

## *Non-Cognitivist Pragmatics and Stevenson's 'Do so as well!'*<sup>1</sup>

MICHAEL RIDGE  
University of Edinburgh  
Edinburgh, UK EH8 9JX

Meta-ethical non-cognitivism makes two claims — a negative one and a positive one. The negative claim is that moral utterances do not express beliefs which provide the truth-conditions for those utterances. The positive claim is that the primary function of such utterances is to express certain of the speaker's desire-like states of mind. Non-cognitivism is officially a theory about the meanings of moral words, but non-cognitivists also maintain that moral states of mind are themselves at least partially constituted by desire-like states to which moral utterances give voice.<sup>2</sup> Non-cognitivists need a plausible account of what distinguishes whims, addictions and cravings from genuinely moral judgments. For while non-cognitivists maintain that in a suitably broad sense moral judgments just are constituted by desire-like states they also insist that

---

1 Thanks to two referees for the *Canadian Journal of Philosophy* and the participants in the Work in Progress series in Philosophy at the University of Edinburgh for helpful comments on an earlier draft.

2 At least, as I am using the term 'non-cognitivism' here this is true. There is a broader sense of 'non-cognitivism' which could include views according to which moral utterances function to express neither beliefs nor desires. For example, the judgment that an action is unreasonable might express a speaker's puzzlement about how anybody could be willing to make the trade-offs involved in such an action. In this respect, the negative thesis is more essential to non-cognitivism. Thanks to an anonymous referee for emphasizing this broader notion of 'non-cognitivism.'

not any old desire constitutes a genuinely moral judgment.<sup>3</sup> Since the challenge is to demarcate what is distinctive about moral attitudes we might usefully call this the demarcation challenge.<sup>4</sup> One common strategy for meeting the demarcation challenge is to focus on desires directed at getting others to share one's own desires and emotions. This strategy has some of its earliest roots in Charles Stevenson's pioneering work.<sup>5</sup> Stevenson argued that there is a 'do so as well!' aspect to moral discourse. On Stevenson's account, in saying something is morally good a speaker not only expresses her own attitude of approval of the object of evaluation but also urges her interlocutors to share that attitude, thereby expressing a desire that they 'do so as well.' More recently, in *Ruling Passions*, Simon Blackburn emphasizes the importance of what he refers to as a 'staircase of practical and emotional ascent':

What kind of thought or feeling is involved when we have a moral reaction to some conduct or some situation? ... We should think in terms of a staircase of practical and emotional ascent. At the bottom are simple preferences, likes and dislikes. More insistent is a basic hostility to some kind of action or character or situation: a primitive aversion to it, or a disposition to be disgusted by it, or to hold it in contempt, or to be angered by it, or to avoid it. We can then ascend to reactions to such reactions. Suppose you become angry at someone's behavior. I may become angry at you for being angry, and I may express this by saying it is none of your business. Perhaps it was a private matter. At any rate, it is not a moral issue. Suppose, on the other hand, I share your anger or feel "at one" with you for so reacting. I may stop there, but I may also feel strongly disposed to encourage others to share the same anger. By then I am clearly treating the matter as one of public concern, something like a moral issue.<sup>6</sup>

Alan Gibbard explicitly refers to Stevenson's work and incorporates a 'do so as well!' element into his own theory.<sup>7</sup> It is easy to see how one might plausibly think that this would help meet the demarcation challenge. After all, most appetites and whims do not satisfy this require-

---

3 Michael Smith has recently emphasized the importance of a plausible non-cognitivist way of meeting this challenge. See M. Smith, 'Which Passions Rule?' *Philosophy and Phenomenological Research* 65 (2002) 157-63 and S. Blackburn, 'Replies,' *Philosophy and Phenomenological Research* 65 (2002) for Blackburn's reply.

4 James Lenman has usefully suggested this label for the problem; see his forthcoming reply to M. Smith, 'Evaluation, Uncertainty, and Motivation,' *Ethical Theory and Moral Practice* 5 (2002) 305-20.

5 C.L. Stevenson, *Ethics and Language* (New Haven, CT: Yale University Press 1944)

6 S. Blackburn, *Ruling Passions* (Oxford: Clarendon 1998), 9

7 A. Gibbard, *Wise Choices, Apt Feelings* (Oxford: Clarendon 1990), 173

ment. When I am hungry I typically do not desire that other people be hungry as well. When I form a fleeting desire to tap my pencil I do not typically desire that others desire that I tap my pencil or desire to tap their own pencils. Moral reactions seem different from appetites and whims. Typically, when we judge that something is wicked we not only disapprove of it; we also want others to disapprove of it in the same way.

Another apparent attraction of meeting the demarcation challenge in this way is that doing so promises to help the non-cognitivist construct a plausible theory of moral disagreement. The non-cognitivist needs an account of moral disagreement that extends to desires as well as beliefs. Dialectically, this issue is an important one because non-cognitivists often argue for their view on the grounds that it is uniquely situated to explain how moral disagreement can persist in spite of agreement on all of the relevant facts. One non-cognitivist strategy for understanding moral disagreement would be in terms of one person's desiring that someone else share her first-order desire. Once again, Stevenson's work picks up this theme. With a series of evocative examples, Stevenson argues that we do have a pre-theoretical and plausible notion of what he refers to as 'disagreement in attitude.' Moreover, Stevenson at one point suggests that two people disagree in attitude when 'at least one of them has a motive for altering or calling into question the attitude of the other' (3).

In spite of its attractions, this strategy for meeting the demarcation challenge and explaining moral disagreement in non-cognitivist terms is misguided. Stevenson's suggestion that there is a 'do so as well!' aspect to moral discourse is plausible enough, but does not help meet the demarcation challenge. For non-cognitivists should understand the desires involved in the Stevensonian 'do so as well!' not as constituting the speaker's moral judgment itself but as a pragmatic component of moral discourse. On Stevenson's account, to say something is morally good is to indicate one's approval of it and urge one's interlocutors to approve of it. In effect, Stevenson's insight is that in claiming that something is morally good a speaker expresses her preference that others approve of it too. It would be easy to move from Stevenson's insight about the function of moral language to the conclusion that to *think* something is morally good one must desire that others approve of it. However, this conclusion does not follow. The attitudes one must have to *think* x is morally good may not, in other words, be identical with the attitudes that one expresses whenever one *says* x is morally good.

For expository reasons, my discussion will be cast throughout in terms of moral value judgments rather than deontic judgments of right and wrong. My considered view is that the same points apply to deontic judgments. However, I do not want to beg any questions about the relation between the right and the good. In particular, I do not mean to assume that the good is conceptually prior in any interesting sense to the

right or vice-versa. It is a good question just how non-cognitivists should distinguish deontic judgments from value judgments. However, for present purposes I must set these difficult issues to one side and simply focus on moral value judgments however the non-cognitivist should understand them.<sup>8</sup>

One function of moral discourse is to persuade others to share one's non-cognitive attitudes, just as one function of descriptive discourse is to persuade others to share one's beliefs. In saying something is morally good, a speaker typically not only expresses her own approval of it, she also exerts pressure on her interlocutors to approve of it too. On Stevenson's account, a moral utterance expresses the speaker's non-cognitive attitude and an imperative that her interlocutors share it. More recently, Alan Gibbard has picked up on this Stevensonian idea and built a similar element into his own version of non-cognitivism. Gibbard plausibly argues that in making a moral utterance part of what one is doing is, 'demanding that the audience accept what he says, that it share the state of mind he expresses' (172). Gibbard glosses the speaker's demand that others share his attitude as an implicit claim of authority (174). I shall not argue for this Stevensonian view here, but simply grant it for the sake of argument. For my main concern is not to defend Stevenson's and Gibbard's positive claims about the function of moral language. Rather, my aim is to defend the negative thesis that we should not suppose that judging that something is morally good or bad essentially involves a pro-attitude in favor of others sharing one's approval (or disapproval) of the thing being evaluated. I note the plausibility of Stevenson's and Gibbard's claims about the function of moral language only to emphasize that those claims, which are claims about the function of moral language, are consistent with the negative thesis about acceptance conditions defended here. Again, we should not exclude *ex ante* the possibility that while *saying* that X is morally good involves the expression of a desire that others share one's attitude, *thinking* that X is morally good need not involve any such desire. There are simply two different though related questions here: (1) What states of mind does a speaker express in saying that X is morally good (or bad)? (2) What makes it true that somebody thinks that X is morally good (or bad)?

We can accommodate the insight embodied in Stevenson's 'Do so as well!' without building a 'Do so as well!' type desire (henceforth, a 'Stevensonian desire') into our account of what it is to think that something is morally good. Here it is worth reviewing some of Paul Grice's

---

8 Thanks to an anonymous referee for highlighting the importance of different moral concepts and their connections here.

rightly influential work in the philosophy of language.<sup>9</sup> Grice argues that the states of mind a speaker expresses when making an utterance may not all figure in the truth-conditions of that utterance. If, for example, you ask me whether Russell Crowe is a good philosopher and I reply by saying, 'Crowe has good handwriting,' then it is clear enough that I have (in some sense) expressed the view that Crowe is not such a good philosopher. Nonetheless, the truth of my utterance does not depend on whether Crowe is a good philosopher. For what I literally said was that Crowe has good handwriting; whether that is true is simply a different question from whether he is a good philosopher. Grice drew this distinction in terms of the explicature of one's utterance — that which provides its truth-conditions — and that which is merely implicated (the 'implicature'). More importantly for present purposes, Grice's examples also suggest a closely related but different point that Grice did not explicitly discuss. For the examples show that we must distinguish between (a) the acceptance conditions for an utterance of a given type and (b) the states of mind expressed by a token utterance of that type. Presumably, for any 'acceptance-apt' utterance — that is, for any utterance for which it is appropriate to talk of someone's accepting it — a person can be rightly said to accept such an utterance if and only if certain facts are true of her. I accept an utterance of 'grass is green' if and only if I believe that grass is green. Let us call the conditions individually necessary and jointly sufficient for someone to count as accepting an utterance (or judgment) the 'acceptance conditions' for that utterance (or judgment). Grice's examples serve to illustrate that the states of mind you express with a given utterance may include states of mind one need not have to satisfy the utterance's acceptance conditions. For while I only need to think that Crowe has good handwriting to accept a token utterance of 'Crowe has good handwriting' my token utterance in a given context might well express a good deal more than my view of his handwriting.

Implicatures are usefully divided into conversational implicatures and conventional implicatures. The former depend heavily upon a context of utterance whereas the latter are a function of some fairly well-entrenched linguistic conventions surrounding particular linguistic devices and hence are not so context dependent. Words like 'but' and 'even' serve to introduce conventional implicatures. For example, an utterance of 'Even Crowe could make that inference' implies but does not literally say that Crowe is not very good at making inferences. One

---

9 P. Grice, *Studies in the Ways of Words* (Cambridge, MA: Harvard University Press 1989)

interesting feature of conventional implicatures is that unlike conversational implicatures they seem not to be cancelable. Whereas you can say, 'Crowe has good handwriting, but I don't mean to suggest that he is not a good philosopher' without infelicity, it would be infelicitous to say, 'Even Crowe could make that inference, though I don't mean to imply that Crowe is not good at making inferences.' The general point that the acceptance conditions for an utterance are distinct from the states of mind it expresses also holds with respect to utterances that are not truth-apt or at any rate do not express beliefs<sup>10</sup> that provide the utterances' truth conditions. So even if we are non-cognitivists about moral discourse we should distinguish the acceptance conditions for moral utterances from the states of mind they express. A failure to note this distinction can easily lead one to suppose that the acceptance conditions for moral utterances of the form 'X is good' include not only the approval of X but also the approval of people approving of X. This might seem tempting just because of the plausibility of the view that moral utterances typically express such Stevensonian desires.

Moreover, the very reasons typically given for supposing there is a 'do so as well!' element to moral language provide us with equally good reason to resist the suggestion that thinking something is morally good essentially involves a Stevensonian desire. To make this point we must first mark an important distinction. We should preserve a distinction between judging that p, and judging that others should share one's judgment that p. In the case of ordinary empirical beliefs, this distinction is reasonable enough. Moreover, this point need not rest on the controversial supposition that there are no conceptual connections between thinking that p and taking oneself to be justified in thinking that p. For the point can be made even if we maintain the following conceptual connection between believing and justification: one can *come to* believe something only if at the time at which one adopts the belief one takes oneself to be justified in so believing.<sup>11</sup> This constraint should preserve

---

10 In a suitably robust Humean sense of 'beliefs'; minimalists about truth and truth-aptness argue that any indicative sentence can serve to express a belief in a very thin sense. Non-cognitivism must be characterized in terms of beliefs in some richer sense. The standard move here is to invoke the idea of states with a particular 'direction of fit' — beliefs are made 'to fit the world' whereas desires are made 'to make the world fit them.' Cashing out the direction of fit metaphor is no easy task, though, and I shall not attempt to do so here. The arguments presented in the text should go through on any of a number of ways of drawing the distinction. For useful discussion, see L. Humberstone, 'Direction of Fit,' *Mind* (1992) 59-83, and M. Smith, *The Moral Problem* (Oxford: Oxford University Press 1994).

11 I do not actually mean to endorse this principle, but simply to indicate that my

as much of a connection between believing and taking oneself to be justified as is plausible. Fortunately, there are examples that respect this constraint but still make the point needed for present purposes. For the constraint just articulated is temporally indexed to the time at which one forms a belief. The constraint therefore is compatible with beliefs sometimes having a sort of 'inertia' once formed. It is consistent to think both that an agent must take herself to be justified in adopting a belief in the first place and yet hold that the belief can then take on a life of its own, outliving the agent's taking it to be justified. Indeed, this sort of phenomenon is not really controversial. For example, religious convictions that outlive their connection with the agent's epistemic commitments provide a fairly clear case of the phenomenon. Often the agent in such cases not only comes to believe there is no justification for her belief, in some cases she actually comes to think she has an all-things-considered justification for rejecting it. Still, the belief has a kind of inertia — she cannot, it seems to her at least, shake herself of it. Furthermore, that she continues to have this belief seems borne out in her behavior (she continues to go to church, prays regularly, etc.). Nor does she seem to secretly believe that these beliefs really are justified in some way she cannot articulate — all of the available evidence indicates that she takes these beliefs to be unjustified. She not only will sincerely tell us that this is her view, we might add that she can give a detailed and persuasive account of the arguments against her view. This sort of thing does happen. One sometimes finds it in extended discussions of the philosophy of religion. Such people often distinguish faith from knowledge but typically do not distinguish faith from belief. Indeed, they often emphasize that faith consists precisely in belief without epistemic justification, characterizing themselves as 'believers' as opposed to 'non-believers.' It is, I suppose, just possible that such people are deluded about their own states of mind, and do not really believe the things they say that believe, but this seems unlikely. Not only are people generally very reliable judges of their own beliefs, such people's behavior often seems best explained in terms of their actually believing what they say they believe — their going to church, praying, obeying certain otherwise pointless rules, etc. For present purposes, the point is that such a person need not suppose she has any authority on religious matters, that her religious belief is justified or that others should share her religious judgments.

These cases are relevant because we should make space for something analogous in the moral case. There are, after all, intuitive cases of this

---

discussion is consistent with it and that any connection stronger than it is almost certainly too strong.

phenomenon in the moral domain too. I might no longer see any justification for thinking masturbation is immoral, for example, but find myself unable to shake loose of the judgment that it is morally bad. Early childhood conditioning sometimes leaves people in such a state. So it seems implausible to suppose that thinking something is morally good (or morally bad) must be partially constituted by a desire that others approve of it. Why would anyone think otherwise? The justification most often suggested is that in judging that something is morally good one also implicitly sees oneself as an authority for that judgment, or at least supposes it is somehow justified and worthy of being adopted by others. The cases canvassed above in the case of ordinary beliefs should cast doubt on all of these theses — one can think that *p* without thinking that one's so thinking has authority, is worthy of being held or is justified. Crucially, this is perfectly compatible with accepting the Stevensonian idea that there is a 'Do so as well!' element to moral language. Generally speaking, if I say, 'Abortion is morally bad,' then I thereby express not only my disapproval of abortion but my desire that you share that disapproval. We rely on moral discourse to influence one another in this way, making various implicit conversational demands. This, however, is consistent with allowing that there is a logical gap between, e.g., judging that abortion is bad and judging that others ought to share that judgment. Allan Gibbard claims that in making a moral utterance part of what a speaker is doing is, 'demanding that the audience accept what he says, that it share the state of mind he expresses' (172). Once we distinguish this from Gibbard's more controversial claim that in making assertions we are claiming to be an authority of some kind, this idea is plausible, both with respect to the expression of ordinary descriptive beliefs and the expression of moral judgments. We can agree with Gibbard but hold that the Stevensonian desire expressed by such a demand is not necessary for someone to count as thinking the object of evaluation is morally good.

In fact, the very reason Gibbard gives for including a 'Do so as well!' element in his analysis of the function of moral utterances provides an equally good reason not to build such an element into our theory of their acceptance conditions. For Gibbard's reason for building a 'do so as well!' element into the pragmatics is that this reflects the fact (simply assuming it is a fact) that one expresses a supposition of authority whenever one issues a moral utterance. The preceding discussion shows why this supposition of authority should not be built into the utterance's acceptance conditions. As we have seen, one can think that something is morally good (or bad) without seeing oneself as any kind of authority at all. Indeed, one can think something is morally good and not even think one's view is justified. So insofar as Gibbard is right that the 'do so as well' element does capture a supposition of authority we ought not to

include any such element in the acceptance conditions for moral utterances. It is useful to return to the non-moral case. Here is Gibbard:

Conversation is full of implicit demands and pressures. Suppose I confidently expound astrology, and you give no credence. The result will be discomfort: in effect I demand that what I say be accepted, and you will not accede. (172)

The reason I only 'in effect' make an 'implicit' demand that what I say be accepted is that this is merely a pragmatic element of what I say. That this is so is clear from the fact that my demand plays no part in the acceptance-conditions of my utterance. Intuitively, someone can accept that astrology works but not give a damn whether anybody else believes in it. That a tendency to make such demands is not part of the acceptance conditions for such utterances can also be seen from the way in which such demands can without logical inconsistency be canceled. 'Astrology is an accurate way of making predictions, but you shouldn't take my word for it, I'm irrationally biased on this issue,' is a logically consistent though odd thing to say. This sentence is vaguely similar to the ones Moore discussed and that are now referred to as instances of 'Moore's paradox' — sentence of the form, 'p, but I don't believe it' or 'I believe that p but not-p.'<sup>12</sup> In both cases there is something like a pragmatic tension that falls short of logical contradiction. Whereas, 'Astrology is an accurate way of making predictions, but its not' is logically contradictory precisely because it tries to cancel the expression of a view that is essential to the acceptance conditions of the utterance. Similarly, in making a moral utterance one typically makes an implicit demand that others share one's attitude, but this is best understood as part of the pragmatics of your utterance rather than its acceptance conditions. I might judge that abortion is morally bad but not think others should share that judgment and therefore keep my mouth shut about it. Or if I say anything, I might say, 'Abortion is wicked, but don't believe it because its even worse to be the sort of person who has views on such matters.' After all, considerations of privacy and anti-moralistic sentiments do sometimes lead people to disapprove of having a moral view on certain kinds of personal questions at all even though they may inadvertently have come to hold such a view themselves. Admittedly, saying that abortion is wicked and in the same breath telling your interlocutors not to believe it is wicked would be very odd, but it is not logically incoherent. Intuitively, the last clause of the utterance cancels any implication that one wants others to share one's judgment about

---

12 G.E. Moore, *Principia Ethica* (New York: Cambridge University Press 1903

abortion. It does raise the question of why the speaker said anything about abortion in the first place, but a suitable story can make sense of this (perhaps he was forced at gunpoint to divulge his considered opinion on the matter).

The expression of Stevensonian desires is more analogous to conventional implicature than conversational implicature but differs from standard examples of conventional implicature in interesting ways. The 'do so as well!' aspect of moral language is only analogous to these phenomena, strictly speaking, because the standard account of conversational and conventional implicatures is as expressing beliefs rather than desires. It is more analogous to conventional implicature because it does not depend much on the peculiarities of a particular context of utterance but is rather a function of the conventional meanings of moral terms when used in certain ways. This makes the 'do so as well!' aspect of moral language an interesting case for the philosophy of language more generally. For it is usually supposed that conventional implicatures differ from conversational implicatures both in that they are not cancelable and in that they are less context-dependent. However, the 'do so as well!' aspect of moral language is both conventional and cancelable. In this way, moral language seems to occupy a kind of Janus-faced relationship to standard ways of thinking about the differences between conventional and conversational implicatures. Like conventional implicatures, the 'do so as well!' element is a function of the conventional meaning of some linguistic device in the utterance in a way that is not highly context-sensitive, but like conversational implicatures the 'do so as well!' element is cancelable. This feature of the 'do so as well!' aspect of moral language should make it less controversial than standard cases of conventional implicature. For part of what makes the standard cases controversial is the idea that though you cannot cancel the implicature it is still not part of what you literally said. This is controversial precisely because a defining mark of one's having literally said something simply is that one cannot cancel the implication without linguistic infelicity.<sup>13</sup> A nice feature of the sort of conventional implicature found in the moral case is that it avoids this problem insofar as it plausibly thought of as cancelable.<sup>14</sup> In any event, the main conclusion being defended here is that we can preserve the insights of Gibbard and Stevenson about moral

---

13 For a detailed case against the very idea of conventional implicature see K. Bach, 'The Myth of Conventional Implicature,' *Linguistics and Philosophy* 22 (1999) 327-66.

14 Thanks to Alexander Bird for useful discussion here.

practice by building the 'Do so as well!' into the pragmatics of moral utterances rather than their acceptance conditions.<sup>15</sup>

## Conclusion

Grice's discussion of implicatures shows that the acceptance conditions for a given utterance may not include all the states of mind such an utterance serves to express. Although Grice focused on truth-apt discourse, the distinction holds with respect to discourse that is arguably not truth-apt as well. In particular, we can accept the non-cognitivist thesis that moral language is not truth-apt and still distinguish the acceptance conditions for moral utterances from the states of mind expressed by such judgments. In the moral case this is just to say we should distinguish, e.g., what states of mind one voices in saying that something is morally good and what states of mind one must have in order to count as thinking it is morally good. A failure to heed this distinction can lead one to suppose that the plausible Stevensonian thesis that there is a 'do so as well!' element to moral discourse entails the implausible consequence that to think that something is morally good (or bad) one must prefer that others approve (or disapprove) of it. Fortunately, this consequence follows only if we fail to internalize Grice's insight and its applicability to forms of discourse that are not truth-apt. Once we see that this insight also applies to moral discourse even if that discourse is not truth-apt we can reject the thesis that moral utterances have Stevensonian desires built into their acceptance conditions. Non-cognitivists can happily accept the Stevensonian thesis that there is a 'do so as well!' element to moral language without building a desire that others 'do so as well' into the acceptance conditions of moral utterances. Since such a theory of the acceptance conditions is independently implausible, this is good news for non-cognitivists. However, this also means that the non-cognitivist cannot explain what is distinctive about the attitudes that constitute moral judgments as such in terms of Stevensonian desires. Nor, given this line of argument, should the non-cognitivist adopt the otherwise tempting strategy of explaining moral disagreement in terms of conflicting Stevensonian desires. For we can have moral disagreements that remain unvoiced. Indeed, if I think that something is morally bad and someone else thinks it is not morally bad then we disagree even if we have never heard of one another, much less

---

15 Gibbard might agree, as he notes that conversational demands are 'revocable' (172) in that a speaker can rescind those demands.

come to want to change the other's attitudes. So non-cognitivists cannot fully explain moral disagreement in terms of speech-acts and the Stevensonian desires they typically express. This is not necessarily such bad news for non-cognitivists. In my view, non-cognitivists can more plausibly meet the demarcation challenge and explain moral disagreement without invoking Stevensonian desires. That, however, is another story.

*Received: February 2003*

*Revised: May 2003*