

Expressivism and Irrationality

Mark van Roojen

Noncognitive analyses of evaluative discourse characterize moral discourse as primarily functioning to express attitudes that are not, strictly speaking, representational in the way that ordinary beliefs are representational. But, since expressivists must explain our practices of making evaluative judgments as we do, they owe us an explanation of the logical relations between these evaluative judgments and other judgments. For it is part of our ordinary evaluative practices to make inferences based upon and leading to evaluative judgments. The most thorny problem for this project has been to explain the logical relations between evaluative judgments and other judgments best expressed using evaluative terms in unasserted contexts, such as clauses embedded in conditionals. Because one may use evaluative terms in such unasserted contexts without expressing the attitudes usually associated with them in asserted contexts, it becomes hard to explain why there should be logical relations between the judgments expressed. The noncognitivist who has given this problem the most sustained attention over the years is Simon Blackburn,¹ and recently he has been joined by Allan Gibbard.²

An earlier version of this paper was read at the 1995 Pacific Division meetings of the American Philosophical Association. Many thanks are due my commentator on that occasion, Frances Howard-Snyder, for her useful comments, and to the moderator, Albert Mosley. I'd also like to thank Robert Audi, Ed Becker, Albert Casullo, Gil Harman, Joe Mendola, David Pitt, Lyle Zynda, and two very kind and helpful anonymous readers for the *Philosophical Review* for useful comments on earlier drafts of this paper. Many thanks are also due Jamie Dreier for much helpful discussion of Geach's problem and for comments on an earlier draft, and Harry Ide, both for useful comments on earlier drafts and for many hours of helpful discussion as I was writing them.

¹The relevant texts are, "Moral Realism," in *Morality and Moral Reasoning*, ed. Casey (London: Methuen, 1971), 101–24; *Spreading The Word* (Oxford: Oxford University Press, 1984); and "Attitudes and Contents," *Ethics* 98 (1988): 501–17. The first and third of these texts were republished as essays 6 and 10 of his *Essays in Quasi-Realism* (Oxford: Oxford University Press, 1993). Blackburn continues to endorse the general approach, most recently in "Gibbard on Normative Logic," *Philosophy and Phenomenological Research* 72 (1992): 947–52.

²Gibbard's treatments are found in his *Wise Choices, Apt Feelings* (Cam-

In a series of papers and in his *Spreading the Word*, Blackburn offers us a series of accounts, each of which builds on the last. They all share two features: First, they try to explain the logical behavior of the contents of such judgments by invoking norms³ governing the co-acceptance of attitudes. Second, they interpret conditionals and the like as involving higher-order attitudes towards the attitudes that would be expressed by their constituents were those constituents to occur in asserted contexts.

Previous critics have argued that the proposals involve a conflation of the notion of logical consistency with some other notion or notions that involve pragmatic but not logical inconsistency.⁴ Blackburn's response has in effect been to concede that he is using "inconsistency" to cover something other than what his critics view as logical inconsistency, but to insist that this usage is not a conflation and generates no philosophical difficulty. In this paper, I offer two sets of criticisms meant to bolster the original criticism: The first, more particular, criticism is that the norms governing attitudes that Blackburn must invoke to explain the "inconsistency" in denying the conclusions of valid evaluative arguments with premises one accepts will be too strong. If these norms are strong enough to rule out denying the conclusions of valid arguments while also accepting the premises, they will also rule out concurrently accepting the premises of seemingly good arguments, and the premises of these good arguments must then also be judged "inconsistent" on the interpretation offered. I argue that this is a recurring problem for the expressivist, which derives from the necessity of expanding the notion of contradiction to apply to attitudes that can be neither true nor false. I try to bolster that claim by showing that Gibbard's account is subject to a similar objection. The second, more general, criticism is that an expanded notion of inconsistency applicable to the contents of attitudes cannot easily

bridge: Harvard University Press, 1990), especially chapter 5, and in his "Reply to Blackburn, Carson, Hill and Railton," *Philosophy and Phenomenological Research* 72 (1992): 969–80.

³By 'norms' I mean rational standards, principles, or regulative ideals that prescribe how a person should operate within the domain in question.

⁴In addition to Schueler, whose paper (cited in note 11) occasioned Blackburn's response, prominent critics include Bob Hale, "The Compleat Projectivist," *Philosophical Quarterly* 36 (1986): 65–84; M. H. Brighouse, "Blackburn's Projectivism—An Objection," *Philosophical Studies* (1990): 225–33, and Nick Zangwill, "Moral Modus Ponens," *Ratio* (1992): 177–93.

be defined in terms of norms governing the rationality or irrationality of accepting certain attitudes. For the rational connections among attitudes do not neatly mirror the logical connections between contents. These conclusions will reinforce the conclusion urged by Blackburn's earlier critics, that such noncognitivist analyses cannot extend what is distinctive about logical implication from representational to merely expressive discourse.

1. The Frege-Geach Problem

The problem for the noncognitivist of accounting for the logical connections between evaluative judgments and other judgments was forcefully raised by Peter Geach in the early 1960s.⁵ Moral claims that purport to be assertions function in logical arguments just as nonmoral assertions do. Thus, the following argument is obviously valid, despite containing moral premises:

- (P1) If tormenting the cat is bad, getting your little brother to do it is bad.
 - (P2) Tormenting the cat is bad.
- Ergo*, getting your little brother to torment the cat is bad.⁶

The problem for the noncognitivist is that the second premise is supposed to be functioning to express a negative attitude, rather than to assert a proposition. But the first premise might well be accepted by someone who has no negative attitudes towards tormenting cats. So whatever the meaning of P1, the contribution of its antecedent to whatever P1 expresses isn't to express that attitude. This creates a puzzle: Our standard accounts of the validity of modus ponens require that P2 express what the antecedent of P1 contributes to what P1 expresses. On the noncognitivist account, the words of the antecedent of P1 and the words of the second premise, P2, are functioning as homonyms. But modus po-

⁵Geach first raises the point in a footnote in "Imperative and Deontic Logic," *Analysis* 18 (1958): 49–56, at 54. It is raised again by John Searle in "Meaning and Speech Acts," *Philosophical Review* 71 (1962): 423–32. Geach develops the point in greater detail in "Assertion," *Philosophical Review* 74 (1965): 449–65.

⁶The example is from Geach, "Assertion," 463.

nens would not validly apply to such premises if the words functioned homonymously. Thus, Geach accuses noncognitivists of ignoring “The Frege Point” that sentences or propositions do not change their meanings depending upon whether they are asserted or merely used in other ways, as when embedded in conditionals.⁷ The noncognitivist owes us an account of why the normal inferences are valid, if the noncognitivist’s account is accepted, given that the usual explanation presupposing the Frege point is not available.

2. Blackburn’s Three Proposed Solutions

Each of Blackburn’s answers to the challenge has involved interpreting statements using unasserted moral terminology as expressing a species of higher-order attitude regarding evaluative judgments. The proposals overlap, but each successive version appears to involve some modification to the previous one. It will be useful to present my criticisms of each proposal in turn. Hence, I should summarize each of the proposals before proceeding to those criticisms.

2.1 *Solution One*

Blackburn’s first proposal is to be found in a paper entitled “Moral Realism”: Conditionals refer to or make claims about the attitudes that seem to be parts of the larger whole.⁸ They are “propositional reflections” of claims about attitudes, that is, assertions that seem to make factual claims about states of affairs, while actually making factual claims about attitudes toward those states of affairs. “‘If courage is intrinsically good, then organized games should be part of school curricula,’ claims that an attitude of approval of courage involves an attitude of approval of such games.”⁹ Similarly, then,

⁷While embedding is the issue here, Geach himself thought this was an example of a general kind of problem for theories of many sorts: “A thought may have just the same content whether you assent to its truth or not; a proposition may occur in discourse now asserted, now unasserted, and yet be recognizably the same proposition” (ibid., 449).

⁸“Moral Realism,” in *Morality and Moral Reasoning*, ed. Casey (London: Methuen, 1971), 101–24, at 119; this paper was reprinted in Blackburn, *Essays in Quasi-Realism*.

⁹Ibid., 120.

the conditional above about tormenting cats claims that disapproval towards tormenting cats involves disapproval of getting one's little brother to torment them. Given this analysis, Blackburn hopes to explain the validity of inferences licensed by accepting the conditional. In circumstances where one expresses an attitude to a thing, and then also asserts that such an attitude involves a further attitude (or belief), one would be logically inconsistent if one did not then hold that further attitude (or belief).¹⁰

Blackburn's explanation for the validity of the inference in Geach's example is that disapproval of tormenting cats, the belief that such disapproval itself involves disapproval of getting one's little brother to do it, and a failure to disapprove of getting one's little brother to do it are logically inconsistent. There is reason to question Blackburn's use of "logically inconsistent" here. For, normally, we think of logical inconsistency as involving belief in propositions that contradict one another, and there is no such inconsistency here.¹¹ But Blackburn's idea seems to be that one is involved in holding attitudes that cannot coherently fit together and that, hence, are irrational to maintain. Thus, he likens the inconsistency to that of someone who approves of giving a tenth of a pound to each of ten people, but does not approve of giving them each two shillings. I will soon come round to arguing that this kind of inconsistency cannot do the work required, and hence to reinforce the complaint that this is not logical inconsistency. But first I should present Blackburn's next two accounts.

2.2 *Solution Two*

Blackburn's account in *Spreading the Word* supplements the earlier account of evaluative logic in two ways. First, it offers an expressivist formalization for translating evaluative utterances, which clearly displays the noncognitive nature of the analysis Blackburn wishes to offer:

Imagine a language unlike English in containing no evaluative predicates. It wears the expressive nature of value judgements on its sleeve.

¹⁰Ibid., 121.

¹¹This point is made by G. F. Schueler, in "Modus Ponens and Moral Realism," *Ethics* 98 (1988): 492–500, at 496.

... It might contain a 'hooray!' operator and a 'boo!' operator (H!, B!) which attach to descriptions of things to result in expressions of attitude. H!(the playing of Tottenham Hotspur) would express the attitude towards the playing. B!(lying) would express the contrary attitude towards lying, and so on. ...¹²

Second, *Spreading the Word* refines Blackburn's proposal regarding higher-order attitudes, so that the attitude expressed towards other attitudes by a conditional with evaluative components is now best characterized not as a belief that one attitude involves another, but rather as an attitude of approval or disapproval towards conjunctions of attitudes. 'If lying is wrong, telling your little brother to lie is wrong' (when sincerely uttered) expresses approval of making disapproval of getting one's brother to lie "follow upon" disapproval of lying.¹³

The logical validity of the standard inferences is now explained via the attitudes one can hold without "clash," or consistent with "the practical purposes for which we evaluate things."

[A]nyone holding this pair [the above, plus the attitude expressed by 'lying is wrong'] must hold the consequential disapproval: he is committed to disapproving of getting little brother to lie, for if he does not his attitudes clash. He has a fractured sensibility which cannot itself be an object of approval. The 'cannot' here follows not (as a realist explanation would have it) because such a sensibility must be out of line with the moral facts it is trying to describe, but because such a sensibility cannot fulfill the practical purposes for which we evaluate things.¹⁴

Logical entailments involving moral judgments are explained as follows: Given the purposes served by moral judgments, a constellation of attitudes that includes the attitudes expressed by the conditional and by the seemingly assertive premises but not those expressed by the conclusion is irrational, because it goes against those purposes.

¹²*Spreading the Word*, 193.

¹³See *Spreading the Word*, 194–95. This idea is already suggested by Blackburn in the earlier article, where he says that to make the claim that one attitude involves another is to make a moral claim ("Moral Realism," 121–22).

¹⁴*Spreading The Word*, 195.

2.3 *Solution Three*

Blackburn's third proposal comes in a response to criticisms of his first two proposals leveled by G. F. Schueler.¹⁵ Schueler complained that Blackburn was unable to show any real inconsistency between accepting the premises and denying the conclusions of the sorts of arguments he was trying to explain. In response, Blackburn offers two related rebuttals. First, he argues that it is not clear that there is just one sort of thing that counts as a genuine inconsistency. Second, he argues that there is a notion of inconsistency of attitude available to an expressivist that does underwrite the claim of inconsistency, but admits that his earlier proposal needs some modification. The essential modification is that Blackburn now treats certain attitudes as consisting in a "commitment" to holding certain attitudes or conjunctions of attitudes, under certain circumstances. But these attitudes no longer necessarily express higher-order moral evaluations, though they may be explained by the presence of such a moral outlook. For example, accepting a disjunction is being committed to accepting the attitude that would be expressed by one of the disjuncts, should the attitude expressed by the other become "untenable." And one may hold that attitude because one has a certain moral perspective. But that perspective is not part of "the content of the conditional."¹⁶

The commitments expressed, then, are commitments to make certain "inferences" should one's attitudes undergo certain changes. Which inferences the proposal licenses depends on the commitments expressed by various evaluative judgments. Blackburn explicates these by defining a "deontic logic" that evaluative judgments must obey, patterning that logic after an earlier proposal of Hintikka's. Blackburn's actual proposal is underspecified, hence ambiguous and hard to explicate.¹⁷ But the feature that

¹⁵Blackburn, "Attitudes and Contents," *Ethics* 98 (1988): 501–17, responding to Schueler's "Modus Ponens and Moral Realism."

¹⁶"Attitudes and Contents," 512. Actually, it isn't clear that it was part of the *content* of the conditional in the second proposal either, given that the judgment expressed that attitude rather than stated that the speaker had it. After all, that is the expressivist proposal regarding moral utterances generally. As a result, it isn't clear how the new proposal distinguishes between expressing the moral attitude and merely expressing the commitment one has because one has the moral attitude or outlook.

¹⁷For example, Blackburn's fourth rule for generating approximations

seems to do the main work is that his proposed rules make it “inconsistent” to express attitudes of approval towards states of affairs that could not coexist (or to be tolerant towards states of affairs that could only exist in the absence of one towards which one expressed favorable attitudes).¹⁸

3. The Particular Objections

I have already noted that Blackburn’s prior critics, Schueler in particular, have criticized him for conflating other notions with that of logical consistency and inconsistency. Blackburn’s response has been to argue that the notions of consistency and inconsistency are heterogeneous, admitting of different kinds of incoherence, so that noncognitive attitudes can also admit of consistency and inconsistency, thereby underwriting logical implications among essentially noncognitive expressions. So, in effect, Blackburn accepts the point that his “inconsistency” is not the same as inconsistency between propositions, but claims that there is no harm in that. The criticisms that follow should undermine this last claim.

to an ideal relative to a set of moral judgments applies only when a set of sentences is itself the “next approximation” to some ideal. But we need to know the content of the fourth rule before we can settle the issue of which sets of sentences are such approximations (see “Attitudes and Contents,” 513–14). Bob Hale does a heroic job, both of explaining the problems with the proposal as presented and of suggesting ways of fixing them, in “Can There Be A Logic of Attitudes?” in *Reality, Representation and Projection*, ed. Haldane and Wright (Oxford: Oxford University Press, 1993), 337–363. He also presses a line of objection to the proposal as amended.

¹⁸At least that seems to be the intended effect of the rules he defines. The actual proposal involves applying the rules to generate sets of sentences that are supposed to be closer and closer approximations to worlds in which obligations are satisfied and tolerable actions brought about. If in any of the descriptions of those worlds a well-formed formula and its negation is propounded, then that set is unsatisfiable. Judgments are inconsistent if every route to a final ideal via application of the logical rules contains a well-formed formula and its negation. Since ‘H!A’ is a well-formed formula if ‘A’ is, unsatisfiability is generated even where an attitude and its negation are expressed. But since the negation of ‘H!A’ is interpreted as T!not(A) (expressing the idea that not(A) is tolerable) the “inconsistency” is generated precisely because A and not(A) are not co-satisfiable. Actually, given the rules as Blackburn actually defines them, things do not always work out like this. ‘H!A,’ and ‘T!A’ (the translation of ‘not(H!A)’) can be made “satisfiable” if considered together with a disjunctive premise. Clearly, Blackburn intended otherwise.

3.1 A Problem for the First Analysis

Consider this argument:

1. If I don't disapprove of Y, then X is wrong. (Premise)
2. If X is wrong, then Y is wrong. (Premise)
3. I don't disapprove of Y. (Premise)
4. Therefore, X is wrong. (From 1 and 3)
5. Therefore, Y is wrong. (From 2 and 4)

The argument is valid, and its premises are consistent. Its premises are inconsistent with the negation of its conclusion (5): Y is not wrong.

Now look at the translation suggested by the noncognitivist and by Blackburn's first proposal that conditionals are disguised claims that one attitude involves another.¹⁹

- 1'. If I don't disapprove of Y, then B!(X).
- 2'. The disapproval of X involves the disapproval of Y.
- 3'. I don't disapprove of Y.
- 4'. B!(X).
- 5'. B!(Y)

Blackburn's explanation of the validity of the argument as translated will be that (2'), (4'), and the denial of (5') are inconsistent in some sense.

But this leads to a problem. Surely (3') and (5') are just as "inconsistent" as (2'), (4'), and not(5') are. That is, they too seem to express clashing attitudes. If (3') and (5') are inconsistent, then the premises (1'), (2'), and (3') must be inconsistent, since they together imply both (3') and (5'). Blackburn has given us an explanation of the validity of the argument only if he has given us an argument that the premises are inconsistent. But the premises of the original argument seem consistent, and support a valid argument. Thus, Blackburn's proposal has failed to meet one of the

¹⁹Since the operators introduced in *Spreading The Word* merely formalize what is already part of the noncognitivist proposal, we can make use of that formalization to capture the idea behind even the earliest proposed analysis.

requirements of an adequate analysis. If the premises of an argument are consistent, the premises of the Blackburn translation of that same argument should be consistent.

3.2 A Similar Problem for the Second Analysis

Consider now this argument:²⁰

- P1. It would be wrong for me to believe ill of my friends. (Premise)
- P2. My parents, father and mother alike, are my friends. (Premise)
- P3. It would be believing ill of a friend to believe that he would be duplicitous with another of one's friends. (Premise)
- P4. If the coded valentine is not a joke, my father is being unfaithful to my mother, and hence duplicitous. (Premise)
- P5. The coded valentine is not a joke. (Premise)
- C1. It is wrong for me to believe that my father is unfaithful to my mother. (From P1, P2, and P3)
- C2. My father is unfaithful to my mother. (From P4 and P5)

The argument for the two conclusions is valid, and its premises are consistent. But the logic of higher-order attitudes that Blackburn uses to explain the validity of evaluative arguments must rule the two conclusions inconsistent with one another. Hence, it must regard the premises as also inconsistent. For the first conclusion might be formalized as B!(My believing that my father is unfaithful to my mother),²¹ and the second expresses just that belief. Black-

²⁰The argument was suggested by a letter to advice columnist Abigail Van Buren, author of "Dear Abby."

²¹At least that seems a perfectly good way to formalize it, given that Blackburn indicates that 'B!' and 'H!' are to "attach to descriptions of things to result in expressions of attitude" (*Spreading The Word*, 193). Immediately following, however, he introduces a special notation for talking about attitudes or beliefs by "putting its expression inside bars: /H! (X)/ refers to the approval of (X)" (*ibid.*, 194). So perhaps the sentence should be formalized as 'B! /Believing that (my father is unfaithful to my mother)/'. Since attitudes are a kind of thing, it isn't clear which way is the preferred translation. In any case, it matters little for the point being made, since either way the attitude in question, as interpreted by Black-

burn had hoped to draw the necessary logical connections between beliefs from the idea that it is inconsistent to have a higher-order negative attitude towards an attitude while also having that attitude. But if that is irrational, and if the logical connections track these rational connections as the proposal suggests, then these two attitudes are rationally inconsistent, and hence logically inconsistent. For the premises in question violate the prohibition on holding a higher-order attitude of disapproval toward some set of lower-order attitudes together with those same lower-order attitudes. And he relied on that prohibition to explain how expressive but not representational attitudes could generate a logical inconsistency.²²

But, of course, they are not logically inconsistent. If a person who thinks it is wrong to believe something, and yet believes it anyway, commits an error, it is not an error of logic. Since the explanation used to explain logical inconsistency is equally applicable here, that explanation commits Blackburn to finding logical inconsistency where there is none.

3.3 Problems for the Third Analysis

It should be clear that the third proposal explains the logical relations entered into by moral terms by stipulating that they are commitments to make inferences in accord with that logic. And this leads to several problems.

First of all, the logic defined incorporates several substantive moral assumptions that, while necessary to explain the validity of standard inferences, generate contradictions where there are none. $H!A$ and $H!(\text{not}(A))$ are unsatisfiable. Yet, it is obviously not contradictory to think that both the truth of a proposition and the truth of its negation are good, or that a state of affairs and its

burn, involves a higher-order disapproval of a lower-order attitude which one oneself holds. And that, he claims, is logically inconsistent.

²²Blackburn might wish to resist this conclusion, but I don't see how he can, given that the argument relies on just the aspect of these attitudes that he uses to count the noncontroversially inconsistent attitudes as inconsistent, despite their nonrepresentational status. It is worth noting that nothing turns on the first-order attitude in question being a belief rather than a pro or con attitude. For we might have constructed a parallel argument for the claim that my father had done something deeply wrong, and for the judgment that it is wrong to believe that one's friends have done something deeply wrong.

alternatives can be good. Perhaps H! is meant to capture only judgments of what is obligatory and not judgments of what is good,²³ and it is admittedly puzzling how one could have an obligation both to bring about a state of affairs and also to not bring it about, but there is no contradiction in thinking that we have such conflicting obligations. If someone is mistaken in believing that there are such conflicts, the mistake is not one of logic.²⁴

Once again, we could exploit this feature of Blackburn's proposal to generate a valid argument with consistent premises that his account would class as inconsistent from the start. But I trust the general point is clear without such an illustration. It thus turns out that the third proposal suffers from the same flaw as the first two. To explain implications between expressions of attitude, Blackburn must invoke a notion of inconsistency strong enough to rule noncognitive attitudes inconsistent, hence extending the notion of inconsistency beyond cases where the truth of one content rules out the truth of another. But the stronger notion invoked rules logically consistent claims inconsistent, precisely because it is stronger.

4. Gibbard's Norm Expressivism

We have noted a recurring problem, that the resources used to explain inconsistency where all agree there is inconsistency generate inconsistency where there seems to be none. And we might wonder whether the problem stems from the nature of the project or is, instead, merely peculiar to Blackburn. I wish to argue that it stems from the nature of the project. So it will strengthen my case if the other leading expressivist theory, that of Allan Gibbard, has similar problems.²⁵

²³In any case, the expressivist owes us an account of the functioning both of expressions regarding rightness and of expressions regarding goodness, a requirement that is generally ignored. Since on many views it can be good that someone did the wrong thing, the task would involve some complication.

²⁴In fact one well-known argument against "realism" and in favor of some kind of noncognitivism starts with the idea that such conflicts are possible. See Bernard Williams, "Consistency and Realism," in *Problems of the Self* (Cambridge: Cambridge University Press, 1973), 187–206.

²⁵Gibbard's analysis is found in his *Wise Choices*.

Gibbard's analysis concerns judgments of rationality first, and judgments of moral rightness and wrongness are analyzed as a species of judgment about rationality. For simplicity's sake, my discussion here will focus on the former. Judgments of rationality are to be analyzed, not as stating facts about the world, but as expressing attitudes towards sets of norms. Norms here can, I think, be thought of as much like a prescription commanding or prohibiting certain naturalistically described kinds of actions. Gibbard then refines the initial analysis until judgments of rationality are equated with expressions of acceptance, and demand for acceptance by others, of both a lower-order norm requiring the action/attitude that is judged rational and a higher-order norm requiring the acceptance of the lower-order norm.²⁶

Gibbard's treatment of Geach's problem involves the postulation of a special sort of content for judgments involving rationality. The contents of normative judgments are sets of ordered pairs of possible worlds and systems of norms. Judgments of rationality express the speaker's "ruling out certain combinations of factual possibilities with norms."²⁷ Judging that a certain action is rationally required rules out combinations of possible worlds in which the object in question has certain naturalistic features with systems of norms that do not recommend objects with those features. These contents are then used by Gibbard to define a logic for the expressions that relatively nicely tracks the logical roles they actually do seem to have.²⁸ A disjunction has the content of ruling out all

²⁶*Wise Choices*, 172–73. Gibbard goes on to discuss the conditions under which one can make conversational demands of others to accept the norms towards which one is oneself expressing acceptance. It isn't entirely clear whether the analysis further requires one to believe oneself in such conditions before one can sincerely express the requisite attitude. If so, the analysis is more complicated yet. As I read the argument, the ensuing section is merely an interesting discussion of when it might be sensible to make, or to heed, such demands; the actual analysis is complete at page 173.

²⁷*Ibid.*, 102.

²⁸As Gibbard formulates the proposal, the norms apply to naturalistically described kinds of objects, so that they will never treat two naturalistically identical objects differently. Failure to respect the supervenience of the normative on the non-normative will, therefore, be treated as a logical error. Since Gibbard could allow norms applying to particulars, this is but an incidental feature of the approach, and does not show that any similar view would illegitimately treat such errors as logical errors when they are not.

the norm-world pairs that are ruled out by both disjuncts, and a conjunction rules out the norm-world pairs ruled out by either conjunct. An argument is valid if the intersection of the sets of pairs representing the contents of the premises (that is, those not ruled out) is a subset of the sets representing the conclusion.

Despite the general adequacy of the world-norm pair apparatus to model the inferences desired, Gibbard's proposal can be used in two ways to generate logical inconsistency where there is none. The first is relatively minor and trades on a feature of the formal system detailing the logic. The second is more important and arises out of the claim that certain kinds of higher-order attitudes entail logical commitments of the kind specified.

To see the first problem we need some more detail about the logical apparatus Gibbard proposes. Gibbard means to restrict the kinds of norms that can form a system, so that a forbidden course of action cannot be required. Thus, the systems of norms that can be elements in the norm-world pairs must be "consistent" in this sense.²⁹ But this consistency requirement raises a familiar problem for the logic. It is a controversial ethical claim whether an agent might not be faced with a choice requiring her to avoid an option that can be avoided only by doing another that is also forbidden (witness the literature regarding moral dilemmas). In fact, Gibbard himself recognizes this when he remarks that a requirement that an action cannot be optional if its alternatives are forbidden would be controversial.³⁰ Those who question whether this must be so are

²⁹As is apparent when he states:

We can characterize a system *N* of norms by a family of basic predicates '*N*-forbidden', '*N*-optional', and '*N*-required'. Here '*N*-forbidden' simply means "forbidden by system of norms *N*", and likewise its siblings. Other predicates can be constructed from these basic ones. . . .

[W]hen a system of norms *N* applies in a definite way to an alternative, that results in an alternative's having exactly one of the three basic properties, being *N*-forbidden, *N*-optional, or *N*-required. We can call a system *complete* if these predicates trichotomize the possibilities. . . . (So long as *N* is consistent, nothing will be more than one of these things.) (*Wise Choices*, 87)

³⁰*Ibid.*, 88. In a footnote he cites van Fraassen and Marcus (no logical slouches) as opponents of the requirement under consideration. Gibbard appears not to have noticed that his own proposal is not really any different, in that it rules out a person's being both required to do and forbidden from doing the very same thing. That seems no different from requiring that there always be an outcome that is permitted.

not thereby displaying a lack of logical facility. They are merely engaging in a dispute over the content of the correct set of norms. Insofar as Gibbard's proposal would assimilate such disagreement to a logical failing by those on one side of the dispute, it is problematic.

I don't want to make too much of this, since it may well be possible to place other restrictions on the permissible normative systems that would yield the correct logical relations between contents, without the suspect requirement.³¹ The more important problem for Gibbard's account comes when we notice that there are other ways besides assertions involving rationality to express one's acceptance of higher-order norms recommending lower-order norms together with acceptance of those lower-order norms.

If Gibbard's analysis is correct, any such expression should function logically just as judgments of rationality do. But do they? Suppose I say, "I hereby express my acceptance of the requirement that I remain silent, as well as a norm requiring the acceptance of that norm, and furthermore prescribe that others do likewise" (call this statement "A"). This statement expresses the acceptance of the same set of norms that a judgment of rationality does, on Gibbard's analysis. And that it functions to express that attitude is not accidentally related to the sentence's sincere use. If we can explain the meaning of judgments of rationality by explaining that they perform a certain expressive function when sincerely uttered, we have just as much reason to count that expressive function as part of the meaning of judgments like this. The statement in question also does more than express an attitude: It states that the speaker is expressing the attitude *by the very utterance of the statement itself*. So in addition to its expressive content, the judgment has the perfectly truth-assessable propositional content of claiming that it expresses those commitments.

³¹Why am I not equally charitable towards the analogous problem in Blackburn's third proposal? Because given that Blackburn wants to derive the inconsistency of the attitudes in question from the purported irrationality of having desires that would be self-defeating if all realized, the commitment is essential to his approach. Since Gibbard is not quite so clear on why the attitudes in question bring with them the relevant logical commitments, it is harder to judge whether the commitment in question is essential to his approach.

Admittedly, it seems odd that one statement could have both contents, but so far as I can see that oddness is part and parcel of accepting the expressivist account. That account says that statements that function to express such attitudes have an expressive content that underwrites their logical functioning. To make an exception here, just because the statement also has clear propositional content, seems to me to be an ad hoc qualification of the theory to avoid uncomfortable consequences.³² In any case, the additional content had by the judgment, over and above its expressive content, should leave undisturbed any inconsistencies deriving from the expressive content alone. Thus, statement *A* will be inconsistent with any statement that is inconsistent with its expressive content. It would (on Gibbard's account) generate an inconsistency wherever the counterpart judgment about rationality would generate such inconsistency.³³

Consider now this argument:

- (1) If I express acceptance of norms requiring that I remain

³²It really does seem odd that a statement could have two kinds of content, both of which contribute to its logical functioning. But that idea has been part of noncognitivist proposals for some time. Stevenson, for example, distinguishes the emotive meaning of expressions from their descriptive meaning, and claims that moral terms have both. Thus, the proto-analysis he offers of judgments of goodness—'I approve of this; Do so as well'—contains both a descriptive and a prescriptive component (*Ethics and Language* (New Haven: Yale University Press, 1944), chap. 2; see also chap. 3, sec. 7). And R. M. Hare similarly distinguishes prescriptive meaning from descriptive meaning. Ethical terms have both; for the most general the prescriptive meaning is primary, whereas for more particular terms (such as 'industrious') the descriptive meaning is primary. See *The Language of Morals* (Oxford: Oxford University Press, 1952), 118–25. It is always a bit of a puzzle how such contents are supposed to be connected to one another, but given the commitments of the noncognitivist, it is not too surprising that the view yields such dual contents. Once one has postulated a special nondescriptive meaning for terms that makes them apt for the expression of attitudes, one has only to note that there are terms with clear descriptive meaning that are apt to express attitudes to wonder whether an expression might have both kinds of meaning.

³³Obviously, the logical behavior of utterances with overlapping content need not be precisely alike. For my argument to work, all I need is that a judgment will be inconsistent with anything not consistent with judgments having only a subset of its content. In other words, adding content to a judgment will not make it consistent with any judgment that it is inconsistent with. And that seems obviously correct.

- silent, as well as some norm requiring acceptance of those norms, and furthermore prescribe that others do likewise, it will be due to a phobia of public speaking. (Premise)
- (2) Acting on norms of which one expresses acceptance due to phobia is a manifestation of irrationality. (Premise)
 - (3) If an action is a manifestation of irrationality, it is not rational to do it.
 - (4) Remaining silent would be acting on norms requiring that I remain silent. (Premise)
 - (5) I hereby express acceptance of norms requiring that I remain silent, as well as some higher-order norm requiring acceptance of those norms, and furthermore prescribe that others do likewise. (Premise)
 - (6) Acting on the norms accepted in 5 would be acting on norms I express acceptance of due to phobia. (From 1 and 5)
 - (7) Acting on the norms accepted in 5 would be a manifestation of irrationality. (From 2 and 6)
 - (8) Remaining silent would be a manifestation of irrationality. (From 7 and 4)
 - (9) Remaining silent is not rational. (From 3 and 8)

This argument is valid, and its premises seem to be consistent. But if judgments of rationality should be analyzed as performing the complex expressive function assigned by Gibbard, the fifth premise, which is our statement *A*, has the same content that 'Remaining silent is rational' has. But that claim contradicts the conclusion of the argument. Since one of the premises of the argument would be inconsistent with its conclusion, then its premises must have been inconsistent all along.

As I said, it does not seem that these premises were inconsistent all along. The argument appears to capture a kind of reasoning a person might go through. And that reasoning does not seem to be a *reductio*, in that while it gives the person reason to conclude that remaining silent is not rational, it does not give the person reason for concluding that the fifth premise (or any other premise) was false. That, anyway, is the judgment of common sense, and expla-

nations of the logical relations among the contents of judgments owe that common sense a certain deference.³⁴

The advocate of the account thus faces a dilemma. Either she needs to say that some expressions of the attitudes in question carry certain logical commitments whereas others do not, or she must bite the bullet and find inconsistency where there seems to

³⁴One kind of reply that won't work is to plead that one is offering reforming analyses, rather than analyses of the actual meaning of the terms as currently used. The idea here might be that if the proposed analysis is not meant to capture what we currently mean, but is meant instead to offer a reform of our current practice that captures most of what is important, then what appear to be counterintuitive consequences of the theory won't be reasons to reject it. They will be counterintuitive only given that we use the words in an unreformed way. Once we adopt the reforms suggested by the expressivist theory, the currently counterintuitive consequences should lose their odd flavor as we adjust to new linguistic practices. (See Gibbard, *Wise Choices*, 32)

The reason this reply won't work is that the expressivist account of logical contradiction does not leave our nonevaluative judgments undisturbed. The account involved two components: (1) an analysis of evaluative judgments as expressions of certain complex attitudes, and (2) an account of logic that counts certain expressions logically inconsistent. A reforming definition is in effect a stipulation that we will use our evaluative terms in the manner suggested by the analysis that constitutes the first component. But the trouble complained of here was really trouble for the second component, the account of what is logically consistent and inconsistent with what. It was because the account must find certain expressions of attitudes logically inconsistent that it could not count the arguments in question consistent. And this was not dependent on the words with which those attitudes were expressed. (To see this, notice that the account will have to count expressions of certain attitudes logically inconsistent, whether they use a form of words including evaluative vocabulary or not.) A redefinition of moral terms should not affect the logical behavior of judgments not employing those terms.

I suppose Gibbard might be proposing a reforming definition of logical terms ('consistent', 'valid', and so on) as well. This would be more ambitious, and hence would come at the cost of likely additional departures from common sense over and above those required by reforming definitions of only evaluative vocabulary. My discussion below of a possible "bul-

be none.³⁵ The first horn seems preferable to the second, but the advocate of that horn owes us an explanation. Why is it that some such expressions of attitude commit us logically whereas others do not? I am not sure what form such an explanation should or could take.³⁶

5. A First General Moral

We are now in a position to take stock and find a moral to the discussion so far. We've examined four different expressivist analyses of evaluative utterances and their explanations of valid inferences. Each of these explanations has committed the account in question to finding contradictions where there seem to be none. Furthermore, the problem seems not to be idiosyncratic to Blackburn's attempts at expressive analyses. We have inductive reasons to suppose that further attempts will suffer the same defect.

I think there is a bit more to say to supplement the inductive argument for the general conclusion that further attempts will be liable to suffer from the same defect. For the defect is due to an essential feature of such accounts. Any expressivist account that wishes to explain how evaluative judgments can be inconsistent

³⁵Does Gibbard's most recent account avoid the problem? There he seems to abandon the analysis of the judgments in question as the acceptance of sets of higher- and lower-order attitudes. Instead, he claims (somewhat cryptically) that they stand to the acceptance of norms as fully factual judgments stand to the factual apprehension (Gibbard, "Reply to Blackburn, Carson, Hill, and Railton," *Philosophy and Phenomenological Research* 51 (1992): 970–71). Until that relation is filled out, it will be hard to tell whether the objection, as formulated here, applies. For without more details it is not clear what the state of mind in question is supposed to be.

In any case, the more recent account seems to make my first criticism of Gibbard harder to elude, in that it ties the inconsistency of conflicting moral judgments to the inconsistency of any fully detailed contingency plan—for what to do, think, or feel for every imaginable circumstance—that the norms might commend (*ibid.*, 973–74.) Thus, ruling out conflicting obligations becomes essential to the approach.

³⁶The explanation very likely would involve explaining why expressions of the attitudes with contents of the sort provided by the analysis must have the "logical form" exhibited by the logic defined in terms of norm-world pairs. Gibbard does not say very much about this. One approach, the one taken by Blackburn, is to argue that one is committed by the contents of the attitudes to making the appropriate inferences, on pain of irrationality. I argue below that this will not work.

with other judgments must extend the notion of inconsistency beyond an application to judgments whose contents are truth-evaluable. So, the noncognitivist's conception of "inconsistency" will be less stringent than our ordinary one, in that it will allow inconsistency where the ordinary account (applying only to truth-evaluable judgments) does not.

Now the expressivist's hope is that despite finding inconsistencies of a sort that the orthodox account would not, the expressivist notion will track what the orthodox account would have said were evaluative utterances interpreted in a realist fashion. Thus, since the realist will find a contradiction between "Lying is always good," and "Lying is not always good," the expressivist will want a translation that counts the judgments as inconsistent in the less stringent expressivist sense. And where two judgments taken in a realist fashion do not contradict, the expressivist will want there to be no contradiction in the broader expressivist sense. Still, the expressivist's sense of contradiction is an extension beyond the orthodox conception. Attitudes that would not count as contradictory on the latter account will count as contradictory on the former.

The objections so far have traded on just this feature of expressivist explanations of valid inference. Once we allow that there are ways to be inconsistent beyond the sort that applies to judgments expressing truth-evaluable contents, judgments without incompatible truth conditions can turn out to be inconsistent in this sense. I see no way for the expressivist to ensure that it will be only expressivist translations of judgments that the realist counts as straightforwardly in contradiction that will have this property.

I conclude that any account that enriches the notion of logical inconsistency so as to allow such inconsistency between non-truth-evaluable contents will be subject to similar objections. If I'm right, this would vindicate the critics' claims that logical consistency is one thing, the consistency invoked in the explanation another, and that it is a mistake to conflate the two. The criticism is that the major attempts to extend the notion of logical inconsistency all class as logical errors what are at best errors in substantive ethical doctrine, and very likely no errors at all.

6. Can The Expressivist Bite the Bullet?

I have claimed that the arguments presented as examples in my objections do not involve the logical contradictions that an ex-

pressivist analysis would entail that they have. And I think that I have ordinary judgment about logical contradiction on my side. But might not Blackburn (or Gibbard) resist the criticism offered by claiming that the arguments in the examples do in fact have contradictory premises, first appearances to the contrary notwithstanding? “Of course,” Blackburn might say, “the premises are logically inconsistent whenever there is a pragmatic incoherence in accepting them. That goes with the strategy. So the premises here *are* logically inconsistent.”³⁷

This is a difficult issue, because it is not clear how many of our pretheoretical judgments about logical implication Blackburn means to preserve, nor perhaps how many he should mean to preserve. But I think that we can make some progress. I’ll start by noting that the issue presents the expressivist with a kind of dilemma. Either the expressivist means to preserve as much as possible of our pretheoretical judgments regarding consistency over the evaluative domain, or she does not. If she does not, it becomes hard to see why we need an explanation of logical implications among evaluative judgments. If there is no commitment to our pretheoretical views regarding the logical relations between evaluative judgments, the expressivist should be biting a different bullet: She should deny that there are such logical relations. Thus, I think that the expressivist is better off grasping the first horn and trying to preserve as much of the pretheoretical “data” as possible.

Of course, consistent with that general goal one might think that jettisoning *some* commonsense judgments might be justified, if those judgments were not too numerous or too important. But it is worth noting that the project of explaining logical validity becomes easier as the expressivist expands the conception of logical inconsistency. The larger the number of statements that count as logically inconsistent, the more explanations we can give. It is the attempt to respect common sense about which judgments are in fact not logically inconsistent that makes the task nontrivial.

Among the pretheoretical judgments that I claim should be preserved is a commonsense distinction between logical consistency

³⁷This response on Blackburn’s behalf was forcefully urged against me by both Frances Howard-Snyder and Joe Mendola on independent occasions. My response to it has benefited enormously from suggestions by a reader for the *Philosophical Review*.

and mere pragmatic inconsistency or incoherence. When we look at arguments such as the kinds I have been using as examples in my arguments, we find ourselves making such distinctions. Take the second argument—the one used to undermine Blackburn's second account—for example. On first inspection, there is an inconsistency between the premises and the denial of the conclusion. We may also think that there is a pragmatic inconsistency between some of the premises and the conclusions, and therefore between some of the premises themselves. But the two kinds of inconsistency appear to be quite different. That, I think, is the most natural pretheoretical view. Biting the bullet, and classing both as examples of logical inconsistency, precludes us from making a distinction that we intuitively wish to make.

If we confined our view just to the evaluative domain, this might look like a standoff. The expressivist could admit that the explanation comes at some cost to our commonsense judgments that evaluative claims like those in the examples are logically consistent, but argue that the cost is bearable. But if we expand our view to nonevaluative discourse, we see that there are further costs as well.

In the domain of nonevaluative judgments, we find ourselves able to distinguish between logical inconsistency and merely pragmatic inconsistency or incoherence. Moore's paradox provides us with a set of claims that are not logically contradictory but that are pragmatically inconsistent: 'It is raining' and 'I don't believe that it is raining'.³⁸ As a general matter, we are able to distinguish logical contradictions from other kinds of incoherence. Thus, it is not only evaluative judgments that can display the kind of pragmatic inconsistency Blackburn uses to explain logical inconsistency.

As a result, biting the bullet and counting the premises of the argument in the evaluative example inconsistent either makes logical consistency in evaluative domains discontinuous with logical inconsistency for nonevaluative domains, or disturbs some commonsense judgments about consistency and inconsistency for nonevaluative domains. The first alternative violates what should be a desideratum for explanations of the logic of moral utterances: that

³⁸See, G. E. Moore, *Ethics* (New York: H. Holt, 1912), 125 (or reset ed. (London: Oxford University Press, 1947), 78) and G. E. Moore, "A Reply to my Critics," in *The Philosophy of G. E. Moore*, ed. Schilpp (Evanston: Open Court, 1942), 543.

this logic be as close as possible to logic for nonevaluative utterances. 'Lying is wrong' and 'I don't believe that lying is wrong' should count as no more logically contradictory than the analogous examples involving beliefs about rain.³⁹ The remaining alternative extends the scope of the conflict with common sense in a different way. Judgments like those involved in Moore's paradox, which we thought were not logically contradictory, will be deemed to be logical inconsistencies after all. We would have continuity between evaluative and nonevaluative domains, but no ability to distinguish pragmatic from logical inconsistency for any domain. Either alternative strikes me as exacting a high cost.⁴⁰

7. A Second Problem with Blackburn's Third Account

Returning now to Blackburn's treatment of logic in his third proposal, we should note that it suffers from a second defect. In explicating the meanings of conditional and disjunctive embeddings of moral terminology, Blackburn presents them as commitments to infer in certain ways, should certain attitudes be taken up, or given up. We are to think of them as analogous to a species of conditional intention. But it would be inaccurate to suggest that we have such commitments wherever we are prepared to accept such disjunctions or conditionals, despite a surface plausibility to that idea. If I believe that all swans are white or all ravens are black, I may have no commitment to infer that all ravens are black if I encounter a black swan. I may instead give up my belief in the original disjunctive claim. In fact, depending on my reasons for accepting the disjunction, it may be irrational not to give it up.

But it isn't hard to see why Blackburn was tempted by the claim

³⁹This particular point is reason to reject the kind of bullet biting that would save Blackburn's first account from the example I employ against it. The general point I'm urging here counts against the sort that would save the other two.

⁴⁰These costs will be costs whether the analysis is offered as an analysis of the concepts we already possess or as a "reforming analysis" of the sort I consider on Gibbard's behalf in note 34. The latter sort purport to offer us reformed definitions for terms like 'logical inconsistency' that capture most of what we could reasonably hope to say using our old unreconstructed vocabulary. The considerations above just underline some of the things we had hoped to say, but which would be barred to us using the reformed terms.

that one who accepts P or Q is committed to inferring Q , should she come to accept $\text{not}(P)$. He is trying to capture logical implication. And P or Q , together with $\text{not}(P)$, implies Q , not the disjunction $\text{not}(P \text{ or } Q)$ or Q .

8. A Second General Moral

The general problem here is that logical implications between contents and rational connections between attitudes with those contents are not always straightforwardly linked. P or Q , and $\text{not}(P)$, together always imply Q , but finding oneself believing both may or may not be a reason to believe Q . In attempting to make the link tighter, as he must in order to use one to explain the other, Blackburn ends up distorting the commitments of one who accepts the contents in question.

A theory that explains implication between contents as falling out of the rational connections between attitudes would not merely distort isolated areas. There are many cases where the inconsistency of contents is not tracked by the irrationality of accepting attitudes with those contents. It is not always irrational to accept every member of an inconsistent set of beliefs. For example, if it takes a large number of one's beliefs to generate the inconsistency, and if one has no reason to focus on any particular belief as less probable than the rest, it may be rational to accept all of them. It is most reasonable to suppose that some of one's beliefs are false, and that for any large set of beliefs one has, some one of them is false. But that belief, together with a large set of one's beliefs, and the belief that these beliefs are a large set of one's beliefs, will be inconsistent. Thus, rationality may require one to have a set of beliefs that are inconsistent.⁴¹

A defender of expressivist accounts may try to sidestep the worry by claiming that the connections between attitudes that underwrite "inconsistencies" between expressive utterances are not rational connections, but rather connections of some other sort. The idea would be to concede that the logical connections between the contents of judgments (or expressions of those judgments) do not neatly track the rational relations among attitudes, but to offer some other kind

⁴¹Gilbert Harman has long urged this and similar conclusions. See for example, *Change in View* (Cambridge: MIT Press, 1986), especially chap. 2.

of relation between attitudes that they do track. The problem with this response is that we need an independent characterization of the relationship between attitudes that underwrites logical inconsistency. That is, the expressivist needs to be able to say what the relationship is that all and only attitudes with “inconsistent” contents share, in a way that is not parasitic upon an independent judgment that the contents are inconsistent. And this relation will need to be of a sort that can figure in the justification of inferences, at least indirectly. Since we have some idea of what rationality requires, characterizing the attitudes as irrational to hold at once satisfies the independence requirement. And since rationality figures directly in the justification of inferences, it also satisfies the explanatory desideratum. But, rationality having been abandoned, a substitute meeting both desiderata must be found.

My suspicion is that none is available. The expressivist might try to work with some notion of “pragmatic inconsistency,” without relying on any claim that such inconsistencies are irrational. But many ways of characterizing such inconsistencies seem to presuppose that the contents in question are truth-evaluable.⁴² Even if that worry can be sidestepped, our first general conclusion remains. Expanded notions of inconsistency will very likely count consistent premises inconsistent.

We thus have two reasons for pessimism about the noncognitivist’s approach to logic. By extending what is to count as inconsistency, the approach threatens to class consistent views as logically inconsistent. And relatedly, rational connections between beliefs and logical connections between their contents are two different things, which impinge upon one another but are yet distinct. These considerations strengthen the suspicion broached by earlier critics of Blackburn’s accounts, that logic is one thing, pragmatic incoherence and the rationality of beliefs and desires another. Attempts to explain one in terms of the other will distort our account of which propositions are consistent with which, or our account of which inferences are rationally warranted, or very likely both.

University of Nebraska, Lincoln

⁴²For example, one kind of pragmatic inconsistency occurs when the very assertion of a claim refutes the truth of the claim, as in “I have written nothing.”