



CHICAGO JOURNALS

Ecumenical Expressivism: Finessing Frege

Author(s): by Michael Ridge

Source: *Ethics*, Vol. 116, No. 2 (January 2006), pp. 302-336

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/10.1086/498462>

Accessed: 25/07/2013 16:40

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*.

<http://www.jstor.org>

Ecumenical Expressivism: Finessing Frege*

Michael Ridge

Metaethical expressivism has many virtues. It can explain the depth of moral disagreement, fits easily into a naturalistic world view, and can explain how moral judgment guides action. Moreover, so-called quasi-realism forms of expressivism can accommodate many of the realist-sounding things we say.¹ However, expressivism seems to have trouble making sense of utterances in which moral predicates occur in unasserted contexts. While we might be able to make sense of “torture is wrong” roughly along the lines of “Boo for torture!” it is hard to see how an account of this sort could deal with utterances like “If torture is wrong then Camp X-Ray should be abolished.” A speaker can accept the latter without disapproving of torture or Camp X-Ray. Moreover, any extension of expressivism to deal with such utterances must accommodate the validity of arguments in which they are premises. Since this problem comes to us through the work of P. T. Geach and is analogous to a problem Frege once posed for certain theories of negation, it is usually referred to as the “Frege-Geach problem.” In this article I articulate a new version of expressivism called “Ecumenical Expressivism,” which can avoid the Frege-Geach problem altogether. A crucial idea is that expressivism can and should embrace the thesis that moral utterances express both desires and beliefs. Before turning to the Frege-Geach problem, we must first do some philosophical spadework to uncover the logical space Ecumenical Expressivism occupies.

* Many thanks to Simon Blackburn, Matthew Chrisman, David Copp, David Enoch, Robin Flagg, Joshua Gert, John Hawthorne, Frank Jackson, Mark Lance, Jimmy Lenman, Huw Price, Robert Johnson, Geoffrey Sayre-McCord, Mark Schroeder, Michael Smith, Nicholas Sturgeon, Jon Tresan, Mark van Roojen, Pekka Väyrynen, Timothy Williamson, and audiences at Cambridge, Glasgow, Madison, and Oxford. Thanks also to two anonymous referees for *Ethics* and the editors of *Ethics*.

1. See Simon Blackburn, *Essays in Quasi-Realism* (Oxford: Oxford University Press, 1993).

I. A TALE OF A FALSE DICHOTOMY

Philosophical folklore tells a tale of a fundamental dichotomy in moral semantics, according to which moral utterances either (*a*) express beliefs but not desires or (*b*) express desires but not beliefs. This tale is a chapter of a larger narrative in the philosophy of language. This larger narrative depicts the meanings of words and sentences quite generally in terms of how they conventionally express states of mind. Perhaps the most famous rendition of this story comes from John Locke.² On any version of this story worth telling, we must distinguish the meaning of a word or sentence from what a particular speaker means by that word or sentence, a distinction famously drawn by Paul Grice.³ For example, suppose someone answers the question “Is Tony Blair an honest person?” by saying “Blair is a typical politician.” In this context, it will be clear that the speaker means that Blair is not very honest. Nonetheless, his sentence does not mean that Blair is not very honest. In Grice’s famous terms, the idea that Blair is not very honest is a conversational implicature of the utterance rather than part of its strict and literal meaning. For a defense of a sophisticated contemporary version of a broadly Lockean approach to semantics, see Wayne Davis’s *Meaning, Expression, and Thought*. Davis builds heavily on some of Grice’s most important insights but also departs from Grice’s own account in many important ways.⁴

The Lockean framework suggests that moral semantics should be understood in terms of the question “What states of mind do moral words and sentences conventionally express?” No plausible Lockean story will hold that all sentences express beliefs and only beliefs. Many nonassertoric utterances express nonbelief states. Utterances of sentences like “Hooray for Hollywood!” and “Boo for Bollywood!” express attitudes of approval and disapproval rather than beliefs. What sorts of states of mind are moral sentences conventionally used to express? A

2. In book 3 of *An Essay concerning Human Understanding*. See <http://www.ilt.columbia.edu/publications/Projects/digitexts/locke/understanding/chapter0301.html>.

3. See H. P. Grice, “Utterer’s Meaning, Sentence Meaning, and Word Meaning,” *Foundations of Language* 4 (1968): 225–42. The literature on this topic is enormous. For a useful bibliography, see Wayne Davis, *Meaning, Expression and Thought* (Cambridge: Cambridge University Press, 2003).

4. The Lockean framework is, of course, not uncontroversial, and on certain alternative approaches to semantics, meaning should be understood in terms of the expression of intensional entities (propositions, questions, imperatives) which are not themselves states of mind. One important challenge for the Lockean approach is to explain how a Lockean account can plausibly explain the distinction between semantics and pragmatics, a challenge I shall not explore here. For a contemporary Lockean take on this difficult challenge, see Wayne Davis, *Non-descriptive Meaning and Reference* (Oxford: Oxford University Press, 2005). Thanks to an editor at *Ethics* for helpfully pressing me on this point.

philosophically innocent answer is that they express moral judgments. This is fine as far as it goes, but we must now ask the question “What is a moral judgment?” On the one hand, we assess people’s moral judgments as true or false, we subject them to epistemic norms, and they can figure in rational inferences. These features suggest that moral judgments are beliefs. On the other hand, moral judgments can guide action without the help of an independently existing desire. Furthermore, intelligible moral disagreement can persist beyond agreement on the relevant facts. These considerations suggest that moral judgments are desires.⁵

These competing considerations have generated two diametrically opposed philosophical camps. Cognitivists insist that moral utterances express beliefs rather than desires, while expressivists hold that moral utterances express desires rather than beliefs. David Brink characterizes the expressivist as arguing that “moral judgments must express the appraiser’s non-cognitive attitudes, *rather than her beliefs*,”⁶ while “cognitivists interpret moral judgments as expressing cognitive attitudes, such as belief, *rather than non-cognitive attitudes*, such as desire.”⁷ This way of understanding the debate is standard. Frank Jackson and Philip Pettit characterize expressivism as maintaining that moral utterances express desires rather than beliefs.⁸ This story seems unlikely to have a happy ending. Two camps of theory builders emerge, with the members of each camp developing increasingly sophisticated theories designed to accommodate the features so easily accommodated by the opposition. Cognitivists try to show how beliefs can play the motivational role more naturally associated with desire.⁹ Expressivists argue that noncognitive attitudes of the right sort can play the intellectual roles more naturally associated with belief. Not surprisingly, these efforts can seem like at-

5. Yet others have been led by these considerations to reject the Humean philosophy of mind which insists that no mere belief can motivate. On these accounts, moral judgments are unitary states which are belief-like in that they are representational yet desire-like in that they can explain action without the presence of an independent desire. These states are sometimes referred to as “besires.” See J. E. J. Altham, “The Legacy of Emotivism,” in *Fact, Science and Morality: Essays on A. J. Ayer’s Language, Truth and Logic*, ed. Graham MacDonald and Crispin Wright (Oxford: Blackwell, 1986), 275–88. Besires are peculiar because intuitively any representation can exist without the motivation allegedly essential to that representation. An adequate appraisal of besire-based theories goes beyond the scope of this article. Here I simply take a broadly Humean philosophy of mind as given.

6. David Brink, “Moral Motivation,” *Ethics* 108 (1997): 4–32, 9; emphasis added.

7. Ibid., 5; emphasis added.

8. Frank Jackson and Philip Pettit, “A Problem for Expressivists,” *Analysis* 58 (1998): 239–51.

9. Some cognitivists also argue that this role is overstated; see, e.g., Sigrun Svavarsdottir, “Moral Cognitivism and Motivation,” *Philosophical Review* 108 (1999): 161–219.

tempts to fit a square peg in a round hole. Our story begins to look like a tragedy or a dark comedy.

Fortunately, our choices are not limited to (*a*) the thesis that moral utterances express beliefs but not desires and (*b*) the thesis that they instead express desires but not beliefs. That is indeed a false dichotomy. For we can instead hold that moral utterances express both beliefs and desires. An approach based on this idea would have certain obvious advantages. After all, if moral utterances express both beliefs and desires then it should be much easier to accommodate both the desire-like and belief-like aspects of moral judgment. Call any view according to which moral utterances express beliefs and desires "ecumenical" to mark the inclusiveness of such views. Would the truth of the Ecumenical View mean that the dispute over cognitivism and expressivism is a tempest in a teapot? This depends on whether we continue to characterize cognitivism and expressivism in exclusive terms. If we cling to these calcified characterizations then the Ecumenical View would entail that neither view is correct. While those weary of endless metaethical debates might welcome this conclusion, we should not embrace it too quickly. For these characterizations fit poorly with some of the main arguments for expressivism.

II. TRANSFORMING THE DEBATE

The Ecumenical View suggests that if there is real dispute between cognitivists and expressivists then it cannot simply be about the action-guiding features of moral judgment. Perhaps we should focus instead on the possibility of apparently intelligible moral disagreement in spite of agreement on all the relevant facts. G. E. Moore's famous (some would prefer "infamous") open question argument is relevant here.¹⁰ Take any proposed naturalistic analysis N of a moral predicate M. Moore's open question argument maintains that it will always be possible for someone without conceptual confusion to grant that something is N but still wonder whether it is really M. If, however, N really was an accurate analysis of M then the question "I know it is N but is it M?" would not be conceptually open. *Prima facie*, the open question argument seems to beg the question. For if the naturalist maintains that a given moral predicate M is equivalent in meaning to some naturalistic definition N then the fact that these questions are not conceptually open is a direct and obvious consequence of her view. However, the open question argument can be understood in a non-question-begging way (I do not claim that this is how Moore himself understood the argument). The crucial move is to understand the argument as an inference to the best explanation. On this interpretation, the main prem-

10. G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903).

ise of the argument is not that the relevant questions really are open but the more modest premise that they seem open to competent users of the terms. This argument does not beg the question insofar as the naturalist can consistently grant that the relevant questions do at least seem open to competent speakers. The issue is how best to explain this appearance. The argument maintains that the best explanation is that these questions really are open. That is, after all, a relatively simple and direct explanation of the sort that we would accept by default in other contexts. For example, we would accept such an answer when considering why the question "I know he is popular, but is he intelligent?" seems open to competent speakers. The fact that competent speakers persist in finding such questions open even when presented with a wide variety of proposed analyses bolsters this explanation's plausibility.

Generations of expressivists have argued that they are the real beneficiaries of Moore's argument. For Moore's nonnaturalism encounters several notoriously difficult problems shared by any antireductionist account of moral properties.¹¹ For example, antireductionist accounts have trouble explaining why moral properties supervene on nonmoral properties. For if moral properties really are irreducible then moral and non-moral properties are in Hume's terminology "distinct existences," which makes it puzzling why a difference in the former must always supervene on a difference in the latter. By contrast, expressivism offers a straightforward explanation of why our practices include a supervenience constraint. Whereas antireductionist realists must explain a metaphysical relationship between distinct properties, the expressivist needs only to explain the sensibility of adopting a supervenience constraint. Since the point of moral discourse is to recommend options on the basis of their natural properties, it is easy to see why such a constraint is sensible. Without such a constraint then it would be hard to see in what sense options were recommended on the basis of their natural properties.

More would need to be said about these lines of argument for a rational assessment of their plausibility. For example, something would need to be said about so-called Cornell-style moral realism which models moral semantics on the semantics for natural kinds developed by Saul Kripke and others.¹² However, the point here is simply to see how we

11. This distinction is problematic and has been glossed in a wide variety of ways; see Michael Ridge, "Moral Non-naturalism," *Stanford Encyclopedia of Philosophy*, available at <http://plato.stanford.edu/entries/moral-non-naturalism/>.

12. For a classic presentation of the Cornell-style view, see Richard Boyd, "How to Be a Moral Realist," in *Essays in Moral Realism*, ed. Geoff Sayre-McCord (Ithaca, NY: Cornell University Press, 1998), 181–228. For some very useful discussions of how Moore's open question argument might be extended to deal with these interesting forms of moral realism, see Terence Horgan and Mark Timmons, "Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived," *Philosophical Papers* 21 (1992): 153–75, and

should define cognitivism and expressivism if these arguments provide the main motivations for expressivism. This suggests that expressivists must insist that moral utterances do not express beliefs which are such that the utterance is true just in case the belief is true. For the open question argument is meant to show that no representation of the facts can commit someone to a moral stance. So perhaps we should understand cognitivism and expressivism as follows:

Cognitivism: For any moral sentence M, M is conventionally used to express a belief such that M is true if and only if the belief is true.

Expressivism: For any moral sentence M, M is not conventionally used to express a belief such that M is true if and only if the belief is true.

This distinction is exclusive but not exhaustive. For there is logical space for a hybrid view according to which some but not all moral utterances express beliefs which provide their truth-conditions.¹³ For present purposes I put such hybrid views to one side, except to note that classifying them as hybrid views is plausible.¹⁴ More important, expressivism as defined above is a negative thesis, but expressivists historically also embrace the positive thesis that moral utterances express attitudes of approval and disapproval. Since part of my point is that this positive thesis can be understood as common ground between expressivists and cognitivists, I have not included it as a defining feature of expressivism.

On this taxonomy there will be cognitivist and expressivist versions of the Ecumenical View. Ecumenical Cognitivism allows that moral utterances express both beliefs and desires and insists that the utterances are true if and only if one of the beliefs expressed is true. Ecumenical Expressivism also allows that moral utterances express both beliefs and desires but denies that a moral utterance is guaranteed to be true just

"Troubles on Moral Twin Earth: Moral Queerness Revived," *Synthese* 124 (1992): 224–60. See also Nicholas Sturgeon, "Moore on Ethical Naturalism," *Ethics* 113 (2003): 528–56.

13. Paul Edwards and David Wiggins have defended such views. See Paul Edwards, *The Logic of Moral Discourse* (Glencoe, IL: Free Press, 1955). Thanks to an anonymous referee for bringing this to my attention.

14. One might argue that such hybrid views should be classified as forms of cognitivism since they vindicate the idea that there are objective moral facts. However, hybrid views (like cognitivism itself) leave open the possibility that all nontrivial moral claims are false. John Mackie's error theory is compatible with cognitivism and indeed with the hybrid view discussed in the text. Moreover, on a quasi-realist account, expressivism is compatible with the idea that there are "objective moral facts" but gives that phrase a deflationary reading. The debate between expressivists and cognitivists is not therefore well construed in terms of the question of whether there are objective moral facts. The debate concerns the question of whether moral utterances express moral beliefs which provide the truth-conditions for those utterances, and on that point the hybrid view is neither fully cognitivist nor fully expressivist. Thanks to an anonymous referee for pressing me on this point.

in case the belief(s) it expresses is (are) true. The logical space opened up by this taxonomy has been largely but not entirely ignored by cognitivists, though here the tide seems to be turning. The most extensive published discussion of the idea behind Ecumenical Cognitivism is perhaps in David Copp's "Realist-Expressivism: A Neglected Option for Moral Realism."¹⁵ On one reading of James Dreier's speaker-relativist theory, that theory is best understood as a species of Ecumenical Cognitivism.¹⁶ Daniel Boisvert develops a form of what I would call Ecumenical Cognitivism in his "Expressive Assertivism and the Embedding Objection."¹⁷ A form of Ecumenical Cognitivism is also developed by Jon Tresan in his "De Dicto Internalist Cognitivism," and Matthew Millar argues for a form of Ecumenical Cognitivism.¹⁸ Jackson mentions the idea in passing.¹⁹

By contrast, the possibility of Ecumenical Expressivism has, as far as I know, been entirely neglected. The contrast between Ecumenical Expressivism and Ecumenical Cognitivism can seem subtle, but the differences are precisely the ones that are relevant from the point of view of Moorean open question intuitions and worries about explaining supervenience. For the Ecumenical Cognitivist, belief has a kind of priority, in that which beliefs are candidates for counting as moral is fixed by their content.²⁰ Either this privileged content is about some irreducible

15. David Copp, "Realist-Expressivism: A Neglected Option for Moral Realism," *Social Philosophy and Policy* 18 (2001): 1–43.

16. James Dreier, "Internalism and Speaker Relativism," *Ethics* 101 (1990): 6–26. Thanks to Joshua Gert for useful discussion of the relevance of Dreier's work in this context.

17. See Daniel Boisvert, "Expressive Assertivism and the Embedding Objection," http://www.csub.edu/~dboisvert/Online_Papers/EA_and_The_EMBEDDING_OBJECTION-Long_Version.pdf.

18. Jon Tresan, "De Dicto Internalist Cognitivism" (unpublished manuscript, University of Florida); Matthew Millar, "Making Moral Judgements: Internalism and Moral Motivation" (PhD diss., University of Edinburgh, 2005).

19. Frank Jackson, *From Metaphysics to Ethics* (Oxford: Oxford University Press, 1998), 161.

20. Actually, some cognitivists may instead claim that moral judgments are irreducibly indexical. A crude version of subjectivism would hold that "I ought to X" just means "Xing would promote the satisfaction of my desires." On this sort of view, no particular content is essential to moral belief as such. My moral beliefs will be about me, while your moral beliefs will be about you. In David Kaplan's terminology, this species of cognitivism insists that moral beliefs as such must have a certain character rather than a certain content, where character is a function from a context of utterance to a content (see David Kaplan, "Demonstratives," in *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein [Oxford: Oxford University Press, 1989], 481–563). Your belief that you are tall and my belief that I am tall are identical in character but differ in content. Whether the cognitivist holds that moral beliefs as such must have a particular content or a particular character, the basic idea remains that we have a constraint on which beliefs can count as moral for a given moral agent which is independent of the agent's motivations. The content

moral property or not. If it is then worries about explaining supervenience arise; if not then Moorean worries are germane. Ecumenical Expressivism instead gives logical priority to desire. The belief that a given moralizer must have to count as judging that X is morally right is a function of that speaker's potentially idiosyncratic proattitudes rather than vice versa. So someone who approves of actions insofar as they maximize utility will count as thinking charity is morally right only if she believes that charity maximizes utility.²¹ However, this is the relevant content only because of the person's attitudes. Since making a moral judgment does not on this account require a belief with any particular content, Ecumenical Expressivism avoids the open question argument. Moreover, Ecumenical Expressivism can embrace the standard expressivist explanation of supervenience canvassed above.

Ecumenical Expressivism is a sort of dialectical mirror image of Ecumenical Cognitivism. Just as Ecumenical Cognitivism can steal the thunder of traditional expressivists by accommodating the motivating power of moral judgment "on the cheap," Ecumenical Expressivism can steal the thunder of traditional cognitivists by accommodating the logical features of moral judgment on the cheap. The real issue between cognitivists and expressivists may therefore have little or nothing to do with judgment internalism or the Frege-Geach problem. Instead, the dispute may hinge on more old-fashioned (and unfashionable) Moorean open questions and supervenience.

In a way, the suggestion that expressivists can make room for the expression of beliefs should not be surprising. In light of so-called thick moral predicates like "cowardly" it has long been clear that expressivists cannot plausibly insist that moral utterances never express beliefs. To say someone is cowardly is at least in part to express the belief that the person is easily motivated by fear, and expressivists have admitted this much. Nor has the expressivist recognition of a role for beliefs been limited to utterances employing "thick" predicates. Even A. J. Ayer's notoriously crude emotivism allowed that an utterance like "You acted wrongly in stealing that money" expresses (in some sense of 'express' anyway) the belief that the speaker's interlocutor stole the money.²² R. M. Hare distinguishes the descriptive meaning of a moral utterance

of the independently fixed candidate belief then determines which motivations are necessary for a given belief to count as moral. Ecumenical Expressivism rejects all such prior constraints on which beliefs can count as moral.

21. Or rather, someone who approves "in the right way"; expressivists must explain what is distinctive of moral proattitudes as such, but this difficult issue must also be put to one side here. For some discussion, see Gibbard's distinction between accepting a norm and being in the grip of one (Allan Gibbard, *Wise Choices, Apt Feelings* [Oxford: Oxford University Press, 1990], 60–61).

22. A. J. Ayer, *Language, Truth, and Logic* (New York: Dover, 1952), 107.

from its evaluative meaning and concedes that moral utterances typically have both kinds of meanings.²³

However, what has not been recognized is that expressivists can provide a thorough and systematic account of the role of belief in moral judgment for both thin and thick moral predicates, and that doing so provides crucial resources with which to disarm the Frege-Geach problem. The role of belief in Ayer's account is too limited to achieve this goal. Atomic moral utterances will on his view often simply be the expressions of a suitable noncognitive attitude. Saying "stealing is wrong," for example, is on Ayer's view not to express a belief at all but simply to express a suitable attitude against stealing. Utterances like "Your stealing that money was wrong" express a belief (if they do at all; even this is controversial) only because of the particular presupposition which happens to be involved in that utterance. Many atomic moral utterances will involve no such presuppositions. Hare, by contrast, seems to suppose that a particular speaker might be known by her interlocutors to accept a certain moral outlook according to which something is wrong just in case it has certain descriptive features. These features might be widely thought throughout the community to be necessary and sufficient for wrongness, and the members of the community may know that these standards are shared. In such a conversational context, moral utterances will conversationally imply (in roughly Grice's sense) that the relevant descriptive features are present.²⁴ In communities in which these associations are widely shared, the descriptive meaning might even come to be a sort of conventional implicature. This sort of descriptive meaning as implicature is interesting but too contingent to figure in an adequate solution to the Frege-Geach problem. For in contexts in which one's interlocutors do not know one's moral outlook or in which the relevant implicature is canceled, this descriptive meaning will be absent.²⁵

23. R. M. Hare, *The Language of Morals* (New York: Oxford University Press, 1952), 146.

24. Hare himself does not distinguish conversational implicature from other sorts of meaning, but this seems like the best interpretation of what he has in mind.

25. It might be objected that the connection between attitude and belief expressed in a single moral judgment is contingent according to Ecumenical Expressivism too, so that if contingency is a problem for Hare's account then it is a problem for my own account as well. The reply is that on my own account it is not contingent that a moral utterance always expresses a belief in which anaphoric reference is made to the property in virtue of which the speaker approves of actions. Moreover, it is this belief that does the work in avoiding the Frege-Geach problem. There is no analogously noncontingent belief involved in moral judgment on Hare's account. On Hare's account there will be cases in which no belief at all is expressed by a moral utterance, in which case the machinery of Ecumenical Expressivism will be unavailable. It is in this sense that the connections Hare draws between moral judgment and belief are too contingent to do the needed work. Thanks to an anonymous referee for discussion.

III. THE FREGE-GEACH PROBLEM

It is a platitude that moral predicates can be embedded in unasserted contexts and valid arguments. The following simple argument is often used to illustrate the point:

- (1) Lying is wrong.
- (2) If lying is wrong then getting one's little brother to lie is wrong.
Therefore, (3) getting one's little brother to lie is wrong.

If 'wrong' in premise 1 expresses disapproval of lying, while 'wrong' in premise 2 does not express disapproval then 'wrong' is used with different meanings in premises 1 and 2. In that case, there is a fallacy of equivocation. However, the idea that arguments of this form must involve such equivocation is absurd. Giving an account of the meanings of moral predicates in unasserted contexts which can explain the validity of such arguments is the heart of the Frege-Geach problem.

Some argue that the Frege-Geach problem is completely disarmed if so-called deflationist accounts of truth and truth-aptness are correct and can be accommodated by quasi-realist forms of expressivism. Whereas early expressivists held that moral utterances are not truth-apt, contemporary expressivists often claim that they are truth-apt but embrace deflationist theories of truth and truth-aptness. Roughly, on a deflationist theory of truth there is nothing more to truth than disquotation. Deflationist accounts of more interesting uses of 'true' such as "Everything she just said is true" are more tricky, but we need not get caught up in those issues here. A deflationist theory of truth-aptness claims, roughly, that there is nothing more to a sentence's being apt for truth than its being such that its substitution into the schema "it is true that p" is well formed.²⁶

The possibility of combining deflationism about truth and truth-aptness with expressivism opens up some interesting logical space. So

26. Perhaps the deflationist about truth-aptness should also require that the sentence be meaningful and not too vague to be truth-apt. This depends on whether being truth-apt is simply being the sort of thing that it makes sense to call true or false or whether being truth-apt also requires being determinately true or false. Deflationists may well have in mind the former. Lewis Carroll's "Twas brillig, and the slithy toves / Did gyre and gimble in the wabe" is perhaps in this deflationist sense truth-apt yet has no determinate truth-value. Someone who heard a recitation of "Jabberwocky" but did not realize that it was a nonsense poem might reasonably wonder whether it was really true that the slithy toves gyred and gimbled in the wabe. Yet if someone heard a poem in which the question "Do the slithy toves gyre and gimble in the wabe?" was asked and this person then wondered whether that question was true then this would betray a kind of confusion—in more old-fashioned terms, a category mistake. Perhaps we should distinguish different conceptions of truth-aptness here, with some more deflationary than others. Thanks to an anonymous referee for pressing this point.

long as the belief expressed by a moral utterance is not semantically guaranteed to provide the truth-conditions for the utterance, the fact that the belief expressed contingently provides the truth-conditions for the token utterance is consistent with expressivism as characterized here. For example, perhaps Jackson and Pettit are right to hold that expressivists must admit that moral utterances express the belief that the speaker has the attitude expressed.²⁷ However, it does not follow from this that the utterance is true just in case the belief is true. For example, perhaps my saying that “X is good” is true just in case Pettit approves of X” just is my saying that X is good if and only if Pettit approves of X and my saying the latter, in turn, may just be the expression of an attitude of approval toward all and only the Pettit-approved stuff. On this view, someone can allow that Pettit’s utterance does express the belief that Pettit approves of X but deny that this provides the truth-conditions for his utterance. For to deny that Pettit’s belief provides the relevant truth-conditions is simply a further first-order remark, and hence (in part) an expression of a further noncognitive attitude. In particular, it is to express attitudes of approval not cued to all and only the Pettit-approved.

Some philosophers, most notably Daniel Stoljar and Huw Price, argue that if both deflationism about truth and truth-aptness are correct then the expressivist can adopt a truth-conditional account of the validity of arguments in which moral predicates figure, in which case the project undertaken here seems unnecessary.²⁸ However, I am not here assuming that deflationist theories of truth and truth-aptness are defensible. The account developed here is neutral on those difficult questions for dialectical reasons. It would be unfortunate if the tenability of expressivism depended on the tenability of deflationism. Of course, anyone willing to “bite the bullet” and hold that moral utterances are not truth-apt might also bite the bullet and deny that they can figure in valid inferences. I shall here try to minimize the number of bullets in the expressivist’s diet.

Moreover, expressivism faces an important Frege-Geach-style challenge which remains even if we accept the Stoljar-Price truth-conditional approach. For given a generally Lockean account of semantics we shall need some account of what states of mind are expressed by utterances with moral predicates in unasserted contexts and an account of what constitutes the acceptance of such utterances. Furthermore, this account

27. See Jackson and Pettit, “A Problem for Expressivists.”

28. See Daniel Stoljar, “Emotivism and Truth Conditions,” *Philosophical Studies* 70 (1993): 81–101; and Huw Price, “Semantic Deflationism and the Frege Point,” in *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*, ed. S. L. Tsohatzidis (London: Routledge, 1994), 132–55.

must meet certain constraints. First, it must explain why anyone who accepts the premises but denies the conclusion of a valid argument is making a logical mistake. This inconsistency must be distinguished from the pragmatic inconsistency found in the kinds of sentences famously discussed by Moore (e.g., “I believe that p but not-p.”) since those so-called paradoxes do not involve outright contradiction. Call this the “inconsistency constraint.”²⁹ Finally, expressivism must also explain how it can sometimes be rational to infer the conclusion of a valid argument using moral predicates. Call this the “rational inference constraint.” The Stoljar-Price solution to the Frege-Geach problem does not show how this constraint can be satisfied. Existing versions of expressivism also fall prey to what Cian Dorr has referred to as the “wishful thinking” problem (discussed below), wrongly classifying rational inferences as mere wishful thinking. So deflationism notwithstanding, we still need an expressivist account that meets these two constraints to meet the challenge posed by the Frege-Geach problem.³⁰

IV. ECUMENICAL EXPRESSIVISM

Consider the sentence “There is moral reason not to eat meat.”³¹ On the Ecumenical Expressivist account, an utterance of this sentence will express both an attitude in favor of certain kinds of actions and omissions and a belief that refraining from eating meat is an omission of the relevant kind. Perhaps I am a utilitarian and approve of actions and omissions just insofar as they promote happiness. However, on an expressivist account maximizing utility cannot be part of the meaning of my sentence. 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

29. Here my discussion dovetails with Gibbard’s, which itself picks up on some nice points developed by James Dreier. See Allan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2003), 61–68; and James Dreier, “Expressivist Embeddings and Deflationist Truth,” *Philosophical Studies* 83 (1996): 29–51.

30. See Walter Sinnott-Armstrong, “Expressivism and Embedding,” *Philosophy and Phenomenological Research* 61 (2000): 677–93, for a more extensive argument that deflationism can solve at most some but not all of the problems in this area.

31. In the text, I suppose for expository reasons that the most basic moral predicate is “is a moral reason,” with other moral predicates like “is morally good” and “is a moral duty” being analyzable in terms of moral reasons. The moves made in the text can just as easily be made in metaethical theories which reject this assumption.

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.³² Returning to our analogy with uses of 'hooray', consider the following sentence (call it "MP"): "Hooray for teams that have a certain property such that Manchester United has that property just in case the Packers have it." Again, this is a very odd sentence; cheers tend to be much less verbose and more catchy. Nonetheless it is good English and provides a useful albeit imperfect analogy. An utterance of MP expresses an attitude of approval keyed to a certain property and expresses a belief involving anaphoric reference to that property—the belief that Manchester has the property just in case the Packers have it. Moreover, the utterance manages to do all of this without inheriting the truth-conditions of the belief it expresses.

One advantage of this approach is that it accommodates the plausible idea that a competent user of a predicate must be committed to applying the predicate to all and only things with certain features. In other words, competence with a predicate presupposes a rough-and-ready conception of necessary and sufficient conditions for the predicate's application. The speaker's appreciation of these conditions may be vague and indeterminate. This requirement for competence with a predicate can sound very demanding, but it is actually rather lax. To emphasize how relaxed this constraint is, note that one could be competent with a predicate on this account simply by supposing that a predicate applies to something just in case it is sufficiently similar to a particular paradigm, where "sufficient similarity" is left completely vague

32. Arguably, the proposed account relies on a broader notion of anaphora than the standard one invoked in the philosophy of language. For often the idea in those contexts seems to be that anaphora must refer back to something to which the speaker has referred previously. It is not clear on the account on offer, though, whether any such reference has occurred. For it is not clear that expressions of attitude in favor of actions with a given property literally refer to those properties. However, it seems clear enough that we can use anaphoric devices to refer back to properties introduced into conversation through expressions of attitude. For example, I can say "Hooray for the property of being red. That property is my favorite property." The second sentence makes anaphoric back reference to the property of being red, and any speaker of English would understand that this is what "That property" picks out. If the first sentence does not literally refer to the property of being red (and I do not mean to concede that it does not but just want to explore the implications of this view) then what follows is that we must broaden our conception of anaphora to fit better with ordinary language and not that ordinary language is mistaken because anaphora is limited to context in which reference has already taken place. Perhaps it is enough for anaphora that a property has been introduced into the conversation in some much more minimal sense that need not involve reference. Developing a full account of how this notion of "being introduced into the conversation" should be understood is an important task but not one I shall undertake here. Thanks to an anonymous referee for pressing me on this point.

and open-ended. Cognitive science models of how people make moral judgments which emphasize the idea of prototypes are compatible with Ecumenical Expressivism.

Once again, the idea is not that a competent speaker necessarily could articulate the property in virtue of which she approves of actions. This is why the device of anaphora is essential. The suggestion that someone has a certain *je ne sais quoi* provides a helpful analogy, since that suggestion implies both a proattitude toward someone in virtue of some of her properties and an explicit recognition that one is not sure just what those properties are. Of course, the analogy is not perfect; a moral speaker may well know exactly what properties she has in mind. The point here is simply that one need not know this to be a sincere and competent user of moral predicates. Another analogy can be found in an utterance of a sentence like "There's something about Mary." Sentences of the form "There's something about X" have become idiomatic and express an attitude of approval toward X in virtue of certain of X's properties even though the speaker may be unsure exactly what those properties are. Note that while we might treat an utterance of "There is something about Mary" as truth-apt, it would be a mistake for a speaker to endorse it as true simply because she agrees that Mary does in fact have some properties (doesn't everything?!). Rather, anyone who endorsed such an idiomatic utterance as true would thereby express suitable proattitudes with respect to Mary in virtue of some of her properties.

A similar case, but with the opposite conative valence, is the idiomatic remark that "Something is rotten in the state of Denmark," which is almost never taken to have what a naive compositional analysis would suggest to be its literal meaning. Normally, such an utterance serves to express the speaker's sense of unease with some feature of a contextually relevant situation even though the speaker may be unsure just what the relevant feature is. Again, I do not mean to suggest that these examples provide exact models but, rather, only rough and suggestive analogies. Also, a speaker's use of a moral predicate to express the belief that an object has a certain property should not be understood as the assertion that it has it. To assert that p is only one way of expressing the belief that p; one can express the belief that p through Gricean implicature, for example.

So perhaps we should understand atomic uses of moral predicates as follows:

"There is moral reason to X" expresses (a) an attitude of approval of a certain kind toward actions insofar as they have a certain property and (b) a belief that X has that property.

"Morally, there is most reason to X" ("Morally, one ought to X")

expresses (a) an attitude of approval of a certain kind toward actions insofar as they have a certain property and (b) a belief that X has that property to a greater extent than any of the available alternatives.

"The fact that X would be P is a moral reason to X" expresses (a) an attitude of approval of a certain kind toward actions insofar as they have a certain property and (b) a belief that being P constitutes that property with respect to X.

A number of points need to be made to clarify this proposal.

First, the beliefs expressed by moral utterances must be understood as not thereby guaranteed to provide their truth-conditions. Otherwise the resulting position would not be an expressivist one. An expressivist might explain this failure to provide truth-conditions in one of two ways. First, they might claim that moral utterances simply are not truth-apt. This would be more in line with earlier forms of expressivism associated with Ayer and Bertrand Russell. Alternatively, the expressivist might hold that moral utterances are truth-apt but deny that their truth-conditions necessarily are provided by the beliefs they express. The main idea here would be to defend some form of deflationism about truth and truth-aptness.³³ Given deflationism about truth and truth-aptness, the expressivist might argue that whether an agent's belief provides the truth-conditions for her utterance will be a substantive first-order question and not a question to be settled by metaethical theorizing. In particular, to say that a given belief provides the truth-conditions for a given moral utterance is to express one's own attitudes as approving of things just insofar as the belief in question is true of them. So I might know that you are a hedonist and that your moral utterances express your approval of pleasure and your belief that various states of affairs involve pleasure. If you say that there is reason to do something then I can without confusion deny that what you have said is true even if I grant that the action would promote pleasure. For my claim that what you have said is not true really just amounts to my expression of my own attitudes as not favoring actions insofar as they promote pleasure. I shall not here argue for one or the other of these two approaches. In what follows, the reader should focus on whichever of the two approaches just canvassed seems more plausible.

Second, the version of Ecumenical Expressivism developed here makes heavy use of the idea of anaphoric beliefs. One interesting objection explores how it might deal with anaphoric utterances.³⁴ Suppose,

33. Here I briefly abandon my stated neutrality about deflationism about truth to explore this version of expressivism; I am not committed to this version, though the argument of my article as a whole remains neutral on whether deflationism is correct.

34. Thanks to an anonymous referee for pressing me on this point.

for example, that you say "Sacrificing the happiness of a few for the greater happiness of the many is right." I, no utilitarian, reply by saying "Sacrificing the happiness of a few for the greater happiness of the many sometimes has that property, but sometimes it does not." On a natural understanding of anaphora, my use of 'that property' refers back to the property to which you referred. By hypothesis, I am no utilitarian, so my utterance should not suggest that I share your tendency to evaluate actions simply in terms of whether they maximize happiness. What should expressivists say about such utterances?

Contemporary expressivists maintain that in a suitably deflationist sense of 'refers', moral predicates refer to moral properties. On Allan Gibbard's view, for example, normative utterances express a speaker's plans. More generally, judgments in which normative predicates figure are in his terms 'plan-laden', which is to say that such judgments essentially involve deciding how to act and in this sense are not purely descriptive. Crucially, claims about the reference of moral predicates are plan-laden on Gibbard's view. As Gibbard puts the point, "Any planner is committed to a Claim of *Factual Constitution*: that there is a factual property that *constitutes* being okay to do. This too will be a plan-laden truth, not a truth of prosaic fact."³⁵ Gibbard then invokes the idea of character, which is a technical idea from the philosophy of language. In Gibbard's sense, "the character of a concept gives its extension—its truth value, what it designates, or the like—as the concept is applied from various standpoints."³⁶ Gibbard explains that "with a plan-laden concept like being okay to do, what its character is will be a planning question, a question of how to live. . . . This contrasts with the case of a naturalistic concept, whose character is settled by the concept alone."³⁷ In which case, the claim that 'right' refers to the property of maximizing happiness will itself be a plan-laden judgment, resting on a plan-laden judgment about the character of 'right'. Last, Gibbard distinguishes signification from "quasi-signification." A predicate signifies its object only if the character of that predicate is fixed entirely by the concept, but a predicate only quasi-signifies its object if its character is partly a matter of how to live.³⁸

35. Gibbard, *Thinking How to Live*, 94.

36. Ibid., 119.

37. Ibid.

38. Ibid., 130. Moreover, it is hard to see how any expressivist who does not follow maverick expressivists like Ayer in dismissing talk of reference to normative properties altogether can avoid a conclusion along these lines. For expressivists aspire to accommodate the intuitions which underlie Moore's open question argument. This makes it hard to see how any expressivist could allow that claims about the reference of normative predicates are not normative claims. For it would then turn out to be an ordinary descriptive fact that 'right' refers, e.g., to maximizing happiness, in which case Moore's

We are now in a position to return to utterances involving anaphoric reference to moral properties.³⁹ When my interlocutor says that sacrificing the happiness of the few is right, his use of ‘right’ quasi-signifies the property of being right, but just which property that is remains a substantive moral (plan-laden, in Gibbard’s framework) question. When I respond by remarking that “sometimes sacrificing the few has that property, but sometimes it does not,” my use of ‘that property’ makes anaphoric reference to whatever property his use of ‘right’ signifies or (as in this case) quasi-signifies. So my use of ‘that property’ quasi-signifies the property of being right, which in my view may well not be the same property my interlocutor takes it to be. So such anaphoric reference will in effect involve substantive first-order judgment and hence will involve the expression of a suitable noncognitive attitude. The Ecumenical Expressivist should join with Non-Ecumenical Expressivists like Gibbard in holding that this is a first-order moral judgment. These views about reference and quasi-reference will inevitably be controversial, and I cannot defend them properly here. However, if there is a problem for expressivism here then it is a problem for all forms of expressivism and not just the ecumenical varieties, as the discussion of Gibbard should make clear.

The third clarification concerns pluralism about reasons. Ecumenical Expressivism holds that to judge that there is more reason to promote one outcome instead of another just is to approve (in the right way) of actions insofar as they have some single property and to believe that the promotion of that first outcome instantiates that single property to a greater extent than the second. Why, though, should we assume that the concept of a moral reason is unitary in this way? Why not instead allow that there might be irreducibly many sorts of moral reasons so far as the concept of a reason for acting goes? To some extent this objection rests on a misunderstanding of the notion of a “property” as deployed in the proposed account. The misunderstanding is to read ‘property’ in a metaphysically heavy way. As used here, ‘property’ should be read pleonastically, so that any predicate or combination of predicates

concerns about open questions will become as salient for the expressivist as they are for any form of reductive cognitivism.

39. In some contexts, such anaphoric reference might be to the property to which one’s interlocutor has speaker referred, although this will be unusual. For example, suppose you say “Wrongness just is being forbidden by the Ten Commandments. Working on the Sabbath is wrong.” I, the village atheist, sardonically reply by saying that “working on the Sabbath has that property, all right, but it isn’t wrong.” In this sort of case it may be clear enough that my anaphora refers back to the property to which you have speaker referred rather than the property to which your use of the word ‘wrong’ refers. However, such cases are exceptional, and the more standard case will be one in which my use of ‘that property’ refers back to the word reference of your use of ‘wrong’.

picks out a property. Wildly disjunctive and even contradictory combinations of predicates pick out a property in this sense.⁴⁰

However, the possibility of such disjunctive conceptions may provide cold comfort to friends of a certain sort of fundamental pluralism. For now return to the analysis of comparative claims. To say that there is more reason to go home than to stay is to express an attitude of approval of actions insofar as they have a certain property and the belief that going home instantiates this property more than staying. Now suppose that we try to accommodate pluralist intuitions by allowing that a given speaker's proattitudes are keyed to some disjunctive property. For example, perhaps I approve of actions insofar as they promote pleasure or love. Now suppose that I am a radical pluralist, and I think that no deeper property unites the properties of promoting pleasure and love. Such a view is indeed intelligible according to Ecumenical Expressivism, so Moorean intuitions are accommodated. However, this sort of pluralism can rapidly begin to look implausible. For suppose I say that there is more reason to X than Y. On the outlook under consideration this means I am expressing the belief that X is more pleasure-conducive-or-love-conducive than Y. How should we understand what it is for X to instantiate this disjunctive property to a greater extent than Y, though? We cannot plausibly analyze "more pleasure-conducive or love-conducive" in this context as equivalent in meaning to "either more pleasure-conducive-or-more-love-conducive." For that would absurdly imply that 'more reason' is not asymmetric—X might be more pleasure-conducive than Y, while Y is more love-conducive than X. In that case it would follow that there is simultaneously more reason to X than Y and more reason to Y than X. We should instead hold that such comparisons in terms of disjunctive properties are determinately true only when the better option dominates the worse option along at least one dimension and is at least equally strong in terms of the other dimension(s).⁴¹ When one option scores better

40. Nothing metaphysical hangs on this; if one does not like the term 'property' here then everything can be recast in terms of falling under descriptions to which the agent's preferences are sensitive.

41. An alternative would be to hold that such comparisons can be determinately true when one option dominates the worse option by a great deal on one dimension and is at most slightly worse on the remaining dimensions. One obvious problem with this suggestion is that it is irreducibly vague. A deeper problem is that even putting such vagueness to one side it seems to have counterintuitive implications in the present context for at least two reasons. First, suppose I hold a pluralistic theory with an extremely wide variety of types of reasons. Now suppose that A scores much better than B along one dimension but that B scores just slightly better than A on a whole host of the remaining dimensions. If the number of dimensions is great enough then we might intuitively want to say that there is more reason for B than A, but that contradicts the proposed semantics. Second, we may think that some sorts of reasons are simply more important so that doing better along some dimensions matters more than doing better along others. For example, the

along one dimension but the other option scores better along another dimension, the claim that there is more reason for either will lack a determinate truth-value. This, however, means that this sort of pluralism entails a very strong and implausible sort of incommensurability. For example, suppose that A would promote a minuscule amount of pleasure but not promote love at all while B would promote huge amounts of love but even less pleasure than A. On the view under consideration, it would implausibly follow that there is no determinate fact as to which is more choice-worthy.

If this were a problem specifically for Ecumenical Expressivism then it should give us pause. However, all one needs to get the preceding argument going is the intuitive idea that for there to be a reason for an option is for the option to be favored by some consideration, and this is an idea that hardy moral realists like Thomas Scanlon and Jonathan Dancy embrace.⁴² Once this much is allowed, it is a small step to the conclusion that for a fact to be a reason for an option is for that fact to stand in the favoring relation to that option. Finally, once we allow that for there to be reason to X is for there to be facts which stand in the favoring relation to X, it is hard to avoid the conclusion that for there to be more reason to X than Y just is for X to be more favored than Y. In that case, though, we need to know more about the favoring relation. Either the favoring relation is disjunctive or not. If it is disjunctive then the argument of the preceding paragraph goes through. If, however, the favoring relation is nondisjunctive then we can avoid incommensurability but only by avoiding a radically pluralist understanding of the favoring relation itself. Since the assumptions needed to get this argument going are minimal and shared by a number of other views, it seems not to be especially a problem for Ecumenical Expressivism.⁴³

fact that my action scores very well indeed along the dimension of preventing mild pains (e.g., it prevents thousands of toe stubblings) may not make it better than an alternative course of action which scores only slightly better in terms of lives saved (the other action saves one life whereas my crusade against toe stubblings would save none). Intuitively, we might well want to say that there is more reason for the second action than the first, but the proposed semantics entails the opposite verdict. Thanks to an anonymous referee for pressing me on this point.

42. T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), 17; and Jonathan Dancy, *Ethics without Principles* (Oxford: Oxford University Press, 2004), 29.

43. One might suppose that nonnaturalist versions of cognitivism could escape this argument, but this is an illusion. For in the relevant sense the nonnaturalist is not a fundamental pluralist. The nonnaturalist presumably thinks there is only one nondisjunctive nonnatural property picked out by ‘good’—the property of goodness. A nonnaturalist could, I suppose, hold the more baroque view that ‘good’ refers to a disjunction of two nonnatural properties— g_{1} and g_{2} . However, in that case the argument given here will go through once again, and the nonnaturalist will be committed to incommen-

However, such “companions in guilt” replies are never fully satisfying, and I admit that the argument from pluralism does pose a challenge to the version of Ecumenical Expressivism outlined so far. One interesting strategy for meeting this challenge comes from an unexpected quarter—the idea of a dispositional version of Ecumenical Expressivism. Dispositional accounts in general analyze concepts in terms of the responses certain sorts of subjects would be disposed to have under certain sorts of conditions. Perhaps the most famous example of a set of concepts thought to be ripe for a dispositional analysis are color concepts. However, an impressive array of philosophers have also argued that normative concepts (including moral concepts) are rightly analyzed in dispositional terms. Historically, the idea can arguably be found in various guises in David Hume’s “common point of view,” Adam Smith’s idea of an “impartial spectator,” and Immanuel Kant’s “Kingdom of Ends.”⁴⁴ More recently, the same idea can be found in Roderick Firth’s Ideal Observer theory, as well as in the work of David Lewis, Mark Johnston, and Michael Smith.⁴⁵ With the possible exception of Hume’s view, these accounts have all been offered as forms of cognitivism. However, in a more ecumenical spirit, we might try to see whether expressivists could not also embrace some of the insights of dispositional accounts. The idea would be to agree with those cognitivist defenders of traditional dispositional accounts that the concept of what we ought to do is best understood in terms of what a suitably idealized subject would want us to do, but disagree with the idea that the concept of a “suitably idealized subject” can be given a plausible reductive analysis. Nor, the argument continues, should we allow that the idea of a suitably idealized subject denotes an irreducible property. Instead we should incorporate a dispositional element into our expressivist account, perhaps as follows:

“There is moral reason to X” expresses (a) an attitude of approval of a certain kind toward actions insofar as they would garner approval from a certain sort of subject and (b) a belief that X

surability after all. Or at least, she will be committed to incommensurability if she accepts the platitude that one thing is better than another just in case it instantiates more goodness. Thanks to Elinor Mason for useful discussion here.

44. David Hume, *A Treatise of Human Nature*, 2nd ed., ed. L. A. Selby Bigge, rev. P. H. Nidditch (Oxford: Clarendon, 1978); Immanuel Kant, *The Groundwork of the Metaphysics of Morals*, trans. Lewis White Beck (New York: Macmillan, 1990); and Adam Smith, *The Theory of Moral Sentiments* (Indianapolis: Liberty Fund, 1984).

45. Roderick Firth, “Ethical Absolutism and the Ideal Observer,” *Philosophy and Phenomenological Research* 12 (1952): 317–45; David Lewis, “Dispositional Theories of Value: II,” Mark Johnston, “Dispositional Theories of Value: III,” and Michael Smith, “Dispositional Theories of Value: I,” *Supplement to the Proceedings of the Aristotelian Society* 63 (1989): 113–37, 139–74, and 89–111.

would garner approval from that sort of subject.

“Morally, there is most reason to X” (“Morally, one ought to X”) expresses (a) an attitude of approval of a certain kind toward actions insofar as they would garner approval from a certain sort of subject and (b) a belief that X would garner the approval of that sort of subject to a greater extent than any of the available alternatives.

“The fact that X would be P is a moral reason to X” expresses (a) an attitude of approval of a certain kind toward actions insofar as they would garner the approval of a certain sort of subject and (b) a belief that that sort of subject would approve of X qua its being P.

How might this dispositional version of Ecumenical Expressivism help make sense of pluralism without implying implausible forms of incommensurability? The key idea is that how much reason there is in favor of a given action is a function of the preference ordering of the relevant sort of subject.⁴⁶ So long as the ideal subject’s preferences are not indeterminate there will be a fact of the matter as to whether any two options are equally preferred or whether one is more preferred than the other. So there need not be any incommensurability. However, there is still room for a kind of pluralism about reasons. For a suitably idealized subject might find actions attractive in virtue of a plethora of different features, in which case there will be a heterogeneous plurality of kinds

46. An alternative approach here would be to maintain that one reason is stronger than another just in case the relevant sort of idealized subject would approve of acting on the former when it conflicts with the latter. This approach faces two problems. First, consider a case in which there are several reasons for an action and only one reason against it. Suppose that the weaker reasons, when taken together, outweigh the stronger reason. An ideal subject presumably would approve of acting on the weaker reasons because of their aggregate weight. In that case, though, the proposed account seems to entail that each of those weaker reasons is actually stronger than the stronger reasons, which contradicts our initial description of the case. Second, there may be reasons which cannot coexist, in which case the proposed account provides no account of how one reason could be stronger than the other. To take an exotic example, suppose that I promise to meet you for lunch so long as there is no reason whatsoever not to do so. We might want to hold that my promise does generate a reason for me to meet you for lunch but only in the extremely rare situation in which there is no reason whatsoever against meeting you. How can we compare the reason generated by this promise with other reasons? The proposed account seems to imply that we imagine a counterfactual situation in which the two reasons conflict and let the approval of an ideal subject in that case decide the matter. By hypothesis, though, the reason generated by such an odd promise can never conflict with any other reasons. This, however, intuitively does not entail that the reason cannot be weighed against other reasons. Intuitively, the reason generated by my promise is much weaker than our reasons not to kill innocent people or cause gratuitous suffering, for example. Thanks to an anonymous referee for pressing me on this point.

of reasons for action. For example, an idealized subject might prefer actions insofar as they promote pleasure and also prefer actions insofar as they foster scientific knowledge. So on a dispositional version of Ecumenical Expressivism, we can have a recognizable pluralism about reasons without implausible forms of incommensurability.⁴⁷

Of course, a dispositional version of Ecumenical Expressivism will run into objections specific to the incorporation of a dispositional element. Most notably, dispositional analyses face a Euthyphro-style challenge according to which ideal subjects prefer options because of their independently normative properties rather than its being the case that those options have those normative properties because they would be preferred by ideal subjects. To some extent this objection rests on realist assumptions that anyone with expressivist leanings would not find persuasive, but there are perhaps more accommodating moves an expressivist can make. However, this is not the place to determine which version of Ecumenical Expressivism is more plausible—the dispositional or nondispositional version. Each species of Ecumenical Expressivism has its own characteristic advantages and disadvantages, but my aim here is to investigate the characteristic advantages of the genus to which they belong. For expository reasons I shall here continue to work in the framework of the simpler nondispositional version, but this should not be taken to prejudice the issue

47. One might have thought that we could capture all of these advantages within the framework of the original nondispositional version of Ecumenical Expressivism. The idea would be that dispositional theories can simply be understood as constituting another substantive normative perspective which someone could hold within the framework of our original version of Ecumenical Expressivism. However, this approach unfortunately would not really preserve pluralism about reasons. The crucial move, when it comes to preserving pluralism in this context, is to avoid being forced to the conclusion that there is really only one reason for action, namely, that such-and-such sort of subject would want one to perform the action. The dispositional version of Ecumenical Expressivism avoids this conclusion by holding that for a fact to be a reason for a given action is for that fact to be a fact to the effect that the action has some feature such that an idealized subject would approve of the action qua its having that feature. However, it is hard to see how a dispositional first-order normative view situated within the nondispositional version of Ecumenical Expressivism could avoid this conclusion. For recall that on the nondispositional version of Ecumenical Expressivism, for a fact to be a reason for action just was for that fact to be a fact to the effect that the action has “that property” where ‘that property’ anaphorically refers to the feature in virtue of which the speaker approves of actions. If, however, the speaker approves of actions just insofar as they would garner the approval of such-and-such a subject, as a speaker would if she accepted a dispositional first-order normative view, then it would follow that on her view there really is only one reason for action—namely, that an action would garner the approval of such-and-such a subject. It is hard to see how to preserve pluralism without building a dispositional element into the analysis of normative concepts itself.

of which version is more promising. Indeed, I actually favor the more complex dispositional version of the theory.⁴⁸

V. AVOIDING THE FREGE-GEACH PROBLEM “ON THE CHEAP”

How might Ecumenical Expressivism deal with logically complex sentences? First, consider negation. Just as with atomic moral judgments, negation is best understood in terms of a belief whose content is a function of the agent’s proattitudes:

“There is no moral reason to X” expresses (a) an attitude of approval toward actions insofar as they have a certain property and (b) the belief that X lacks that property.

One virtue of this account is that it can easily accommodate scope distinctions that have caused problems for traditional Non-Ecumenical versions of expressivism. In particular, we must distinguish someone’s thinking that there is no moral reason to X from thinking that there is moral reason not to X. Non-Ecumenical forms of expressivism have had trouble with this distinction because they work only with desires. For Non-Ecumenical Expressivists seem stuck explaining what it is to think that there is no moral reason to perform a given action in terms of indifference to the performance of that action. However, indifference also is the Non-Ecumenical Expressivist’s best shot at explaining

48. Another interesting issue is what, on the dispositional version of Ecumenical Expressivism, we should say about someone who approves in the right way of more than one observer at the same time. The possibility of such complex normative psychologies is in one sense an advantage of the dispositional account. For it provides a sort of structure with which we could try to capture what Nagel memorably calls the “fragmentation of value” and Sidgwick’s “dualism of practical reason.” However, the possibility of such normative psychologies also raises a problem for the semantics. For now we can no longer comfortably hold that normative utterances express beliefs which make anaphoric reference back to the relevant sort of observer—that may simply be a failed definite description in the case in which a speaker approves of more than one observer in the right (and nonderivative) way. Fortunately, it seems to me that the dispositional account can be amended in a natural way to deal with this sort of phenomenon. For we can instead hold that normative utterances express (a) approval (of the right sort) of a set of observers (this set may or may not include more than one member) and (b) a belief which makes anaphoric reference back to what the set of all such observers would converge in approving of, disapproving of, insisting on, or whatever (depending on what the normative predicate being analyzed is). Such an analysis seems to fit well with how we could understand the all-things-considered judgments such a fragmented agent might nonetheless make and will also allow us to avoid the Frege-Geach puzzle in pretty much the same way developed in the text; we will get contradictions in belief in the right cases in just the same way, and the only difference will be in the contents of the beliefs which are contradictory. However, developing the details of the needed semantics would take us too far afield, and so for present purposes I must set these issues to one side. Many thanks to Mark van Roojen for useful discussion.

what it is for a speaker to suspend judgment as to whether there is any moral reason to perform a given action. Non-Ecumenical Expressivism therefore runs a serious risk of blurring the distinction between thinking that there is no reason to X and having no view as to whether there is reason to X. For Ecumenical Expressivism, this distinction is easy. To think that there is no moral reason to X is to approve of actions in virtue of their having a certain property and to believe that Xing does not have that property. Whereas not to have a view as to whether there is moral reason to X would be to approve of actions in virtue of their having a certain property but to suspend judgment as to whether X has that property.⁴⁹

However, one might object that this account of negation implausibly makes a certain sort of radical nihilism impossible. For on this account even to think that there is no moral reason to perform an action, one must suppose it lacks some property such that any action which did have that property would be one there is moral reason to perform. So the radical nihilist position according to which there could not possibly be any moral reason to do anything seems ruled out by linguistic fiat. This seems problematic; radical forms of nihilism may be mistaken, but they should not be ruled out of court so quickly.⁵⁰ Here it is essential to emphasize that ‘property’ as deployed in the present account is used so broadly as to allow that logically contradictory descriptions pick out properties in the relevant sense. As long as this is allowed, we can understand how someone might intelligibly think that nothing could possibly be a moral reason for action. For a given speaker might approve of actions insofar as they are commanded by God but also believe that the very idea of God is somehow logically contradictory. In that case the agent would think not only that nothing is worthwhile but that nothing could be worthwhile. It remains true that anyone who thinks that a given consideration is not a moral reason must have some idea of what would have to be true of a consideration in order for it to be a moral reason, even if the relevant condition is logically contradictory. However, this is not objectionable. For it is plausible to suppose quite generally that the ability to judge that something is F requires some

49. Actually, there is an important caveat to my claim in the text. For someone could suspend judgment in normative matters more globally, not taking on any substantive first-order normative views at all. I discuss this interesting possibility in the appendix.

50. Walter Sinnott-Armstrong raises a similar objection to Simon Blackburn’s version of expressivism. See Walter Sinnott-Armstrong, “Nihilism and Scepticism about Moral Obligations,” *Utilitas* 7 (1995): 217–36; Nicholas Unwin, “Quasi-Realism, Negation, and the Frege-Geach Problem,” *Philosophical Quarterly* 49 (1999): 337–52.

conception of what would be involved in its being F. Someone with no idea how to distinguish Fs from non-Fs has no conception of F-ness.⁵¹

The final ingredient in the proposed solution to the Frege-Geach problem is a suitable definition of validity. Standard accounts hold that an argument is valid just in case it is impossible for its premises to be true while its conclusion is false. Since some versions of Ecumenical Expressivism hold that moral utterances are not truth-apt, we cannot take this account over as it stands. We can, however, adopt a close cousin:

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.⁵²

It is crucial to the tenability of this definition that it ranges over all possible believers. Suppose I am a utilitarian, so I approve of actions insofar as they maximize utility. In that case it would be contradictory for me to think that an action maximizes utility yet is not morally right. However, the inference “X maximizes utility, therefore X is morally right” had better not be valid, on pain of contradicting Moorean open question intuitions. Fortunately, on the account offered here this argument is invalid. For while it is true that a utilitarian who believes both that an action maximizes utility and that the action is not morally right is thereby caught in an inconsistency, it is not true that any possible believer who believes that a given action maximizes utility and that the action is not morally right is thereby guaranteed to be caught in an

51. Actually, I do not need the strong thesis that someone must have some idea of what would be sufficient for something to be F in order to have some conception of F for purposes of the argument in the text. For in order to count as a nihilist someone needs only to judge that some necessary condition for there being a reason to perform an action could never be met. So in order to accommodate the possibility of nihilism I need only to suppose that having a conception of F involves at least some conception of necessary conditions on something’s being F. While I think that the stronger thesis is plausible, I do not need to defend it for present purposes. Thanks to an anonymous referee for pressing me on this point.

52. The ‘thereby’ in this definition can be read in stronger or weaker ways. On a strong reading, both the premises and the denial of the conclusion must figure in the explanation of why the person has beliefs with inconsistent contents. This reading of the ‘thereby’ yields a logic closer to traditional relevance logics than classical logic, since it will not entail that a contradiction entails absolutely everything. On a weaker reading, it is enough that accepting the premises and denying the conclusion are sufficient for having beliefs with inconsistent contents. This reading of ‘thereby’ yields something much more like classical logic. Thanks to Peter Milne for bringing this to my attention.

inconsistency. Anyone not committed to utilitarianism can accept the premise and reject the conclusion without inconsistency.

Before seeing how this definition of validity, when combined with Ecumenical Expressivism, allows us to avoid the Frege-Geach problem, it is worth pausing to consider how someone might object to the proposed definition of validity itself. Perhaps the most interesting objection contends that the definition implies the validity of arguments which go from a premise with a conventional implicature to a conclusion which asserts that implicature; for example:

1. Sidney is a philosopher yet funny.

So, 2. It is in some way unexpected or surprising that a philosopher would be funny.

The objection insists that this argument is not valid, yet anyone who accepts the premise and at one and the same time denies the conclusion is thereby caught in a contradiction. The short answer to this objection is that in the sense of ‘accepts’ in play in the proposed definition of validity someone can count as accepting a premise in which a conventional implicature word is used even if they do not accept the associated implicature. So in the intended sense of ‘accepts’ someone who believes that Sidney is a philosopher and funny thereby counts as accepting premise 1. For purposes of the definition of validity on offer, we can even stipulate that the sense of ‘accepts’ in use here is a technical one, so that if it does not fit well with absolutely all central aspects of ordinary usage that need not be an objection. So long as the resulting definition of logical validity is plausible, the fact that the notion of “accepts” in play in that definition is to some extent a technical term need not be seen as a problem.

In fact, the proposed reading of ‘accepts’ actually fits well with at least one of its meanings in ordinary English. Semantic intuitions on this question seem (in my admittedly anecdotal experience) divided over whether one must accept the conventional implicature of a sentence to count as accepting it. Suppose that Hilary thinks that it is not at all surprising that a philosopher is funny, and I also happen to think that Sidney is a philosopher and is funny. Now suppose that I ask, “Does Hilary accept that Sidney is a philosopher yet funny?” Some people react to such questions by saying, “Yes, but he wouldn’t put it that way,” while others insist that Hilary simply does not accept that Sidney is a philosopher yet funny. Most people think that it is a peculiar question and are not sure what to make of it, which suggests that intuitions here are not deeply held. The Ecumenical Expressivist account emphasizes

those strands of ordinary thought which favor the answer “Yes, but he wouldn’t put it that way.”⁵³

In any event, putting technical concerns about conventional implicature to one side, how might the proposed account of validity help us deal with the Frege-Geach problem? Consider how this account of validity applies to arguments employing moral predicates. Begin with the simplest form of argument, reiteration—“p, therefore p.” Let ‘p’ be an atomic moral utterance such as “There is moral reason not to have abortions.” On the proposed conception of validity, the argument is valid just in case any agent who accepts the premise but denies the conclusion would thereby be guaranteed to be caught in an inconsistency in belief. Since the denial of the conclusion would simply be “There is no moral reason not to have abortions” the question is whether anyone who accepts A, “There is moral reason not to have abortions,” and who accepts not-A, “There is no moral reason not to have abortions,” is thereby caught in an inconsistency. On the model proposed here, any possible agent who accepts A and accepts not-A both thinks that refraining from having abortion has a certain property and thinks that abortion does not have that property. This clearly is an inconsistency of a familiar kind—inconsistency in belief.

It is straightforward to see how this account can be extended to deal with other logically complex sentences. The general scheme for any logically complex sentence in which ‘moral reason’ appears is as follows. Let ‘p’ stand for a logically complex sentence in which “there is a moral reason to X” (but no other moral predicates) appears. An utterance of ‘p’ expresses (a) the agent’s approval of actions insofar as they have a certain property and (b) the agent’s belief that p*, where p* is identical to p save that all occurrences of the form “there is a moral reason to X” are replaced by “Xing has that property,” where ‘that property’ is to be understood as making anaphoric reference to the property mentioned in a. The same strategy can be extended to other sentences in which ‘moral reason’ is used in fairly obvious ways. For example, consider a logically complex utterance ‘q’ in which a

53. It is worth pausing to clarify the notion of “accepts” in play here. One natural but misguided reading of the relevant notion is that someone accepts a sentence ‘p’ just in case they are poised for sincere utterance of ‘p’. However, the discussion of conventional implicature reveals that this cannot be the right reading of ‘accepts’ in this context. For if someone says “even philosophers get this joke” and believes that philosophers get the joke but also does not think that it is surprising that philosophers would get a joke then this smacks of insincerity. Yet in the sense of ‘accepts’ in play here one need not accept the implicature to count as accepting that even philosophers get the joke. Simply believing that philosophers get the joke is enough. The idea of acceptance in play is one we can get via disquotational and asking whether someone believes that which is disquoted. Putting indexicals to one side, someone accepts a given sentence ‘p’ just in case they believe that p.

phrase of the form “the fact that X is p is a moral reason to X” occurs. Any such utterance expresses (a) the agent’s approval of actions insofar as they have a certain property and (b) the agent’s belief that q*, where q* is identical to q save that all occurrences of the form “the fact that X is p is a moral reason to X” are replaced by “the fact that Xing has that property,” where ‘that property’ once again is to be understood as making anaphoric reference to the property mentioned in a.⁵⁴

It should be clear by now how this account can explain the validity of arguments with moral predicates quite generally. Consider the standard case of modus ponens:

1. There is moral reason not to lie.
2. If there is moral reason not to lie, then there is moral reason not to encourage your little brother to lie.
3. Therefore, there is moral reason not to encourage your little brother to lie.

On the proposed account, the acceptance of premise 1 requires the belief that there is some not lying that has “that property,” where the use of ‘that property’ denotes the property in virtue of which the speaker disapproves of actions quite generally. The acceptance of premise 2 involves the belief that if refraining from lying has that property then refraining from getting your little brother to lie also has it. To deny

54. One further complication which I ignore in the text arises over contexts in which normative predicates figure in the contents of a propositional attitude attributed to someone (e.g., when I say “she believes that abortion is wrong”). The point is that in these contexts we are not typically assuming that the person to whom we attribute the propositional attitude associates the same cluster of descriptive properties with a given normative predicate that we do. For example, when I, as a utilitarian, say that Jones believes that abortion is wrong I need not be presuming that Jones believes that abortion fails to maximize happiness. So we should understand such attributions in terms of the attribution of a suitable belief-desire pair without taking a position on whether the speaker shares our conception of the good. So when I say that she believes that abortion is wrong I am making a purely descriptive claim, namely, that she has the belief that abortion is wrong. It turns out (though a given speaker may not realize this) that the belief that abortion is wrong is really in another sense a belief-desire pair—a general proattitude of the right kind and a belief (in a strict Humean sense of “belief”) which makes suitable anaphoric reference back to the content of that proattitude. The only real difficulties emerging for Ecumenical Expressivism on this front arise when we combine ascriptions of normative beliefs with claims about the truth of what the subject believes, which should allow us to infer a normative conclusion. For example, we have inferences like “She thinks abortion is wrong, and everything she thinks is true, so abortion is wrong.” However, I shall not here go into the details of how Ecumenical Expressivism is best extended to deal with these further cases. For this would require a full theory of truth (for a start) and would therefore take us too far afield from an outline of the basic ideas and advantages of the Ecumenical approach. I hope to return to this important cluster of issues elsewhere, though. Thanks to Timothy Williamson and John Hawthorne for pressing me on this point.

premise 3, though, would be to believe that not encouraging little brother to lie does not have that very property. This is an inconsistent set of beliefs, and so the argument is logically valid on the proposed account. The general strategy works across the board in an elegant way, no matter how complicated the judgments. In contrast with a “logic of attitudes” approach, there is no pressure to generate increasingly sophisticated higher-order attitudes to model increasingly complex sentences. The Frege-Geach problem simply does not arise in the first place.

VI. MEETING THE TWO CONSTRAINTS

We should now see whether Ecumenical Expressivism does better than its rivals in meeting the two constraints canvassed at the end of Section III (the inconsistency constraint and the rational inference constraint). Expressivists have often held that we must broaden our conception of inconsistency and rely on a logic of attitudes to make sense of the validity of moral arguments. Critics have argued that such accounts inevitably confuse moral mistakes with logical mistakes. Furthermore, such accounts arguably conflate the kind of inconsistency involved in sentences associated with Moore’s paradox (e.g., “p but I don’t believe it” and “I believe that p but not p”) with logical inconsistency. Moorean sentences do not involve logical inconsistency—they could be true. If, however, we analyze moral arguments in terms of a logic of attitudes then this distinction becomes problematic. Mark van Roojen makes this point. Here is a useful example from van Roojen’s discussion.⁵⁵ The general idea of the logic of attitudes approach is that it is inconsistent to have a higher-order attitude toward an attitude while nonetheless having that attitude. The following sentence is inconsistent in just that sense: “It is wrong for me to believe that my father is unfaithful to my mother but my father is unfaithful to my mother.” Clearly this sentence is not logically inconsistent, but it expresses both an attitude of disapproval toward another attitude and that very attitude (the infidelity belief) and so comes out as inconsistent on a logic of attitudes approach.

The same basic problem infects the more recent “commitment-based” approaches to the Frege-Geach problem. For if to say “Believing that p is bad” is to express one’s commitment to avoiding the belief that p, then avowing “Believing that p is bad” will wrongly come out as inconsistent in the relevant sense with p, which in turn implies that “The belief that p is bad, so not-p” is valid. Gibbard’s version of the commitment-based approach holds that an argument is valid when accepting the premises and rejecting the conclusion is incompatible with becoming fully decided on all matters (factual and normative) without flouting one’s

55. See Mark van Roojen, “Expressivism and Irrationality,” *Philosophical Review* 105 (1996): 311–35, 320–21.

own commitments. The point is that on Gibbard's own account, "Believing that p is bad" expresses a commitment to avoid believing that p.⁵⁶ Clearly, that commitment can never be fulfilled so long as the agent believes that p. So anyone who accepts that believing that p is bad in Gibbard's sense yet believes that p cannot live up to her commitment not to believe that p without changing her mind. So Gibbard's account implies that "believing that p is bad" is logically inconsistent with "p," which in turn incorrectly implies that "believing that p is bad, so not-p" is valid.

By contrast, Ecumenical Expressivism avoids these problems from the outset. Ecumenical Expressivism does not require a separate logic of attitudes and does not conflate logical inconsistency with Moorean paradox.⁵⁷ For those problem cases arose only because inconsistency was being understood in terms of disapproving of one's own attitudes or failing to live up to one's commitments. Ecumenical Expressivism eschews these approaches and understands the relevant inconsistency as ordinary inconsistency in belief. Moreover, unlike its rivals, Ecumenical Expressivism can meet what I earlier called the "rational inference constraint." Consider the following argument:

1. Lying is bad.
2. If lying is bad then God did not lie.
3. Therefore, God did not lie.

The argument's validity is straightforward. Accepting premise 1 commits a speaker to believing that refraining from lying has a certain property (the one to which his approval is keyed), accepting premise 2 commits a speaker to believing that if refraining from lying has that property

56. In his more recent work, Gibbard holds that atomic normative judgments like "believing that p is bad" express plans, but it seems reasonable enough to suppose that a plan not to believe that p amounts to a commitment to avoid believing that p.

57. An anonymous referee suggests that while Ecumenical Expressivism does not need a logic of attitudes to deal with the sorts of arguments canvassed in the text, it still will require such a separate logic for arguments like "I have most reason to keep my promise, I cannot keep my promise if I go fishing, so I do not have most reason to go fishing." If we read 'most reason' here as allowing ties, so that I can have most reason to keep my promise and most reason to do something else then the argument is not valid on any account. So the idea must be that we should read 'most reason' as precluding ties. In that case, though, the argument is valid without any recourse to a logic of attitudes. For someone who accepts the premises and denies the conclusion of this argument must, according to Ecumenical Expressivism, simultaneously believe (a) that keeping his promise is an action available to him which has more of a certain property than any of the other actions available to him, (b) that going fishing is an action available to him which is incompatible with keeping his promise, and (c) that going fishing has the property in question to a greater extent than any of the other alternatives available to him. These beliefs entail that keeping his promise simultaneously instantiates the property in question more and less than going fishing, which is surely contradictory. So the argument is valid according to the account laid out in the text.

then God did not lie, while denying premise 3 commits a speaker to believing that God did lie. These three beliefs are inconsistent in a perfectly familiar sense. Moreover, we can imagine circumstances in which it might be reasonable to infer premise 3 from premises 1 and 2. This is an important advantage of Ecumenical Expressivism over traditional forms of expressivism. On Blackburn's most recent account and the sort of account favored by Price and Gibbard, embedded contexts express the speaker's "commitments."⁵⁸ For example, accepting a disjunction is being committed to accepting the attitude expressed by one of the disjuncts if the attitude expressed by the other disjunct becomes "untenable." These commitments underwrite the making of various inferences. Gibbard offers a similar account. Gibbard suggests that deciding what to do and deciding what one ought to do are really one and the same thing. In light of this identification he constructs the idea of a "hyperplan," in which one has a universal plan of life, a plan that "covers any occasion of choice one might conceivably be in, and for each alternative open on such an occasion, to adopt the plan involves either rejecting that alternative or rejecting rejecting it."⁵⁹ He then suggests that logically complex evaluative and normative utterances express one's commitment to rejecting certain universal plans of life. For example, to say that I ought to X or Y is to express my commitment to rejecting universal plans of life in which I neither X nor Y. Gibbard explains logically complex contexts in terms of commitments to rejecting certain combinations of attitudes.

The idea is that the inconsistency that underwrites the validity of moral arguments is the inconsistency of accepting a commitment not to have certain combinations of attitudes but having them anyway. Now consider again this argument:

1. Lying is bad.
2. If lying is bad then God did not lie.
3. Therefore, God did not lie.

58. For Blackburn's earliest account, see Simon Blackburn, *Spreading the Word* (Oxford: Oxford University Press, 1984). He later switches from talk of what one attitude "involves" to a theory employing higher-order attitudes in *Essays in Quasi-Realism*. For his most recent account, see Simon Blackburn, *Ruling Passions* (Oxford: Oxford University Press, 1998). See also Price, "Semantic Deflationism and the Frege Point," as well as Gibbard, *Thinking How to Live*. For criticisms of Blackburn's accounts, see Robert Hale, "The Compleat Projectivist," *Philosophical Quarterly* 36 (1986): 65–84, and "Can There Be a Logic of Attitudes?" in *Reality, Representation and Projection*, ed. J. Haldane and Crispin Wright (Oxford: Oxford University Press, 1993), 337–63; G. F. Schueler, "Modus Ponens and Moral Realism," *Ethics* (1988): 492–500; van Roojen, "Expressivism and Irrationality"; and Nicholas Zangwill, "Moral Modus Ponens," *Ratio* 5 (1992): 177–93. Blackburn responds to many of these worries in developing and revising his account in all of the works by him cited above.

59. Gibbard, *Thinking How to Live*, 56.

Suppose for the sake of argument that the commitment-based approach can explain the validity of such arguments. As Dorr has argued, a further problem remains.⁶⁰ Suppose I accept premise 2 and suspend judgment with regard to premise 1. I then come to accept premise 1 in a way that is not irrational. I then infer premise 3 from premises 1 and 2. It seems like this could be a perfectly rational inference. However, the reason that the inference would go through would seem to be because I have adopted a noncognitive attitude that gives me a practical reason to accept the conclusion. For on Non-Ecumenical accounts, the acceptance of an atomic moral sentence like premise 1 just is the adoption of a suitable noncognitive attitude—roughly, a desire that people not lie or a preference that people internalize norms forbidding lying or something along these lines. The inference would have to go from an attitude of disapproval of lying and a commitment not to combine that attitude with a belief that God lied to the conclusion that God did not lie. It is hard to see how this inference could be epistemologically faultless. For the fact that I have committed myself not to accept that God lied while disapproving of lying and the fact that I disapprove of lying do not provide me with any epistemically respectable reason to infer that God did not lie. Beliefs that are motivated by one's attitudes and commitments as to what attitudes and beliefs one is willing to hold are paradigm cases of wishful thinking. Epistemologically, beliefs should be based on perceptual evidence or other beliefs, not the result of attitudes or commitments. Nor is this problem a minor or localized one; it seems that on the commitment-based account any inference of the logical form “p, if p then q, so q,” where ‘p’ is an evaluative thesis and ‘q’ is a descriptive one, will be such that anyone who makes such inferences is guilty of wishful thinking. Ecumenical Expressivism simply does not face this problem. For on an Ecumenical Expressivist account the conclusions of such arguments are inferred from other beliefs.

VII. CONCLUSION

Metaethical debates have become mired in disputes about the extent to which moral judgments are belief-like and the extent to which they are desire-like. This debate rests on a false dichotomy. Ecumenical Views insist that moral utterances express both beliefs and desires. Just as Ecumenical Cognitivists can and should agree with expressivists that motivation is “internal” to moral judgment, Ecumenical Expressivists can and should agree with cognitivists that belief is “internal” to moral judgment. Doing so immunizes expressivism against the Frege-Geach problem without provoking any further unfortunate side effects.

60. Cian Dorr, “Expressivism and Wishful Thinking,” *Nous* 33 (2002): 558–72.

Appendix

What, given Ecumenical Expressivism, should we say about someone who suspends judgment on all claims of the form “X is required” and “X is not required”? Such a person might well not approve of actions insofar as they had the relevant property—indeed there may in their case be no “relevant property.” For the point is that such a person has no fixed view about what would or would not make actions required or permissible, or indeed about what would suffice for a pro tanto reason to perform them. Most ordinary agents do not suspend judgment in this way, but someone impressed by some forms of philosophical skepticism and convinced of a suitably realist metaethical analysis of normative terms might well suspend judgment in just this way. This would create no unique problems were it not for the fact that such a person might also make various conditional moral judgments. For example, such a person might well judge that if passive euthanasia is sometimes required then active euthanasia is sometimes required too. The trouble is that on the most natural extension of the account developed in the text, I must hold that to make such conditional judgments already constitutively involves approving of actions insofar as they have some property or other, which itself seems to be sufficient to hold that in some sense the person does hold various substantive (nonconditional) first-order views about which actions are (or, at any rate, would in the right circumstances be) required. But that seems absurd; making such conditional judgments should not presuppose any substantive first-order views about when actions really are or would be required.

Technically, the ascription of such substantive first-order views does not quite follow from making such conditional judgments on the proposed view. For an agent might approve of actions insofar as they have certain properties but have no beliefs about when the relevant properties are (or would be) actually instantiated, instead holding only conditional beliefs about the relevant property (e.g., the belief that if passive euthanasia has it, then active euthanasia has it too). Since a substantive moral belief on my account requires both proattitude and a suitable anaphoric belief, this preserves logical space for someone to embrace the relevant conditional theses without taking on any substantive first-order views whatsoever. For example, I might believe that if passive euthanasia has the relevant property then so does active euthanasia without having the slightest idea whether passive euthanasia actually does have it, and indeed without having any idea whether anything does or does not have it.

This reply has the feel of a cheat, though. For there does seem to be a sense in which such a speaker does have a normative outlook, and it is given by his proattitudes; it is just that he is unaware of the details of his normative outlook. Perhaps we could put the point in terms of the agent’s having a conception of the good without being willing actually to deploy that conception. Fortunately, we can meet the objection more directly, albeit with a revision of the account developed in the main body of the text. I am inclined to hold that the possibility of global

agnosticism about substantive first-order questions suggests that non-atomic moral judgments are multiply realizable in an interesting way. More specifically, there are two ways to judge, for example, that if passive euthanasia is sometimes required then so is active euthanasia. First, there is the standard way in which most ordinary people actually make the judgment. To make the judgment in this way is (on the ideal observer version of the theory) to approve of actions insofar as they would garner the approval of a certain sort of observer and to believe that if passive euthanasia would sometimes garner the approval of such an observer then so does active euthanasia. Second, though, there is the way in which someone who suspends judgment about all substantive nonconditional first-order moral judgments would make such a conditional judgment. Here, I suggest that it is most plausible within the framework of Ecumenical Expressivism to understand such an agent as taking a stand against the approval of certain sorts of observers—those observers who would simultaneously approve of passive euthanasia but at one and the same time not also approve of active euthanasia, say. In the Ecumenical framework, this will amount to the agent's adopting a perfectly general noncognitive attitude, here an attitude of refusal—refusal to approve of an observer unless it has certain features and the belief that such features (once again we have a belief with anaphoric reference back to the content of a noncognitive attitude) preclude simultaneously approving of passive euthanasia while not also approving of active euthanasia. The fact that making conditional judgments will still involve such anaphoric beliefs provides the resources with which to solve the Frege-Geach problem in much the way proposed in the text. The inclusion of a belief will still allow the proposed definition of validity to work in much the same way, without any reliance on a logic of attitudes or the like, or any of the notorious problems associated with other Non-Ecumenical approaches.

The only problem, so far as the technical details of the solution to the Frege-Geach puzzle go, would arise if it were possible for someone to accept a conditional premise in the way characteristic of someone who is agnostic on all substantive nonconditional first-order normative claims, while at one and the same time accepting a nonconditional substantive first-order premise in the more standard way. For in this sort of case, if it were possible, the belief expressed in the major premise would not "hook up" logically in the right way with the belief expressed by the conditional premise to explain the validity of the argument. However, such cases are not possible on the theory on offer here, properly understood. For if someone does have a normative outlook at all, as they must to accept an atomic judgment like passive euthanasia is right, then they can only count as making the relevant conditional judgment if they have the right sort of belief about that observer. Refusing to approve of certain sorts of observers can play a role in conditional (and other nonatomic) moral judgments only when someone lacks a normative outlook. Once someone adopts a general normative stance by approving of a certain sort of observer, it is plausible to hold that

this is dominant in determining their normative judgments, including their conditional judgments, and that they therefore simply do not count as judging, for example, that if passive euthanasia is right then so is active euthanasia unless they believe that the observer they take to be ideal would approve of the former only if he also approved of the latter.

Of course, even if the proposed account does technically avoid the Frege-Geach problem, one might well raise independent problems with the disunity of an account which allows for the multiple realizability in this way of normative judgments. Multiple realizability is a familiar and plausible idea in the philosophy of mind, but this is typically understood in terms of multiple physical realizability of one and the same higher-order functional (or some such) property. The sort of multiple realizability on offer here is, admittedly, rather different and more radical than this. In Sellarsian terms, we here have multiple realizability within the "manifest image" itself (assuming psychological facts are within the frame of the manifest image, anyway), as opposed to multiple realizability of manifest image concepts within the framework of the subvening scientific image. Nonetheless, I think that such multiple realizability is defensible and may even have plausible analogues elsewhere. For example, there may be two ways to accept that quarks exist—the way that I (as someone with very limited knowledge of physics) do and the way that a theoretical physicist does. The former may make reference to what experts believe in a way that the latter does not. I do not want to commit myself to defending such a view as correct; that raises some very difficult issues in the philosophy of language and science, and indeed I am not really sure I even think that it is correct. However, I do want to emphasize that it is not a completely crazy or unmotivated view, and it could well turn out to be right. Such a view cannot plausibly be dismissed, in my view, simply because it posits a sort of multiple realizability within the manifest image when it comes to accepting, for example, that electrons have negative charge. In which case, the account sketched in this appendix also cannot reasonably be dismissed simply because it involves such multiple realizability. Of course, more needs to be said here, but further discussion would take us too far afield from the main points under discussion in the text. For example, one might wonder whether we could not simply adopt the "refusal to endorse" approach as a global analysis of normative judgments, instead of restricting its scope to cases of agnosticism about nonconditional first-order judgments. The short answer to this worry is that such an account would not do proper justice to the action-guiding features of nonconditional first-order normative judgments, but I lack the space here to explain that point in the detail it deserves. I hope to return to this issue in more detail in another context. Many thanks to Michael Smith for drawing me out on the issues discussed here.