

## Fanaticism and Hare's Moral Theory

Robert K. Fullinwider

*Ethics*, Vol. 87, No. 2 (Jan., 1977), 165-173.

Stable URL:

<http://links.jstor.org/sici?sici=0014-1704%28197701%2987%3A2%3C165%3AFAHMT%3E2.0.CO%3B2-2>

*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 creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).



Robert K. Fullinwider

State University of New York at Albany

In this paper I take up the persistent problem posed for R. M. Hare's moral theory by the existence of the "fanatic." I chronicle Hare's various attempts to solve this problem. My motive for focusing on the fanatic is the light to be shed on two of Hare's theses—that an adequate moral theory can be extracted from the logical features of moral language alone and that this moral theory is utilitarian in character. Tracing Hare's treatment of the fanatic since *Freedom and Reason (FR)*<sup>1</sup> shows how Hare has added to his initial formalist description of his theory the further view that it is essentially utilitarian. I will argue that Hare does not successfully meet the problem of the fanatic, either in terms of the formalist elements of his theory or in terms of its utilitarian elements. And I will show that Hare does not successfully connect the two: he fails to establish that utilitarianism is entailed by the logic of moral language.

I

In order to describe the problem of the fanatic and to follow the ins and outs of Hare's various attempts at solving the problem, it will be necessary to say something about his theory of moral reasoning. Hare believes that, on the basis of the logical properties of moral language, a strong theory of moral reasoning emerges, one that is adequate to deal with real moral problems. Recently Hare has come to claim that these formal properties of moral language entail a substantive moral view—utilitarianism.<sup>2</sup> The features of moral language that underlie Hare's theory are the prescriptivity of moral judgments and the universalizability of moral reasons. Together these features require that when an agent (morally) justifies his doing *A* for reason *R* he must concede that, for anyone who has reason *R* to do *A*, that person's doing *A* is also, to that extent, justified. (Call this the universalizability condition.) The fanatic arises as an embarrassment to Hare's

1. R. M. Hare, *Freedom and Reason* (New York: Oxford University Press, 1965); hereafter cited as *FR*. Throughout the paper, numbers in parentheses refer to pages in the work mentioned.

2. See Hare's recent criticism of Rawls: "Critical Study: Rawls' Theory of Justice," pts. 1 and 2, *Philosophical Quarterly* 23 (April and July 1973): 144–55, 241–52, esp. 154–55; and "Ethical Theory and Utilitarianism" (paper read to Creighton Club, New York, October 1974 and forthcoming in H. D. Lewis, ed., *Contemporary British Philosophy*, IV, esp. pp. 10, 15; hereafter cited as *ETU*).

theory because it appears that he can maintain his fanatical position without logical error; consequently, the theory would seem to lack rational power to attack the fanatic. Such impotence is a serious liability for any moral theory that claims adequacy.

The fanatic is one who, according to Hare, is willing to sacrifice the interests of others, even his own mundane interests, in order to realize some ideal. In one of Hare's examples the fanatical Nazi, in support of his ideal of racial purity, is willing to demand the extermination of Jews. If he can do this without logical inconsistency, is his position not, from the perspective of Hare's theory, theoretically impregnable? Does a moral theory which appeals solely to the logic of moral language have resources for offering a moral reason for the consistent Nazi not to act in furtherance of his ideal?

In expounding his theory, Hare in effect describes three tests the moral reasoner must meet. He believes the three tests are required by his account of moral language. Test 1 I shall call the "reversal test." In discussing the fanatic in *FR*, Hare requires that the Nazi answer the question, "What if I were a Jew, and this killing I wish for Jews were to happen to me?"<sup>3</sup> If he prescribes death for all Jews, but hesitates to prescribe his own death were he to turn out to be a Jew himself, he fails the universalizability condition and his position is indefensibly inconsistent. However, if the fanatical Nazi, in thrall to his ideal, willingly demands his own death should he turn out to be a Jew, his position meets the universalizability condition and his position is consistent.

If any of the three tests is required by the universalizability condition, then the reversal test seems most surely so, since it enforces consistency in judgment. To meet the universalizability requirement (and the reversal test), the Nazi must adopt a consistent stance, *either* by giving up his prescription to kill Jews *or* by overriding his desire not to be killed if he is a Jew. The reversal test does not determine a particular moral stand but only *sets* of consistent stands. This is what we should expect from a formal test.

Test 2 I call the "exchange test." This requires not that the Nazi imagine what he plans for the Jew happening to him but that he imagine himself as the Jew with his (the Jew's) desires and interests.<sup>4</sup> Whereas the reversal test (test 1) was a test without a particular outcome, the exchange test, by itself, seems to yield no outcome at all. The Jew has no desire to see Jews exterminated, so when the Nazi subjects himself to the exchange test he imagines himself not having a desire that he in fact has. This imagining lacks instructive value for any decision about what he ought or ought not to do. To imagine oneself as not having a desire does not tell one whether it is permissible to act on that desire if one has it.

In fact we must see test 2 not as an independent test, like test 1, but as an element in test 3: the "ideal observer test." Suppose an act I plan will have

3. *FR*, p. 130. See also "The Argument from Received Opinion," pp. 117-35 ("Would you want that trick played on you?" [p. 107]), and "The Practical Relevance of Philosophy," pp. 98-116, both in R. M. Hare, *Essays on Philosophical Method* (London: Macmillan Co., 1971).

4. *FR*, chap. 7, esp. p. 113, and *ETU*, p. 7.

5. See *FR*, pp. 115-24, and *ETU*, p. 3.

adverse effects on X, Y, and Z. I put myself in the place of X, and imagine having his desire against the planned act (test 2); and similarly for Y and Z. By then taking up the role of ideal observer, *I imagine myself as an individual possessing all the desires* regarding the planned act—my own original positive desire and the negative desires of X, Y, and Z. In so imagining, “I” will choose to forego the planned act because doing so is in accord with “my” strongest desire.

The ideal observer seeks rationally to satisfy his strongest preferences. Since his preferences are the sympathetically imagined preferences of many individual persons, Hare’s third test amounts to deciding to act on the basis of a utilitarian summing of positive and negative desires. (I leave aside until Section III, below, consideration of how Hare thinks this substantive test emerges from test 1 or the universalizability-prescriptivity features of moral language.) In taking up the reversal test, the Nazi fanatic had to weigh his desire for a pure race against his desire to stay alive (were he a Jew). In applying the third test and taking up the standpoint of the ideal observer, the Nazi fanatic must weigh his desire for a pure race against the desires of all threatened Jews to stay alive.

## II

At the beginning of chapter 9 in *FR*, Hare sees clearly the potential damage to his theory caused by the impregnability of the fanatic. He begins by declaring that “it would indeed be a *scandal* if no arguments could be brought against a person who, in pursuit of his own ideals, trampled ruthlessly on other people’s interests . . .” (*FR*, 157; my emphasis). Yet, given this clear recognition that a scandal looms for his theory if it cannot muster arguments against the fanatic, Hare’s ensuing solution to the problem of the fanatic is no solution at all. After showing that there are arguments against inconsistent fanatics, his solution consists in *admitting* that the Nazi fanatic who is willing to universalize his prescriptions is *immune* to telling arguments against him. Fortunately, Hare tells us, “really intractable Nazis are perhaps rarer than might be thought” (*FR*, 171).

Allowing that there are fanatics who can pass the reversal test and then suggesting that, fortunately, there are few of them, in effect concedes the scandal. The blatant unsatisfactoriness of the “solution” in *FR* has led Hare to return to this issue several times in the past decade. In “Peace” (*P*, 1966),<sup>6</sup> he again recognizes that a fanatic can apparently pass the reversal test, holding a position which is “unassailable in argument” (*P*, 80, 82; see also 78–79). Again, Hare seeks to undercut the force of this fact by arguing that truly universalizing fanatics are rare (*P*, 80). Most actual fanatics are not clear headed, not fully informed, not factually correct, and so on. *They* and their positions are not “unassailable in argument.” True fanatics who are clear headed and fully informed—these are the fanatics that are “unassailable in argument,” and they are rare.

Relying upon the resources provided by the logic of moral language (i.e., the reversal test), this appears to be the best solution available to Hare. Yet it is an

6. In R. M. Hare, *Applications of Moral Philosophy* (London: Macmillan Co., 1972), pp. 71–89. Dates given in parentheses in text refer to the date of first publication or reading.

inadequate solution in two ways. First, to establish the rarity of the true fanatic Hare has to rely upon the dubious empirical claim that most fanatics rest their ideals upon false beliefs and confused thought. But, more importantly, the existence of *any* Nazi Jew killer "unassailable in argument" is a scandal to that mode of argument. Hare needs to remove the fanatic from the scene altogether. Yet the reversal test is clearly inadequate to this task. In subsequent writings, Hare increasingly relies upon the exchange test and the ideal observer test to dispatch the fanatic. The utilitarian elements of the theory in *FR* are moved to the forefront.

This is noticeable in "Reply to Katz" (*RK*, 1970),<sup>7</sup> where Hare claims to have an argument that will demonstrate "that in order to be a fanatic in the sense required . . . a man would have to have properties which are really *practically impossible*" (my emphasis). Though the logic of moral language and universalizability are the starting points in this essay, when Hare takes up the fanatic it is the ideal observer test that plays the crucial role. The fanatic must put himself "in the situation of his victim with his victim's desires" (exchange test) and, as ideal observer, act on the strongest (summed) desire (*RK*, 50). In order for the fanatic to meet this test

the Nazi would have to have a desire to be rid of the Jew greater, not than the Nazi's own desire not to be put in a gas chamber, but than the Jew's desire not to be put in a gas chamber. Now if we suppose that there are many more Jews than Fanatical Nazis, it becomes practically impossible that the Nazis should have such strong desires to be rid of Jews that they, in sum, outweigh the desires of the Jews not to be got rid of. By applying this utilitarian-style felicific calculus to my theory I can thus show the practical impossibility of there being any real fanatics. [*RK*, 51-52; my emphasis]

The explicit use of the "utilitarian-style felicific calculus" to solve the problem of the fanatic reveals the growing tendency of Hare to identify his theory with utilitarianism, a tendency which is even more explicit in later essays. (Here a question arises about Hare's applying the felicific calculus to his theory: Is the application gratuitous, an ad hoc way of getting rid of the problem of the fanatic, or is the application generated by, and organic to, his theory of moral reasoning? An answer to this question about the relation of the ideal observer test to the reversal test is postponed until the next section.)

Hare repeats essentially the same argument in "Wrongness and Harm" (*WH*, 1972).<sup>8</sup> The fanatic, in order to get through Hare's net, will have to be a person who possesses a desire of (impossibly) enormous intensity. Furthermore, Hare admits, "It is obvious that both the utilitarian and I will have to say that such a heroic desire, if it occurred, *ought to be satisfied*" (*WH*, 107; my emphasis). This admission is particularly illuminating of the way Hare, by this time, has come to view his theory of moral reasoning. In *FR*, where the focus is on universalizability, test 1 plays the major role. Although Hare there notes connections between his theory and utilitarianism, he does not equate the two. Test 1 would never compel Hare to admit that the fanatic ought to have his desire satisfied, whatever the

7. In B. Y. Khanbali et al. (eds.), *Jowett Papers 1968-1969* (Oxford: Basil Blackwell, 1970), pp. 44-52.

8. In R. M. Hare, *Essays on the Moral Concepts* (Berkeley: University of California Press, 1972), pp. 92-109.

intensity of it. In *FR*, Hare is quite prepared to endorse and stand by the liberal principle of tolerance, a principle inconsistent with the fanatic's unchecked zeal. The fanatic's position may be "unassailable in argument" from the perspective of test 1 but, by the same token, so may Hare's opposition to the fanatic. In *FR*, although Hare's position might not enable him to argue effectively *against* the fanatic, it certainly does not require him to argue *for* the fanatic under any circumstances!

By 1972 Hare has come to see his theory and utilitarianism as practically equivalent. Consequently, test 3 (with test 2) plays the leading role in Hare's later discussions of the fanatic. Whereas test 1 identifies consistent sets of moral beliefs, the application of test 3 yields a particular and determinate moral choice for any given set of preferences. Utilitarianism always requires the outcome that maximizes net satisfaction of desires. Hare, in coming to view his moral theory as utilitarian, has moved to a position where he must positively endorse the Nazi's actions under certain circumstances. It becomes even more crucial, as a result, that Hare show that these circumstances can never obtain. Otherwise, Hare's theory is indeed vulnerable to scandal.

The argument in *RK* purports to demonstrate the practical impossibility of the requisite circumstances. It succeeds if we accept its supposition. But need we? The supposition is that the fanatics are outnumbered by their intended victims. Surely we need not suppose this state of affairs to be always the case. Hare recognizes this problem in *WH*, where he admits that he is in trouble "if, although no single person has a desire of such [practically impossible] strength, there is a large number of people whose desires, though not so great severally, add up in total to an intensity which is sufficient to outweigh all the other desires of themselves and other people" (*WH*, 107). This is indeed a possibility that will put Hare in trouble. It undermines his argument that the true fanatic is a practical impossibility. Assuming that he belongs to a large group of like-minded fanatics with a sufficiently limited class of victims, the fanatic whose desires in combination with the desires of those like him, outweigh the combined desires of the victims, is a very real possibility. Moreover, in such circumstances, when the fanatic takes up the perspective of the ideal observer and sympathetically imagines all the desires of all the persons involved as his own desires, his "strongest" (summed) desire will be to kill the victims; and he thus finds justification for his aims. How does Hare respond to this?

"If I may speak for myself," he says, "my confidence that real-life situations of this sort will not occur is based on my factual beliefs about what the world, and the people in it, are like" (*WH*, 107-8). Thus, the supposition upon which Hare's entire argument rests is supported only by Hare's reading of the world and its people. This is unconvincing and unacceptable as a defense of such an important, and artificial, supposition. In effect, Hare's move here amounts to a return to the earlier position of *FR* and *P*. Normal people are not like Nazis; fanatics will be rare; it is unlikely they should outnumber their intended victims. We have already seen that it is not enough for Hare to show that fanatics of the requisite sort are rare, or that large numbers of fanatics are unlikely. This is especially the case in light of the possibility that the utilitarian (Hare) must *endorse* the fanatics' plans

in circumstances of precisely the sort where the combined weight of their desires (due to the ratio of fanatics to victims) outweighs the desires of their victims. Hare needs to establish that these circumstances *never* obtain.

In "Ethical Theory and Utilitarianism" (*ETU*, 1974), Hare adopts a new-appearing strategem for dealing with the fanatic. The existence of the universalizing fanatic poses a counterexample to claims of adequacy for his theory. If certain kinds of counterexamples could be ruled out of bounds on procedural grounds and the case of the fanatic is among these kinds, then Hare can avoid scandal for his theory since it will not even have to face the fanatic's challenge. The relevant sort of counterexample Hare proposes to rule out of bounds is the "fantastic" case.<sup>9</sup> Hare notes that "the only sort of fanatic that is going to bother us is the person whose ideals are so intensely pursued that the weight that has to be given them, considered impartially, outbalances the combined weights of all the ideals, desires, likings, etc., that have to be frustrated in order to achieve them" (*ETU*, 8). To entertain such a possibility is to hypothesize a "fantastic" case. I take it that by calling it fantastic Hare means to emphasize that no human being, as we know them, could have a desire so intense, except in fantasy. In other words, Hare is making in a new guise the same point that he inaugurated in *RK*—that a person possessing such an intense desire is a practical impossibility.

But, calling the case of the fanatic a fantastic case depends, then, upon reinstating the already discredited assumption that there will always be few fanatics and many victims. Thus, despite Hare's confidence that in real life fanatics will never sufficiently outnumber their victims, it is not fantastic to suppose such a case, and Hare's theory cannot avoid dealing with one. Although Hare prefaces his discussion of the fanatic in *ETU* with the remark that "my present formulation enables me to deal in an agreeably clear way with the problem of the fanatic, who has given me so much trouble in the past," this essay, like the others, fails to move the problem of the fanatic beyond the unsatisfactory solution in *FR*. As of 1974, Hare's troubles with the fanatic are still not behind him.

### III

The fanatic can meet test 1. Hare recognizes this. He has tried to show that the demands of test 3 are so stiff that there could be no actual true fanatic or group of true fanatics. He has failed to show this. In the right circumstances the fanatic can pass the ideal observer test as well as the reversal test. And the right circumstances are not fantastic.

9. Hare makes a distinction in *ETU* (p. 9) between level 1 principles and level 2 principles of moral reasoning. "Level-1 principles are for use in practical moral thinking, especially under conditions of stress. . . . Level-2 principles are what would be arrived at by leisured moral thought in completely adequate knowledge of the facts, as the right answer to specific cases." According to Hare, "fantastic" cases are directed against level one principles, and this is illegitimate since level one principles are not designed to handle such cases. Why does Hare believe that fantastic cases are being directed at level 1 principles rather than level 2 principles? The reason evidently is contained in this calumny: "the anti-utilitarians have usually confined their own thought about moral reasoning . . . to what I am calling level-1 . . ." (p. 10). For discussion of principles and fantastic cases in *FR*, see pp. 42 ff.

Even so, we might ask: Why must the fanatic submit to test 3 (with test 2)? To put the matter a different way, how does Hare justify the utilitarian component of his theory? This is an especially important question, since Hare claims to base his theory of moral reasoning "entirely on the formal properties of the moral concepts as revealed by the logical study of moral language; and in particular on the features of prescriptivity and universalizability . . ." (*ETU*, 3). If Hare's theory is based entirely upon the logic of moral language, then test 3 must be derivable from test 1, if test 1 indeed reflects the universalizability condition. Test 3 requires that a moral reasoner take into account *all* the desires of others that bear on his decision (via test 2) and that he maximize the satisfaction of the set of desires (including his own) taken into account. How can this be a requirement of logic or consistency?

In *ETU*, Hare takes the line that universalizability requires *impartiality*, and that impartiality and the principle of utility are practically equivalent. In morally prescribing in some situation, "I am prescribing universally for all situations just like the one I am considering, and thus for all such situations, *whatever* role, among those situations, I myself occupy. I shall therefore give equal weight to the equal interests of the occupants of all the roles in the situations; and since any of these occupants might be myself, this weight will be positive" (*ETU*, 4). Hare's conclusion in this passage does not follow. Why, if I use only the reversal test (reflecting the universalizability condition), will I give equal weight to all the interests? Moreover, even if I did, why must I go on to maximize the satisfaction of those interests (employ the ideal observer test)?

Consider X, who knows that he desires  $a, b, c$ , but does not know which of roles  $R_1 \dots R_n$  he will play in the situation he is imagining. In such circumstances, if he is self-interested he will not favor a principle like, "Let  $R_1$  be favored in satisfying  $a, b, c$ "; instead, X will subscribe to a principle like, "Give  $R_1 \dots R_n$  equal chances to satisfy  $a, b, c$ ." Here, we are viewing  $R_1 \dots R_n$  as open slots, to be filled by X with his desires and by other individuals with their desires. In the imagined situation, X does favor a principle which gives equal weight to equal interests—equal weight to interests in satisfying  $a, b, c$ . He does *not* take into account *all* interests that might be possessed by anyone holding a role in a situation such as he is prescribing for. In other words, X can prescribe universally without undertaking the operations of test 2 (which is a necessary preparation for test 3).

Thus, suppose Y and Z turn out to occupy  $R_1$  and  $R_2$  while X occupies  $R_3$ , and that Y and Z possess interests in  $d, e$ , and  $f$ . Without logical inconsistency, X gives weight to the interests of Y and Z only if they share interests which he weighs positively in his own preference set. In subscribing to the principle, "Give  $R_1 \dots R_n$  equal chances to satisfy  $a, b, c$ ," X meets the requirement of universalizability without being constrained to give weight to all the desires of other individuals (such as Y and Z) who may occupy a role in  $R_1 \dots R_n$ . The situation X is prescribing for is one of a person desiring  $a, b, c$ ; and he prescribes an equal chance to satisfy *these* desires for *any* role player; he subscribes to a principle which is *role neutral*, but not *interest neutral*. Hare is partly right in saying that universalizability requires impartiality. It requires the impartiality of role-neutral principles. But this

is not enough for Hare's position. Tests 2 and 3 conjointly require interest-neutral principles. There is a gap between the impartiality provided by universalizability and the impartiality needed by the ideal observer.

By allowing X to know what his desires are, we shall never get from him the interest impartiality needed by Hare. One way of getting the full impartiality Hare wants is to do as Rawls does and drop a veil of ignorance over X. We would accordingly view X as not knowing what his desires are, and knowing that he will play one role in  $R_1 \dots R_n$ , where attached to each role is a set of desires which will be X's desires should he play that role. However, Hare objects strongly to this device.<sup>10</sup> But even if Hare availed himself of the veil of ignorance, he would have to show how this is required by universalizability; that is, he would have to bridge the gap that has been shown to exist between role impartiality and test 3 and bridge it with resources derived solely from the logic of moral language. Can Hare gain the desired impartiality without resorting to a device like the veil of ignorance?

Suppose X has been instructed to reason thus: "I am to imagine myself with the preferences of X (who is in fact me); I am to imagine myself as Y with the preferences of Y; I am to imagine myself as Z with the preferences of Z; and so on" (test 2). In going through this supposed mental operation, X will be "impartial" with respect to all the desires of Y, Z, etc., in that he will imagine, in turn, having the preferences of Y, Z, etc., giving these imagined preferences the weight given them by their actual possessors. But how is this impartiality retained when X, on the basis of this imaginative operation, comes to make a choice of action (subscribe to a principle)? He is further instructed to take up the stance of ideal observer, and in taking up that stance X, *qua moral reasoner*, is to view X (himself), *qua interested party*, as merely one among those on the list X, Y, Z, etc.; X's knowledge of who he is on the list is irrelevant in his (as ideal observer) making a moral decision (*ETU*, 8). This, however, amounts to smuggling the veil of ignorance in by the back door; for although X is allowed to know his desires, as he would not behind the veil, this knowledge is not allowed to have an effect on his choice.<sup>11</sup> Rawls excludes X from knowledge of his interests. Hare gives X this knowledge, but excludes it from having any effect; it is as if he did not have the knowledge. Hare objects to Rawls's exclusion, but his own exclusion stands equally in need of justification.

To answer the question posed two paragraphs ago, Hare can gain the impartiality he needs only by imposing two instructions upon the moral reasoner. These instructions go beyond what is required by the reversal test; and Hare's belief that universalizability requires impartiality and that impartiality is reflected in the perspective and judgments of the ideal observer, and thus that universalizability

10. "Critical Study," pp. 151-52; *ETU*, p. 3.

11. As ideal observer, X is allowed to know X's preferences, just as he knows Y's and Z's. That knowledge of X's preferences has effect on the decision. What is excluded from having effect is X's knowledge that X's preferences are *his* preferences. So in saying that he allows X to know his own desires, in contrast to Rawls, Hare misleads us, and himself, into thinking he allows something which Rawls does not.

implies utilitarianism, rests upon confusing two kinds of impartiality. The two instructions that Hare must give the moral reasoner are not shown to be connected to the universalizability condition; and the reversal test is not shown to imply the ideal observer test or the exchange test. The instructions given by Hare are only two among many possible, logically consistent instructions. They are not mandated by the logic of moral language. From the point of view of the logic of moral language alone Hare's utilitarianism is gratuitous. The fanatic does not have to subject himself to test 3.

IV

The generalized condition expressed by the Nazi's behavior is "overriding the interests of others in pursuit of an ideal." This condition can occur in a variety of degrees and manners. Although Hare claims it is implausible that the Nazis' desires to kill Jews could ever outweigh the Jews' desires to live, he cannot seriously maintain that, say, a large majority's ideal of decency could never outweigh the desires of a limited minority to see pornographic movies. Neither the reversal test nor the ideal observer test stands insurmountably in the way of such a majority imposing its ideal at the costs of others' interests. But we would not call the majority fanatics. Moreover, nonutilitarian moral theories declare certain kinds of interests or desires to be themselves immoral and thus as requiring to be overridden or frustrated in the name of justice. Unless Hare proposes to designate the moral theories of, for example, Kant and Rawls, as instances of fanaticism, it is apparent that we are misled in being offered the Nazi monster as the representative challenge to Hare's theory. The Nazi is only the most acute challenge.

The enduring theme in Hare's moral philosophy is his belief that a strong system of moral argument can be constructed out of the logic of moral language. The case of the Nazi fanatic puts extreme pressure on Hare's system, and under the pressure Hare has come to modify his theory so as to turn it into a substantive moral theory. I have shown that there is an illegitimate step in Hare's effort to ratify this move and derive utilitarianism from universalizability. Hare expects too much from the logic of moral language, and under the pressure of this expectation has tried to squeeze content from form. This attempt is surprising from the foremost defender of the fact-value distinction in recent times.

We cannot expect to forge strong tools for dealing with moral issues by talking only *about* morality. Moral problems must be faced from *within* morality, that is, from within the framework of an articulated and defended substantive moral theory. Utilitarianism is one such theory. There are others. All are capable of consistent and coherent development as systems of moral thought. None, including utilitarianism, is entailed by the logic of moral language, but all can stand on that logic.