvard University Press.
- _____ [2003]. *Thinking How to Live*. Cambridge: Harvard University Press.
- Jackson, Frank [1999]. ‘Non-Cognitivism, Validity, and Conditionals.’ In Dale Jamieson, ed., *Singer and His Critics*, 18-37.
- Joyce, Richard [2006]. *The Evolution of Morality*. Cambridge: MIT Press.
- Lenman, James [2003]. ‘Noncognitivism and Wishfulness.’ *Ethical Theory and Moral Practice* 6: 265-274.
- Ridge, Michael [2006]. ‘Ecumenical Expressivism: Finessing Frege.’ *Ethics* 116(2): 302-336.
- _____ [2007]. ‘Ecumenical Expressivism: The Best of Both Worlds?’ *Oxford Studies in Metaethics* 2: 51-76.

¹³ Special thanks to the participants in my graduate seminar at the University of Maryland in the spring of 2006.

Schroeder, Mark [forthcoming a]. 'How Expressivists Can and Should Solve their Problem with Negation.' Forthcoming in *Noûs*.

_____ [forthcoming b]. *Being For: Evaluating the Semantic Program of Expressivism*. Forthcoming from Oxford University Press.

_____ [unpublished]. 'Hybrid Expressivism: Virtues and Vices.' Unpublished paper.

van Roojen, Mark [1996]. 'Expressivism and Irrationality.' *The Philosophical Review* 105(3): 311-335.

_____ [2005]. 'Expressivism, Supervenience, and Logic.' *Ratio* (new series) 18(2):190-205.