



Expressivism, Morality, and the Emotions

Justin D'Arms; Daniel Jacobson

Ethics, Vol. 104, No. 4. (Jul., 1994), pp. 739-763.

Stable URL:

<http://links.jstor.org/sici?sici=0014-1704%28199407%29104%3A4%3C739%3AEMATE%3E2.0.CO%3B2-G>

Ethics is currently published by The University of Chicago Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/ucpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Expressivism, Morality, and the Emotions*

Justin D'Arms and Daniel Jacobson

A. J. Ayer first put forward the doctrine of expressivism in his notorious 1936 manifesto, *Language, Truth, and Logic*. There he wrote, "In saying that a certain type of action is right or wrong, I am not making any factual statement, not even a statement about my own state of mind. I am merely expressing certain moral sentiments."¹ Since then expressivism has, along with its competitors, suffered the vagaries of metaethical fashion. Two arguments have been deemed most damaging to it: the problem of embedded contexts and a phenomenological argument concerning the defensibility of values.² The recent rehabilitation of expressivism is buoyed largely by the efforts of two proponents of similar views, Simon Blackburn and Allan Gibbard, to meet these challenges.³ The extent to which they succeed is, of course, a matter of contention; and these issues are no doubt worthy of the extensive discussion they have been receiving. But that discussion bears on the expressivist strategy quite generally, in abstraction from issues specific to moral discourse. Too little attention has yet been paid to the specifically ethical aspects of Gibbard's extremely fertile and far-reaching expressivist theory. That will be our task here.

We want to return to an older problem, for which Gibbard's norm expressivism offers the first comprehensive solution. Ayer held that

* We would like to thank Elizabeth Anderson, David Copp, Stephen L. Darwall, Heidi Li Feldman, David Hills, Peter Railton, and an anonymous referee for their helpful comments on an earlier draft, and especially Allan Gibbard, for many patient and generous discussions.

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

2. The locus classicus of each is, respectively, Peter Geach, "Assertion," *Philosophical Review* 74 (1965): 449–65; and David Wiggins, "Truth, Invention, and the Meaning of Life," in *Needs, Values, Truth: Essays in the Philosophy of Value* (Oxford: Blackwell, 1987).

3. Compare Simon Blackburn, *Spreading the Word: Groundings in the Philosophy of Language* (Oxford: Clarendon Press, 1984); and Allan Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, Mass.: Harvard University Press, 1990).

Ethics 104 (July 1994): 739–763

© 1994 by The University of Chicago. All rights reserved. 0014-1704/94/0404-0001\$01.00

value judgments are expressions of sentiment—but just what mental state is it that they express? Where do expressivism's terms of art, the pro and con attitudes associated with evaluation, fit in a general economy of the mind? Three considerations sharpen this worry: (1) A principal attraction of expressivism is its promise to account for the magnetism or normativity of values as a natural phenomenon. To make good on this promise, though, the theory must cohere with our best scientific-cum-philosophical picture of the mind. *Sui generis* "moral sentiments," which for the expressivist must neither be beliefs nor mere desires, lack any obvious place in that picture. (2) Even Ayer recognized that value judgment has to be more than an expression of the speaker's approbation, or our conflicts in attitude would give no point to ethical dispute. Charles Stevenson glossed this additional component as an injunction to others to "do so as well." Whatever mental state is expressed in a value judgment must manifest some prescriptive, outward-looking aspiration. (3) Finally, as 'good', 'right', 'beautiful', et al. are not synonyms, expressivism must distinguish between the different varieties of evaluative sentiment. The expressivist needs both to identify a family of mental states expressed in value judgment and to single out those family members implicated in specifically moral evaluation.⁴

Norm expressivism is distinctive in offering a naturalistic picture of these mental states which allows for disagreement in attitude and distinguishes moral evaluation by associating it with certain familiar emotions. This constitutes a significant advance for ethical noncognitivism; but with this greater specificity comes increased danger that the theory will fail to jibe with the practices it seeks to describe and explain. We propose to investigate the theory's adequacy to morality.

A CHALLENGE TO NORM EXPRESSIVISM

According to norm expressivism, all normative judgments are expressions of a single, psychologically basic type of mental state: norm acceptance. Any state that could meet the expressivist's requirements must inhabit the murky region between the cognitive and the conative. Too much like a belief and moral discourse looks to be descriptive after all. But too much like a desire and it will fail to display the discipline of actual moral dispute, with its demands for consistency and for interpersonally available justification. The state of accepting a norm is difficult to characterize, and Gibbard seeks to point to its place in human life more than to describe it. Perhaps the most illumi-

4. Compare, respectively, Wiggins, "Truth, Invention, and the Meaning of Life"; Charles Stevenson, "The Emotive Theory of Ethical Terms," *Mind* 46 (1937): 14–31; and David Wiggins, "A Sensible Subjectivism," also in *Needs, Values, Truth*.

nating thing to be said is that to accept a norm is to judge something *warranted*. The notion of warrant in play here is the familiar, epistemic one; but Gibbard applies it broadly, to assess the warrant of feelings and desires as well as beliefs. It is an inherently normative notion: to accept a norm, say for a belief, is to bestow on it a kind of endorsement, one of rational justification. Yet this need not issue in an all-things-considered judgment of what to believe, since considerations of warrant do not always settle that question.

To call a judgment of fact warranted, then, is to express acceptance of ordinary epistemic norms for the warrant of belief. The traditionally evaluative judgments, in contrast, express acceptance of norms for the warrant of other mental states: emotions and desires. (These could be termed quasi-epistemological or “epistemological” norms, to indicate that they trade on the same notion of rational—as opposed to moral or pragmatic—justification.) Like Blackburn’s quasi-realism, this theory distinguishes facts and values sharply, but only at the metaphysical level; they are strongly assimilated epistemologically. Since all normative discourse is reconstructed in terms of norm acceptance, including norms for the warrant of belief, “from the point of view of their justification, they are on a par; factual beliefs and normative judgments stand or fall together.”⁵

Gibbard urges that, for several reasons, norm acceptance deserves a place in the naturalist’s ontology of mental states. First, it allows us to make sense of certain everyday phenomena (such as weakness of the will) which don’t find a comfortable place within commonsense psychology’s belief-desire model of the mind.⁶ Second, the account gives a powerful, unifying treatment of all normative discourse. Finally, Gibbard suggests that the state of norm acceptance plays an important explanatory role in the most plausible picture of human evolution.⁷ Furthermore, in its appeal to the crucial role of “conversational pressures” in normative discourse, Gibbard’s evolutionary speculation offers a novel account of the prescriptivity of value judgment. All this is, at first blush, quite compelling. But whether the speculative psychological and evolutionary stories are acceptable will depend largely on how convincing we find the underlying theory of normative discourse.

5. Gibbard, p. 34.

6. *Ibid.*, pp. 56–61.

7. Gibbard, throughout, but see especially p. 57. Explanation of language evolution in humans often adverts to our need to communicate about features of our material surroundings. But another function of language is to allow coordination of attitudes. In order to live cooperatively, we had to be able to influence and be influenced by the feelings of others. Norm-expressive language and the state of norm acceptance evolved together because they facilitated the kinds of coordination crucial to complex social life.

All norms, we've said, are norms of warrant—they're about what it makes sense to think, to want, or to feel. Gibbard writes about what is "rational," but he grants that this term has a misleadingly learned flavor, and he uses the more vernacular expression "makes sense" synonymously. That locution is not entirely satisfactory either, though, since it's used to explain as well as to justify. It makes sense to be nervous when speaking to large groups, one says, without meaning to recommend the feeling. Perhaps these uses can be distinguished by saying it makes sense that you were nervous, but it did not make sense to be nervous. We're interested in the justificatory "makes sense to" rather than the explanatory "makes sense that." Thus, we'll stipulate, following Gibbard, that to judge that an attitude (whether a belief, a desire, or an emotion) makes sense is to endorse it in a specific way: it is to express acceptance of norms that find the attitude warranted.

What unites such disparate evaluative judgments as those concerning goodness, rightness, and beauty, then, is that each expresses acceptance of some type of norm. In a broad sense, all of these are ethical judgments, but Gibbard is especially interested in a subset of these judgments: those concerning our notions of right and wrong. These judgments constitute what he calls morality "in the narrow sense"; they express norms for the emotions centrally implicated in our moral system: guilt and anger. Other cultures might key their moralities to different emotions. Gibbard allows that there are cultures with shame-based moralities, for instance, in which complex quasi-epistemological lore surrounds the rationality of that emotion. Thus Gibbard's central analysis refers to a notion of wrongness which is to some degree culturally specific:

What a person does is morally wrong if and only if it is rational [makes sense] for him to feel guilty for doing it, and for others to resent him.⁸

The suggestion that feelings and desires—the traditionally "noncognitive" attitudes—can be more or less warranted by different occasions will strike some readers as odd. (Indeed, we'll raise some doubts about this rationalism of the sentiments presently.) But it is important to see that common sense does ordinarily distinguish justified and unjustified fear and hope—as well as guilt and anger. Still we need to ask if the norm expressivist analysis does justice to morality. Consider the following scenario.

Mother has grown older, and grown mentally ill. She makes increasingly exigent demands on the family. Her illness is degenerative.

8. Gibbard, p. 42. In fact this is only the "crude version" of his analysis, and his final version adds several refinements, some of which are irrelevant to present purposes. For simplicity, we will ignore them all here, raising relevant qualifications as they arise.

She always feared being “put away”; you know she wants to stay at home, but you have real doubts about your ability to care for her. And you also see the effects of the tension, pity, and finally resentment on your family. Your spouse, who has been patient and helpful throughout, is beginning to show the strain. The children are restless. On any plausible normative picture, there are ample reasons for the conclusion you’ve been avoiding. In the end, you decide to put her in a nursing home. You’re convinced this is the right thing to do, though you know you’ll feel guilty for doing it. But does your guilt make sense? If you think it does, and yet that your action was not wrong, there are problems for the analysis.

Certainly it makes sense *that* you feel guilty; guilt is to be expected under the circumstances. But our concern, recall, is with what it makes sense *to* feel—with the justificatory rather than the explanatory usage. If guilt is merely explicable here, and not also justified, then the norm expressivist analysis remains unscathed.

But surely there is some plausibility to saying that it makes sense to feel guilty. How else does it make sense to feel about what you’ve done? It’s clear that some negative emotion is called for. You’ve deliberately violated the strongest wish of someone utterly dependent upon you, someone whom you owe most profoundly. Nor does the irony of the role reversal escape you. When the power and responsibility were hers, she nurtured you. Now the dynamic has changed, and you’re choosing, with the cold comfort of good reason, to sacrifice her interests for the sake of the family. Guilt is what you do feel, and surely an emotion in the neighborhood—a negative emotion directed at one’s own action—is warranted; still you aren’t forced to admit that what you actually feel is what you should.

One may want to reply that it doesn’t make sense to feel guilty because it would be counterproductive. After all, lingering guilt will only harm your family further, with no obvious benefit. It’s tempting to evaluate emotions for the good they will do, rather than for whether they are warranted by the situation. But this approach confutes the rationalism of the feelings upon which norm expressivism is predicated. Feelings and desires are supposed to be, in their justification, analogous to beliefs.⁹ Hence pragmatic considerations about what to feel in a given circumstance have to be taken as they would for what to believe: as considerations of the *desirability* of the attitude. What it

9. Furthermore, it’s crucial for expressivism to secure a strong distinction between the epistemic justification of feelings and their pragmatic justification. If the account is to capture our actual normative practice, it must make sense of the tendency to defend normative judgments with reasons, and of the pressures toward consistency in normative discourse. If some normative judgments are judgments of what feelings make sense, then that must be settled by norms of warrant, not pragmatics.

is desirable to believe may not be what is justified—as is routinely demonstrated, in philosophical fiction, of credulous but terminally ill husbands with unfaithful wives. What it is desirable to believe or to feel, writes Gibbard, is “a different judgment from the ones I want to pursue. The judgments I want us to consider are ones of warrant.”¹⁰ We can accept that guilt is not desirable under such circumstances, indeed we can try to expiate it, but these are not grounds for calling it unwarranted.

Nor is the question what to feel, all things considered, nor what it would be right to feel—not even if such feelings were more volitional than we think. Few people who agree that you’ve done the right thing will criticize you for encouraging the family to forgive yourselves. We can consistently hold that sometimes there is nothing wrong with not actually feeling even warranted guilt. But these considerations too are irrelevant to whether the guilt is warranted. The primary question here is not what it’s right to feel, but what it’s right to do. This distinction is easily obscured because the analysis of the judgment that an action is right makes use of what it would make sense to feel about it. But we must keep the issues apart. If our feelings too can fall under moral evaluation, despite their being so impervious to the will, then norm expressivism glosses questions about the rightness of *actually* feeling guilty like this: “Would it make sense to feel guilty, and for others to resent you, *for not feeling guilty* in this situation?” (i.e., having put mother in care). This is not equivalent to the original question, which directly concerned the act. Our own opinion is that it isn’t wrong not to feel even warranted guilt, when no wrong was done and guilt would be counterproductive. But agree with us or not, this doesn’t bear on whether the guilt for your action makes sense; that question must be settled on grounds of warrant.

So far, we’ve described a scenario and tendered a normative judgment upon it: that it makes sense for you to feel guilty while thinking that you haven’t done wrong. This judgment poses some problems for Gibbard’s theory. However, the philosophical genre is such as to raise the expectation that our scenario is being put forward as a counterexample to the norm expressivist analysis; but that isn’t quite the case. Although we’ll argue that our favored way of describing and judging the scenario is the most compelling, the success of the argument does not turn on our ability to convince the reader to share our intuition. We can accommodate normative disagreement here; we won’t dig in our heels against all who deny that guilt is appropriate. Instead, we’ll argue that their judgment rests on a thought that is unavailable to the expressivist. Even were our description convincing,

10. Gibbard, p. 37.

though, the case would not be a counterexample, because there is another twist to the analysis which might have sufficed to rescue it. (We'll ultimately argue that it cannot.) But finally, we don't need the case to be a counterexample. It is enough that our intuition is coherent—that it isn't confused—for the expressivist analysis to be jeopardized.

Again, some caution is necessary. The open question argument has become nearly as familiar a strategy to moral philosophy as the counterexample. G. E. Moore noticed that for any property P, we can sensibly ask, "Given that an object X has P, is X good?"¹¹ Moore's official procedure was one of conceptual comparison. We might have examined the concepts Wrong and Apt-for-Guilt, declared them distinct, and scuttled the analysis. However, the versions of the open question that remain viable do not trade in examination of the conceptual realm. Instead, they gloss Moore's question as querying the objective prescriptivity (or "to-be-pursuedness") of descriptive properties. For any such property P (even one that moved us with metaphysical necessity), the question whether P is worth pursuing will always be a meaningful one, since, as the existentialists have shown, we can alienate ourselves from even our most impervious motives sufficiently to doubt their merit.

Noncognitivists have seized on this observation because it is, for them, at the heart of all evaluative discourse. The fact of disagreement over what, all things considered, to do is nothing less than what gives point to normative dispute. If that is right, then they have motivated a theoretical role for a term that functions purely to express endorsement. The open question argument is then called upon to demonstrate that no term which serves this role can have any descriptive semantic content, because the descriptive judgment and the endorsement can always come apart.¹² The analysis, having precluded a substantive and meaningful claim as trivial or incoherent, is rejected. Were this our strategy, we'd have an argument almost as short as Moore's. The (simplified) expressivist analysis identifies the judgment that an act is wrong with the judgment that it warrants guilt. We've offered a scenario where these judgments come apart—for us, at any rate. It seems

11. G. E. Moore, *Principia Ethica* (New York: Cambridge University Press, 1903). Of course, he would have excepted one nonnatural property, Good; but it seems that nonnatural properties too are susceptible to the argument. Moore's question leads to the paradox of analysis—how do you know Good when you find it? The critique of "naturalism" is thus even more of a red herring than Moore himself admitted. (On this last point cf. William Frankena, "The Naturalistic Fallacy," *Mind* 48 (1939): 464–77.)

12. You might call this the Gloria Vanderbilt theory of normative discourse: No evaluative term can be too descriptively thin, or too motivationally rich. (This theory has been romantically linked to a semantic partner she met at the Non-Cognitivists' Ball, the Joe Friday theory of meaning: Just the facts, ma'am.)

we don't even have to persuade you, but merely to convince you that we're not linguistically incompetent. The norm expressivist analysis thus closes a meaningful question of its own, over how it's appropriate to feel about certain morally permissible acts; hence, the analysis fails.

But contemporary philosophers would be justly suspicious of our bald insistence that a meaningful question is being closed here. Any philosophical analysis strains its concept—this thought has quickly become a platitude of metaethics. Modern metaethical analyses are typically put forward as reforms of the natural language concept they purport to analyze. If the noncognitivist has successfully motivated a theoretical role for the perfectly thin value term, independent of the requirements of his theory, there need be no actual term available that fits the job description without strain. The account must then be judged not by dated methods of conceptual or linguistic analysis, but by its fruits and the fruits of the surrounding theory—by how much it can explain and render intelligible, and how much of our old notions it requires us to dismiss as unintelligible.

This is Gibbard's strategy as well, when he criticizes descriptivistic analyses. As he puts it, their difficulty is "not merely that every time one loophole is closed, others remain. It is that a single loophole remains unpluggable by any descriptive analysis."¹³ (That is, the question whether to endorse the verdict of some privileged epistemic position is made unavailable.) Furthermore, Gibbard thinks this unplugged loophole is of the essence of normative thought. As the battle lines have traditionally been drawn, expressivists reject descriptivism for closing this existential query, while others suggest that expressivism's costs are even more extreme: it requires us to reconstrue the very nature of normative speech acts, which on their surface look for all the world like property ascriptions.¹⁴ This can easily seem a stand-off—at any rate, we find it difficult to commensurate such costs. Instead, we propose trying a different approach.

Norm expressivism closes questions about how it makes sense to feel toward some actions one does not think wrong. This doesn't queer the theory forthright, but we suggest that it poses a substantial problem. Can rational guilt and anger really be teased into the shape of moral judgment? Which features of ordinary thought about these emotions can the expressivist exploit, and which must he explain away? To assess the costs of Gibbard's theory fairly, we need to take a good look at how morality actually engages the emotions. We'll now investi-

13. Gibbard, p. 33.

14. Compare, e.g., Sabina Lovibond, *Realism and Imagination in Ethics* (Minneapolis: University of Minnesota Press, 1983), p. 26, who deploys Wittgensteinian arguments to conclude that "if something has the grammatical form of a proposition, then it is a proposition."

gate guilt and anger, the putatively moral emotions, hoping to persuade the reader that Gibbard's attractive analysis comes at a high price.

GUILT AS A MORAL EMOTION

Maybe it isn't guilt that you think warranted in our scenario, but some cognate—another negative emotion, directed at one's own action. There are a variety of terms in the neighborhood: remorse, compunction, regret. If the norm expressivist can persuade us that when we think we've done the right thing we only judge one of these other emotions warranted, the pressure on the analysis is relieved. But how plausibly can we even distinguish these emotions, much less the norms we accept for when they are apt? The question is particularly troublesome for norm expressivism, which demands a compatible theory of the emotions, one that treads a delicate line between rationalism and cognitivism. The theory is committed to rationalism, the view that emotions (at least, those implicated in evaluation) can be more or less justified—that they are amenable to norms of rationality. But it must deny cognitivism, the thesis that (these) emotions are essentially constituted and individuated, at least in part, by their propositional content. Normative judgments are not supposed to be propositional; the defining expressivist thought is that value judgments are not apt for truth or falsity, because they are not assertions.

In fact, the burden on norm expressivism is even greater. The theory cannot allow normative judgments, even if they aren't propositional, to figure into the constituents of the crucial emotions. If they did, norm expressivism's derivation of evaluation from the acceptance of norms governing these emotions would be circular. Thus Gibbard rejects even the weaker thesis he calls "judgmentalism" about the emotions. He tests the compatibility of theories of emotion with norm expressivism this way: "Do they let us explain particular emotions like anger without invoking normative concepts like being to blame? Unless they do, the norm-expressivist account of moral judgments is circular: it invokes anger and guilt to characterize moral judgments, but we must understand moral judgments already if we are to characterize guilt and anger."¹⁵

A philosopher of a different temperament might think progress was made by identifying moral judgment even circularly with a particular emotional response, but expressivism carries ambitions of noncircularity. If the program is to convince us that statements appearing to ascribe an elusive property (objective prescriptivity) are in fact merely expressions of a sentiment that makes no (erroneous) claim to truth,

15. Gibbard, p. 148.

then the sentiment in question must be demonstrated not to include a judgment. A circular account of the sentiment offers no such demonstration. Indeed, one recent view employs a circular relation between property and response to vindicate the idea that evaluative judgments have truth conditions.¹⁶ Thus Gibbard is committed to arguing for rationalism and against cognitivism, and he does so powerfully. We fear, however, that his arguments sometimes draw on conflicting considerations which norm expressivism cannot simultaneously endorse. Our strategy is not fully to embrace cognitivism nor reject rationalism, but to question whether the delicate balancing act that norm expressivism is committed to can be pulled off. (This strategy should be especially congenial to those, like us, who are skeptical of the prospects for any grand theory of the emotions.)

Can expressivism plausibly claim that it isn't guilt we think apt, but remorse, compunction, or regret? The first two seem, in ordinary usage, to be virtual synonyms of guilt. At any rate, guilt is explicitly implicated in any definition one is likely to find in a dictionary. Still, there may be room for some philosophical sharpening of these concepts. To our ears, compunction suggests a less intense reaction than guilt. We have pangs of compunction; we are plagued by guilt. But this is not a difference that the theory can use. An action does not fail to be wrong just because it is less than momentous; if the difference between guilt and compunction is merely one of degree, then the two emotions should be analogously related to wrongness. While some philosophers distinguish between guilt and remorse, Gibbard explicitly eschews doing so, and uses the terms synonymously.¹⁷ So the best prospect for norm expressivist purposes may be regret.

To make this move, the theory must hold that it is regret, not guilt, we think makes sense when we judge that we acted rightly, yet find some negative emotion to be warranted. The trick is, we must be persuaded of this while resisting the suggestion that what distinguishes guilt from regret is precisely that guilt is about blaming yourself, while regret is not. For that would be cognitivism (or judgmentalism anyway, which is just as bad). Gibbard leans heavily on one argument against cognitivism, which he takes to be decisive. He expresses it in two ways, which are on his analysis equivalent: (1) We can feel guilty for something we've done and yet think we were right to do it. (2) We can feel guilty and yet think it makes no sense to feel guilty.

These possibilities are supposed to tell against the cognitivist thesis that guilt is essentially about blaming yourself. To a cognitivist, an emotion is unwarranted only when its belief component is unjustified.

16. Compare Wiggins, "A Sensible Subjectivism."

17. Gibbard, pp. 296-97.

In the third person case this is unproblematic; to judge someone else's guilt unwarranted is simply to think the belief on which it rests is irrational. What's more difficult is for an agent to acknowledge his own emotion to be unwarranted, for according to cognitivism this commits him to two thoughts that are in tension. He believes something and yet thinks it unjustified. But what grounds do we have for attributing the unjustified belief to the agent? Gibbard is suspicious of the half-beliefs, or beliefs of the heart but not the mind, that he grants we're tempted to ascribe. He asks, "What do any of these things mean, though, but that [you feel guilty]?"¹⁸ Nevertheless, much recent literature on the philosophy of emotion has defended some kind of cognitivism, in the face of these objections. It may seem peculiar to ascribe a belief in some cases of unwarranted emotion, but it's surely not unintelligible. Indeed, something similar occurs in another context, when we display the phenomenal or behavioral symptoms of beliefs that we resist embracing. We find ourselves still expecting to meet a friend somewhere, though we know our reason for forming that expectation has been belied.¹⁹ Do we really still somehow believe our friend will be there? If not, we need another way to understand our recalcitrant expectations. To make this admittedly peculiar feature of unwarranted but felt guilt tell strongly against cognitivism, Gibbard must offer a more plausible alternative.

What is needed is a full-fledged theory of the emotions, and in fact Gibbard has one to offer. He calls it the adaptive syndrome theory.²⁰ On this view, the various emotional states are the product of different, highly specific, adaptive mechanisms. They will normally involve different typical causes, different manners of expression, and different tendencies to action. Thus we can in principle individuate emotions by pointing to disparate adaptive mechanisms underlying the various emotional states. We should not expect, though, that there will be a one-to-one correspondence between physiological mechanism types and emotion types. "What unites the various specific physiological states that constitute anger may not be a similarity that would strike a physiologist who had put all thought of function aside, but a similar-

18. *Ibid.*, p. 130.

19. This case is drawn from Ronald de Sousa, *The Rationality of Emotion* (Cambridge, Mass.: MIT Press, 1987), p. 154; cf. also Malcomb Budd, *Music and Emotion: The Philosophical Theories* (London: Routledge & Kegan Paul, 1985); William Lyons, *Emotion* (Cambridge: Cambridge University Press, 1980); and Robert C. Solomon, *The Passions* (Garden City, N.Y.: Doubleday/Anchor, 1976).

20. He also discusses another view, the attributional theory. For brevity's sake, we leave that view out of our discussion. It's clear that Gibbard's sympathies are not with it, and it doesn't seem to present a helpful alternative with respect to the matters we raise here.

ity of biological function."²¹ So if this approach is to succeed, we must be able to point to the various adaptive syndromes by picking out discrete evolutionary functions for them. Gibbard does just that, telling some plausible stories about the different roles played by various emotions.

Guilt is said to be the expression of a mechanism that exists in human populations because feeling guilty was historically an adaptive response to anger; it opened the door to reconciliation, promoting cooperation between parties rather than (potentially expensive) conflict. Whereas other animals possess a fairly limited range of responses to aggression (submission, flight, combat), humans needed to be able to mitigate anger in socially useful ways. Guilt evolved to meet that need. If the picture is right, then we should find that we ordinarily feel guilt in many situations where others are or could reasonably have been justified in being angry with us. And since our norms for how it makes sense to feel accord a strong presumption of warrant to the ways we actually do feel, we should also find ourselves tending to think guilt warranted in those circumstances we think warrant anger from others. These predictions seem to be borne out in our ordinary thought.

With this schema in hand, the expressivist can perhaps finesse the question of whether we have a name for the sought-after emotion that contrasts with guilt. But we in turn can demand to be shown a genuine theoretical role for that emotion (which we'll call regret) other than that of salvaging the theory. If the adaptive syndrome theory is to allow norm expressivism a distinction between guilt and regret that is not simply ad hoc, the suggestion must be that regret played some distinctive adaptive role in human evolution. Were this true, we could in theory distinguish guilt from regret by pointing to the different underlying adaptive mechanisms that produce them.

It will still be difficult to show that this difference underlies our first-personal judgments about occurrent emotions. Though the difference between guilt and regret might in fact consist in their being expressions of different adaptive syndromes, that will not be the sort of difference we typically appeal to when trying to decide what emotion we are feeling, or what emotion we think warranted. Still, if ordinary emotion talk successfully refers to the operations of some underlying adaptive mechanisms, we should find that the sorts of consideration we adduce in trying to decide whether it is guilt or regret that we feel reflect differences in the functions of these adaptive syndromes.

What is the best argument for the conclusion that it makes sense to feel regret, not guilt, in our scenario? Unfortunately for expressiv-

21. Gibbard, p. 134.

ism, the most compelling argument seems to be that you didn't do anything wrong. Yes it's sad, we're inclined to say, but you shouldn't blame yourself. This way of talking urges that the basic difference between thinking guilt or regret warranted is that the former includes a judgment of fault that the latter does not. But this would sneak moral judgment into judgments of what it makes sense to feel, which the expressivist must not do.

What he can say is that, though we may be tempted to defend the claim that guilt is unwarranted by proffering a moral judgment, we don't need to do so. We could offer purely descriptive considerations.²² We could say, for instance, that the reason it doesn't make sense to feel guilty is that mother's presence was upsetting the kids. Here we've offered a reason why it isn't guilt that is warranted, which doesn't tell against the warrant of regret and avoids any obviously normative concepts. If it worked, that would presumably be because we accept some further judgment to this effect: it does not make sense to feel guilty if you've acted to spare the children's feelings. But in fact we accept no such judgment. If you fire your foreign housekeeper because his presence makes the kids unhappy, it makes sense to feel guilty for what you've done, and doubly so for pandering to the children's naive bigotry. It's clear that the children's feelings alone can not settle the warrant of guilt. Perhaps further considerations such as these would convince you that it is not guilt that's warranted. The trouble, though, is that it seems the success of such an argument depends on its convincing (or reminding) you that you hadn't done wrong. And this is not a reason expressivism can offer.²³

It's not that we're ready to embrace a cognitivist account of guilt after all. To do so would be to impute a tension to the thought that you feel guilty while judging that you did the right thing, and would make contradictory the suggestion that the guilt may in fact be apt.

22. Gibbard says something like this at p. 149.

23. This is tricky. Since the expressivist identifies thinking something wrong with thinking guilt (and anger) warranted, it is open to him to urge you not to feel guilty because you didn't do wrong. This would just be to urge you not to feel an unwarranted emotion. What it means to say the act wasn't wrong, according to norm expressivism, is that guilt isn't warranted. Our strategy has been to argue that it's natural to think it makes sense to feel guilt here; we're considering how the expressivist can persuade us that we actually think regret, and not guilt, warranted. In order to keep his analysis informative, the expressivist must avoid resting his argument on the fact that we judge ourselves not to have acted wrongly. If the fact that we think we acted rightly is a reason for thinking guilt is not warranted (as opposed to simply another way of describing the very same thought), then we cannot accept as informative any analysis of wrongness that appeals to guilt. So the expressivist can utter the words "It doesn't make sense to feel guilty because you didn't do anything wrong," but this will only be an exhortation; it will not be a reason.

And this seems wrong. Furthermore, to assert the cognitivist thesis against Gibbard would be question begging. We're even willing to accept *ex hypothesi* the adaptive syndrome theory, to this extent: Grant that there is an emotion that serves the functional role of palliating anger. And grant too that guilt—the commonsense psychological concept—is sufficiently proximate to warrant a philosophical analysis associating it with this functional role. To do this, though, is to forswear trading on aspects of the concept's role in ordinary language that are incompatible with the analysis. The fine distinctions between guilt and its neighboring emotions, *which norm expressivism needs to make if it is to sustain its rationalist tenet that guilt cannot in this case be warranted*, seem to require discriminations that can be made only by trading on the cognitivist tendencies that do run deep in our ordinary thinking.

The norm expressivist has two avenues of response. He can produce a robust account of regret (or whatever is the alternative emotion) and, upon producing it, convince us that only this emotion could possibly be warranted in our case.²⁴ Or he can concede that guilt may be warranted in this case and look for other ways to avoid the damning conclusion. We'll consider two such arguments.

One might hope some progress could be made by introducing the notion of *prima facie* wrongness. There are really two options here. First, the expressivist could say that the judgment that it makes sense to feel guilty over an act is a judgment merely that it is *prima facie* wrong; then he can grant that guilt is appropriate in our scenario, though no all-in wrong is done. But all-in wrong is the concept expressivism is committed to analyze, and this suggestion just reopens the old problem: What kind of mental state is expressed in those judgments? (Worse, the most likely suspect, guilt, is already committed.)

Second, the expressivist could continue to identify the judgments that an act is wrong all-in and that it warrants guilt. Our temptation to say guilt makes sense in this case—where no wrong is done—must then be explained away. Only an *aspect* of the situation warrants guilt (it was putting mother away, after all), but guilt is not appropriate all things considered. The act then is only *prima facie* wrong. This suggestion raises a similar question. Is the judgment that an aspect of a situation warrants guilt—that is, that it's *prima facie* wrong—another kind of normative judgment? If so, what is its theoretical role, beyond that of salvaging the theory? Notice that what the expressivist can't say is that this notion captures a cognitivist's likely catalog of the typically wrong acts (e.g., betrayals or lies). For in fact we feel no

24. We've registered our doubts about whether this can be sustained. But we can't rule out the possibility by exhaustion, only examine candidates as they emerge.

temptation whatsoever to feel guilty over lying to the murderer Kant famously imagined, who knocks at our door, ax in hand, inquiring as to the whereabouts of our friend.²⁵

The other option is to say that the reason you should feel guilty is that you really have done wrong. If so, norm expressivism would be off the hook. We've assumed that you could not have acted wrongly because all other available options were worse. But this is somewhat controversial: some philosophers deny that ought implies can, precisely because of "tragic cases" such as ours, where they say you can't help but do wrong. The denial of "ought implies can" does not help norm expressivism, though, for two reasons. First, because tragic cases are just one place where the problem we've been developing arises. It comes whenever you've done something that hurts someone toward whom you have a special responsibility, but where there are strong grounds for denying that you've done wrong. Just one such ground, albeit perhaps the sturdiest, is that anything else you could have done was worse.

Second, norm expressivism is not supposed to dictate any particular normative conclusions, but to describe what we're doing when we make normative judgments. It would be odd for the analysis to yield the result that no competent moral agent accepts that ought implies can. This would be a very substantive result, where none is wanted. And it would be difficult to make such an argument cogent, since the plausibility of the maxim being denied will seem to many more secure than any grand theory of normative discourse could be. In any case, this route does not tempt Gibbard. Actually, he takes it as a benefit of his analysis that it offers an interpretation of the dispute over "ought implies can" that leaves matters unresolved, but more perspicuous: "Are there situations in which, no matter what the agent does, it will make sense for him to feel guilty for having done it, and for others to be angry at him for having done it? Guilt here must be more than mere compunction, and part of the problem may be whether a distinction between compunction and full-fledged guilt can be drawn."²⁶

We think this interpretation too even-handed—that it makes too much sense of the denial of the maxim. While it's true that there has been some controversy recently over whether ought really implies can, the maxim has struck many as an ethical platitude, and the counter-arguments can seem vainly heroic. The issue is too complex to be

25. Compare Immanuel Kant, *Critique of Practical Reason and Other Writings in Moral Philosophy*, trans. and ed. Lewis White Beck (Chicago: University of Chicago Press, 1949), pp. 346–50. Of course Kant thought it was wrong to lie to the murderer, but so much the worse for Kant.

26. Gibbard, p. 53.

adequately handled here; but we can briefly illustrate how our analysis of the case at hand suggests a straightforward solution to the problem that should be congenial to many. We make no purport here to respond to objections or to argue for our position; that will be the work of another day.

Why else is it silly to say "You ought to end world hunger," except that it is patently beyond your power? "Ought implies can" is most plausibly denied about some rare and particular situations, most notably, the tragic cases.²⁷ But what is a philosopher who says this really getting at? What is the force of calling an action wrong that everyone, *this philosopher included*, ultimately recommends? Even those who persevere in this approach do not hold that, because in a tragic case you must do wrong, you oughtn't do anything, nor that it doesn't matter what you do. They admit that the agent's moral obligation is to grit his teeth and choose the lesser evil. What is this insistence on wrongness, then, save a totemic application of a term of considerable emotive meaning, used to consecrate the misfortune of the circumstance?

Obviously, we don't find this line of thought congenial. But neither do we think that the philosophers who make much of tragic cases are misguided. They just misdescribe their insight. It seems more appropriate to us to characterize the genuine pathos of tragic cases in terms of it making sense to feel guilty, though you've done nothing wrong. Ought does imply can, even where mere regret or compunction is an inadequate response. When Gibbard writes that "part of the problem may be whether a distinction between compunction and full-fledged guilt can be drawn," we would urge an emendation: it is not part of the problem of making sense of the "ought implies can" controversy but, rather, a problem specifically for norm expressivism's prospects of making sense of what is in fact a broader ethical phenomenon. Moreover, if only it is granted that guilt can appropriately be felt over actions that aren't wrong, several of the intuitions that challenge the maxim can be brought into harmony with it.

So far, we've argued that guilt is warranted, and urged that nevertheless the action need not be wrong. But the norm expressivist analysis is conjunctive: to judge an act wrong is to think *guilt and anger* warranted. This suggests that Gibbard can grant both our claims, so long as he can persuade us that anger is not warranted. (Actually, as we'll see presently, the claim must be about a peculiarly moral kind

27. There is also an argument that denies the platitude in virtue of cases where an agent has himself brought it about that an (erstwhile?) obligation of his cannot be met. Compare Michael Stocker, "'Ought' and 'Can,'" *Australasian Journal of Philosophy* 49 (1971): 303–16. But this seems to us a minor problem, which suggests at most that ought implies could.

of anger.) In fact, though, such a response would be very uncomfortable for Gibbard. Guilt and anger, recall, are supposed to be reciprocal emotions, in that they are governed by the same norms.²⁸ To forsake this reciprocity would jeopardize the adaptive syndrome theory of guilt as a mechanism of reconciliation.

Worse yet, decoupling the warrant of guilt and anger robs the analysis of what strikes us as one of its most attractive features: its naturalistic account of the relation between the recognition that one is under a moral requirement and the recognition that one has reason to act appropriately. To think that society has a legitimate grievance against me is to think moral anger warranted. If in granting that, I must also think guilt on my part is warranted, then the very same state of norm acceptance that constitutes recognizing the grievance would give me some tendency actually to feel guilty. That's because our judgments about what it makes sense to feel influence the ways we actually do feel. Since guilt functions to adjust my motives, in feeling guilty I will also have some tendency to be moved to behave differently. Thus the reciprocal relationship between guilt and anger, and the insistence that they are governed by identical norms, form the cornerstone of a compelling account of how moral requirements generate reasons and motives for those who fall under them.

Of course, our scenario shows at most that guilt can be warranted when moral anger isn't, not vice versa. And it's the possibility that one could grant moral anger apt, and yet deny the appropriateness of guilt, that threatens the norm expressivist picture of the relation between obligation and motivation. However, if the emotions are not reciprocal, and yet guilt is warranted whenever moral anger is, then it's puzzling why the analysis should be conjunctive. Moreover, the success of the theory will then rest entirely on whether this promissory note for a nonjudgmental account of "moral anger" can be cashed. (We'll argue that it cannot.) Hence, though we'll now consider the possibility of rescuing norm expressivism with the claim that, although guilt makes sense here, the appropriate sort of anger does not, we doubt this suggestion would be attractive to Gibbard. If our arguments to this point are conceded, the theory is already in jeopardy.

ANGER AND IMPARTIALITY

It is easy to think mother's anger appropriate. She sacrificed much for you, and wanted very badly not to spend her last years in an institution. You might try to persuade her that the considerations guiding you override her wishes, but there is no particular reason to suppose that this would assuage her anger, nor even make her think

28. Gibbard, p. 139.

it unwarranted. Of course, the primary question is about you: whether you think mother's anger warranted. But it seems quite possible that you do. The better you can empathize with her, the more likely you'll find yourself accepting norms that call for anger from her position. This possibility, at any rate, is surely intelligible. You could think, then, that both guilt on your part and anger on hers make sense and yet judge that you acted rightly. Does this doom the norm expressivist analysis?

Not yet. The proper question is not about mother's anger specifically, but whether anger on the part of "others"—specified generally—is warranted. So what about others? Consider first mother's bridge club, her closest friends, who are also dependent on the care of their families. They too are angry, as it were "on her part." As we know, anger is elicited by perceived slights—not only those of moral import—typically, though not always, slights against ourselves or people close to us. And we know too that people have soft spots, sensitivities that when irritated provoke anger more readily than greater affronts to better protected areas of the psyche. Our disposition to anger, like so much of our personality, reveals our anxieties. In this sense we're peculiarly partial in anger. Whether we actually get angry, and how angry we get, depend strongly on what the situation is *to us*—both depend on aspects that seem morally irrelevant. This is so even when the slight isn't against us; it still makes a difference where we project ourselves imaginatively into the scenario. Who is slighted? And what are they—even symbolically—to you? Of course, you can attempt to correct for the idiosyncrasies of your responses; that is one of the functions of normative discourse. But what you do feel exerts a powerful, albeit resistible, influence on what you think you should—on this we agree with Gibbard. Indeed our actual dispositions to feeling provide our firmest handle on what we think about its warrant. Let's explore some things one might think about anger here.

Mother's friends in the bridge club are angry, but they differ in the norms they accept. Aubron thinks his anger unwarranted, because he considers that were he in your place, he would have done the same. Bertram thinks that only the anger of certain participants is warranted. It makes sense for mother to be angry, and her friends as well, because of their position. Bertram's thought that it would not make sense to be angry, were he not dependent, thus affords no pressure of inconsistency, since his participatory status would then be different. When Bertram examines his rationale, he finds no reason that he deems impartially acceptable. We sometimes acknowledge that our reasons are partial, and abjure trying to convince others to be angry even when we feel our own anger justified, because we don't think we can offer them reasons they should accept.

But there is another attitude one might take toward the scenario. Cecilia judges anger warranted more generally. The norms she accepts say it makes sense for impartial others to be angry. Only Cecilia makes a moral judgment, according to norm expressivism. And isn't that just right? Bertram's admission that it makes no sense for impartial others to be angry with you seems tantamount to denying your moral culpability. But he did think some anger warranted. Thus not all norms for guilt and anger are moral norms. In order to tease rational anger into the shape of morality, norm expressivism needs to circumscribe it carefully; to start, the theory must restrict itself to impartial anger.

Morality is impartial, and its norms are norms for what it makes sense to feel from an impartial standpoint. In our scenario, we can think guilt and some anger warranted even while we judge no wrong to have been done. After all, it would be inappropriate for outsiders to be angry with you; they would be resenting you not merely for doing something permissible but for failing to do wrong, by the very nature of the case. Hence we're prepared to grant—now that we've seen what kind of anger is relevant—that our scenario does not warrant such anger. So our case is not exactly a counterexample (as we previously warned it wouldn't be). However, we're now in a position to see that, in attempting to circumscribe anger to fit morality, norm expressivism actually runs afoul of a host of cases. That's because even our norms for the rationality of impartial anger do not seem to recapitulate moral judgment. The problem is that there's a tension between norm expressivism and the verdicts of commonsense morality, which the theory needs to alleviate.

The tension is rooted in a fact we've noticed about anger: what you actually feel depends very strongly on where you're placed. Most of us don't often find ourselves getting angry at injustices that are unrelated to our own concerns, even when we grant them to be worse, morally, than the local transgressions that sometimes so enrage us. Hume noticed this phenomenon and treated it as a potential objection to his sentimentalist account of moral judgment. Our moral *sentiments*, as he puts it, inevitably "vary according to the distance or contiguity of the objects," even when our moral *judgments* do not. The tension persists when we consider what it's appropriate for impartial others to feel about an acknowledged wrong. Surely it makes no sense for you to be angry about a theft happening halfway around the world; that would be absurd. Gibbard agrees that "how angry it makes sense for me to be depends on what the theft has to do with me."²⁹ But this does not mean that what happens far away is of no moral import. On the contrary, thefts are equally wrong no matter where they occur.

29. *Ibid.*, p. 126.

Thus the theory needs a final wrinkle. Moral norms are norms for what it makes sense to feel from what Gibbard calls an *impartial but fully engaged* standpoint. Though it makes no sense to get angry at the distant theft, anger would be warranted were we fully engaged. But while impartiality has a widely recognized (though not wholly uncontroversial) place in moral reflection, full engagement does not. Is there a plausible motivation for this constraint on which judgments are to count as moral? In order to recapitulate commonsense morality, Gibbard seems to invoke a peculiar standpoint with no precedent as a criterion of moral judgment.

In fact, though, Gibbard's solution is similar to Hume's own. Hume suggests that in order to correct for the effects of proximity on our responses and to prevent what he calls "continual contradictions," in moral reflection "we fix on some steady and general points of view; and always, in our thoughts, place ourselves in them, whatever may be our present situation."³⁰ Thus, for example, though a beautiful face may not provide so much pleasure when seen at a distance, Hume says we "correct its momentary appearance" and judge it no less beautiful, because we know what effects it would have on us if brought closer. He goes on to observe that "such corrections are common with regard to all the senses." Concerning matters of empirical fact, it's a straightforward observation that certain conditions are favorable for making judgments: these are just the conditions under which they're most likely to be true. In his talk of corrections, and his analogy to conditions for good sensory observation, Hume seems to suggest that there is some position from which our sentiments are more likely to get things right.³¹ This is a compelling view for those who want to combine cognitivism and sentimentalism; it allows us to say that moral judgments are grounded in our human dispositions, and yet to maintain their objectivity.³²

Of course, Hume's metaethical suggestions are notoriously in tension with one another; there are other places where he seems to deny cognitivism. Most modern commentators prudently shrink from trying

30. David Hume, *A Treatise of Human Nature*, ed. L. A. Selby-Bigge (Oxford: Oxford University Press, 1978), pp. 581–82 ff.

31. There are other ways to read these passages, of course—but this isn't Hume exegesis. One could understand Hume's version of the engagement requirement, like Gibbard's, as a restriction on the attitudes that count as full-fledged evaluative judgments. To judge an act wrong is to judge that it calls for disapprobation from those close enough to see it in all its hideousness. That seems plausible enough, but now we're back to expressivism's original difficulty: what is "disapprobation" but a promissory note?

32. Wiggins wants to do just that, and finds precedent in passages like these from Hume. Compare "Categorical Requirements: Kant and Hume on the Idea of Duty," *Monist* 74 (1991): 83–106, and "A Sensible Subjectivism."

to impose a univocal interpretation on these passages.³³ Gibbard, though, seems to be borrowing a page from the cognitivist Hume in service of an avowedly noncognitivist metaethical program. In taking up this Humean suggestion, does Gibbard too employ cognitivist intuitions that conflict with other aspects of his theory?

There is no straightforward inconsistency, because norm expressivism has different aspirations for its privileged standpoint than do ideal observer theories and other cognitivism. Cognitivists want to supply conditions under which moral judgment is (either likely to be, or stipulated as being) correct. Expressivists, though, seek an account of what makes a judgment a moral judgment at all. Norm expressivism's unique contribution is its association of moral judgment with the acceptance of norms for certain familiar emotions, particularly anger. So the question for its privileged standpoint—full impartial engagement—is whether we think anger from this standpoint is warranted in all and only those cases in which we judge an act wrong.³⁴

If norm expressivism is to retain the rationalism that is central to moral discourse, it must restrict itself to norms for what to feel from positions that justify, rather than merely explain, our responses. Accordingly, Gibbard defines full engagement as “vivid awareness of everything generic that would affect one's feelings toward a situation,” combined with “undistracted contemplative leisure.”³⁵ These conditions may indeed seem to justify the feelings they give rise to—but do they give rise to the right feelings for norm expressivism? Consider how you actually respond to a moral transgression when you're as close as you come to adopting the privileged standpoint. An appropriate test case would be a wrong that is not at your own expense (so you're impartial), but that you contemplate vividly, considering everything generic that would affect your feelings (you're engaged). Moreover, we need to think about a broad range of such cases, since certain transgressions are bound to affect us peculiarly, due to their special salience to us. What response do you find yourself having when you reflect on wrongs of this sort?

If you're like us, you'll find that this question has no univocal answer. There is some negative attitude common to the cases, but it seems too cognitive to help the expressivist, too close to a judgment itself—something we're tempted to call moral disapproval. In judging an act wrong we surely commit ourselves to thinking some such response appropriate, but need it be anger? Gibbard himself here slips

33. Compare, e.g., Jonathan Harrison, *Hume's Moral Epistemology* (Oxford: Clarendon Press, 1976), p. 110; and Barry Stroud, *Hume* (Boston: Routledge & Kegan Paul, 1977), p. 265, n. 9.

34. Not just when we correctly judge an act wrong.

35. Gibbard, p. 127.

into calling for outrage, an attitude that seems to have more intimate cognitive ties than does anger. But to speak of outrage or indignation would be to smuggle judgment back into emotion. Hume's discussion, by contrast, did not focus on anger at all, but on "approbation" and "disapprobation." This gains in plausibility; it's more tenable to say that to think an act wrong is to think it would make sense to *disapprove* of it from an impartial, informed position. But this gain is expressivism's loss, for disapproval is already a moral or quasi-moral notion.

There need be no such similarity between our less judgmental emotional responses to wrongs. Certainly we get angry, sometimes, but we're also moved to horror, grief, curiosity, astonishment. Nor are we inclined to think that one feeling rather than another is particularly apt. When we and a hellfire moralist are once again confronted with the evil humans commit, is his anger a more appropriate response than our awe?³⁶ As you pore with morbid fascination over the grim details of the latest serial killing, gaining vivid awareness of anything generic that might affect your feelings, including the familiar deprivations and cruelties that the killer himself suffered, do you find yourself thinking anger peculiarly warranted?³⁷

For our part, we find that the further we get toward actually adopting the imagined standpoint of full impartial engagement, the more our feelings about these acts tend toward the more existential emotions—horror, resignation, angst—and away from the more participatory, anger. Thus norms for the warrant of anger from expressivism's privileged standpoint fail to mimic moral judgment. Nevertheless, a good author—or a shrewd lawyer—can rouse us to righteous anger at an act that initially seemed unconnected to us and our concerns. But this is an art; its techniques are subtle and involve more than an inclusive recitation of the facts. The narrative, far from including

36. You might try saying that anger does not have to be the most appropriate response, but merely not inappropriate. We think it can even be that—anger in these situations can seem self-righteous and officious. Granted, it is never wholly misplaced as a response to wrongdoing, but that much is guaranteed by the intimacy between "anger"—the ordinary language concept—and moral judgment, an intimacy which expressivism must forbear trading upon. (This intimacy can be traced from "anger" to "outrage" to "moral indignation.") If the theory has something new to tell us, it must be that anger—i.e., the mechanism identified by the adaptive syndrome theory—is a peculiarly appropriate response to wrongs.

37. In answering this, we need to guard against information that would diminish culpability, for in such cases wrongness and blameworthiness come apart, and the analysis is more appropriate for the latter judgment. (Compare Gibbard, pp. 42–47.) Yet we propose that people's dispositions to anger at mass murderers (and to thinking anger at them apt) have less to do with their judgments of culpability than they do with more personal and idiosyncratic factors. These responses flow from deep and contingent aspects of our character, and they are often responsive to morally irrelevant details of the case.

"everything generic that would affect our feelings," is orchestrated to encourage the audience to project itself imaginatively and partially into a certain participatory role.

Our tendencies to get angry and to think anger warranted are insufficiently dependent on the facts to recapitulate morality. This is not a general thesis about the emotions, however; the warrant of some other negative emotions may be fixed by our factual judgments in just the way the norm expressivist would like anger's to be. Fear is perhaps most directly responsive to the circumstances; its rationality seems strictly proportional to the rationality of an associated belief about danger. Pity is sensitive to the judgment that someone's situation is wretched. Even the city dweller who knows it would be debilitating actually to pity very broadly recognizes that it would be apt from the standpoint of full impartial engagement. After all, were you to consider vividly the circumstances of so many people's lives, it would be hard to resist. More information could extinguish pity, it seems, only in suspect ways. We might become exhausted and unable to feel. Or we might find ourselves so appalled by someone's character as to lose our good will toward him. But doesn't his baseness make him all the more wretched, and hence, pitiful? To think so is to think pity less partial and participatory an emotion than is anger.³⁸

Despite its unruliness, anger is the only emotion (other than guilt) that is a suitable candidate for an expressivist analysis centering on wrongness. And as we've seen, norm expressivism must tie morality to the norms we accept for anger from the impartial and fully engaged standpoint. But what is the force of the claim that someone who does not think it makes sense, even from that standpoint, to be angry at all moral transgressions—but merely to discourage and punish, to disapprove of them—is not making moral judgments? This person can grant that anger sometimes makes sense and sometimes doesn't; he acknowledges a close connection between anger and judgments of wrongness, which sometimes coincide in righteous anger; and he may disapprove of exactly the same actions as does anger's enthusiast. In general he may be, as we fancy ourselves, fully competent with the concepts of morality, and passably motivated by its requirements.

Norm expressivism makes righteousness too internal to morality. In accepting an identification of moral considerations with grounds

38. Aristotle, however, would have disagreed. He held that pity was pain at the vividly entertained thought of disaster for someone like oneself in power and susceptibility (see Aristotle, *Rhetoric*, bk. 2, chap. 8, ed. Richard McKeon [New York: Random House, 1941], p. 1396). On this view, vivid awareness of the wretchedness of someone's life gives no guarantee that it makes sense to pity him. We must also find him sufficiently familiar, which more awareness might vitiate. Is Aristotle's view a piece of psychological realism or another instance of antiquity's ruthlessness to the weak?

for anger in particular, we lose some of our grip on the kind of reason that moral reasons are. They emerge bearing too much of the flavor of a particular emotional disposition, albeit one to which a common brand of moralist has been prone (as did not escape Nietzsche).

Simon Blackburn writes, in his justly laudatory review of Gibbard's book, that the complaint of an impassable "Fregean gulf" between moral disapproval and any conative state is not available to any philosopher of naturalist tendencies.³⁹ While we agree that the existence of such a gap must not be claimed to be discoverable a priori, it's still a difficult matter to find a state that can do justice to the motivational component of moral judgment while also exhibiting the discipline of moral discourse. The central naturalist thought in ethics is that the normative force of values must ultimately be located in motivational force, if it's not to be convicted of systematic reification error. The puzzle, of course, is how then to make sense of the seeming lack of contingency upon motive that moral judgments possess.

Descriptivist theories have certain costs; they inevitably close the question of whether to endorse the values held from whatever privileged epistemic position the theory seizes upon. Expressivism captures normativity through the motivational force expressed in judgment. It too must identify a conative state as of the essence of valuing. Gibbard's singular contribution is to adduce states of norm acceptance, to account for our complex justificatory lore—but just what norms are expressed in moral judgment? The most likely candidate seemed to be norms for the warrant of emotions. And Gibbard's norm expressivism is right to fix upon anger and guilt as singularly apt to capture at any rate our culture's morality. Norm expressivism would tie the sphere of moral reasons directly to the emotions of guilt and anger: reasons to judge an act wrong are reasons to think guilt and (a particular kind of) anger warranted. Moral reasons are tied only indirectly, through the motivational manifestations of these emotions, to what to do. We think an indirect connection is just right; nonmoral considerations also bear on what we have most reason to do.

However, the analysis is forced into some crucial contortions to avoid conflicting with ordinary thought. In its reliance on fine-grained distinctions between guilt and regret, and on a special standpoint for anger, norm expressivism trades on features of commonsense lore that belie its noncognitive theory of emotion. While the emotions are unquestionably a rich resource for moral psychology, they are recalcitrant to theoretical treatment and require case-by-case investigation. Gibbard's work marks a significant advance in its sensitive treatment

39. Simon Blackburn, "Wise Feelings, Apt Reading," *Ethics* 102 (1992): 342–56.

of relations between emotion, warrant, and belief. Were guilt and anger to rest on a solid middle ground between representation and undisciplined feeling, the norm expressivist approach would succeed brilliantly. If, though, as we've argued, these nearly moral emotions have one foot uneasily in each camp, then they will not be able to sustain what is needed: a rationalism without cognitivism.