

## UTILITARIANISM AND CONFLICT RESOLUTION

Terrance C. McCONNELL

A *genuine moral dilemma* is a situation in which an agent ought to do each of two acts, both of which he cannot do. That is, he ought to do A and he ought to do B, but he cannot do both A and B. There can be no doubt that there are many situations that appear, at least at first glance, to be genuine moral dilemmas. Let us call these situations ought-conflicts or cases of moral conflict. Whether an adequate moral theory must rule out the possibility of genuine moral dilemmas (and hence provide, in principle, a solution to ought-conflicts) or whether an adequate moral theory must allow for their existence is a matter of controversy among ethical theorists. Some claim that if a theory allows for the possibility of genuine moral dilemmas then it is in some way incoherent; others argue that if a theory does not allow for the existence of genuine moral quandaries then it is unrealistic. One of the contexts in which this dispute has arisen concerns the assessment of a particular moral theory, viz., (act) utilitarianism. We may characterize utilitarianism as follows:

(U) An agent is morally obligated to do A in circumstances c if and only if his doing A in c will bring about a greater (at least as great a) balance of good over evil than (as) his doing any alternative action open to him in c.

So characterized, utilitarianism is a theory about what objectively ought to be done; that is, it is a theory about what an agent ought to do in circumstances as they actually are, not merely as he believes them to be. Those who hold that an adequate ethical theory must rule out the possibility of genuine dilemmas claim that according to this criterion, at least, utilitarianism is a strong moral theory. Critics have argued, however, that utilitarianism is an unrealistic theory because it does not allow for genuine moral dilemmas.<sup>(1)</sup> The utilitarian, in effect, is claiming that all values are commensurable, and hence there are in theory no irresolvable moral conflicts. By contrast, the oppo-

nents are asserting that there are in fact incommensurable values.

How in principle this dispute about the adequacy of utilitarianism is to be settled seems simple enough. What is needed is an examination of the question of moral dilemmas independent of any particular moral theory. If arguments can be advanced to show that an adequate moral theory should rule out the possibility of genuine moral dilemmas, then it would seem that the utilitarian's handling of this question is a definite strength of his theory. For the purposes of this paper, assume that such arguments can be given; that is, assume that genuine moral quandaries must (in theory) be eliminated.<sup>(2)</sup> Both utilitarians and their critics have thought that if this is granted, then utilitarianism will fare at least as well as, and probably better than, its deontological competitors. My purpose is to challenge this widespread assumption made *about* teleological (or consequentialist) theories in general, and utilitarianism in particular. I shall *not* argue that utilitarianism cannot in principle eliminate genuine moral dilemmas; rather I shall claim that the way such irresolvable conflicts are eliminated leads to *internal* problems for the utilitarian.

## I.

Assuming that an adequate moral theory must rule out the possibility of genuine moral dilemmas, let us see why it has been claimed that the utilitarian theory has a distinct advantage over its deontological alternatives. The utilitarian's argument is as follows.<sup>(3)</sup> Since utilitarianism is concerned with the maximization of a single moral value, there is no possibility of duties coming into conflict. But this feature is lacking in deontological theories. If a theory postulates a number of independent *prima facie* duties, such as the duty to tell the truth, the duty to keep one's promise, and the duty to prevent needless suffering, then it must acknowledge the possibility of such duties coming into conflict. For example, it may be that in a given situation the only way that one can prevent suffering is to break a promise. How can a deontological theory resolve such conflicts? One plausible criterion for resolving such conflicts is for one to do that act which has the best consequences in that situation. If there is no conflict, then one should do one's *prima facie* duty; but in cases of

conflict, one should appeal to the utilitarian principle. But this way of resolving ought-conflicts is not open to the deontologist, as the following line of reasoning shows. (A) If maximizing the good (or general benevolence) is among one's duties, then there will always be a conflict of duties when some other *prima facie* duty bids one to do something that will not maximize the good in a particular situation. And if the utilitarian principle is the criterion for resolving conflicts, then the proposed deontological theory can never yield a moral result contrary to utilitarianism and so is not a genuine alternative to it. (B) But if the theory denies a place to the duty of general benevolence, then it is implausible. Moreover, the problem of conflicts among the duties that it does recognize remains unsolved. (C) And if the theory does recognize a *prima facie* duty of benevolence but does not permit it always to take precedence over other duties in cases of conflict, then again the problem of resolving conflicts among *prima facie* duties remains. As Mill says of (non-utilitarian) moral theories positing independent moral duties, «[T]here is no common umpire entitled to interfere between them.»<sup>(4)</sup> To say, as Ross does, that one is to intuit which of the conflicting *prima facie* duties is one's actual duty is to provide no solution at all. So utilitarianism seems to have an advantage over deontological theories concerning the resolution of ought-conflicts.

This argument makes two important assumptions. *First*, it assumes that the deontologist has no plausible criterion other than the utilitarian one for resolving conflicts among moral rules. The inadequacy of Ross's discussion is often cited as evidence to support this assumption. *Second*, the argument assumes that any theory concerned with the maximization of a single moral value (i.e., any teleological theory) has an easy way of resolving moral conflicts. An agent is always to act on that alternative which maximizes the good, and if there are several *optimific* acts open to the agent, he ought to perform one of the acts (assuming that he can perform any one, but only one, of the *optimific* acts). This idea is set out in the following principle:

- (P) If several alternative actions open to person *p* are each maximal in utility, then *p* is obligated to do any one of them.

So while cases of conflicting moral rules present a difficult problem for a deontological moral theory, the analogues of these situations in a utilitarian schema are not only not a problem, but they are situations in which the utilitarian finds himself blessed with an abundance of riches. I shall deal only briefly with the first assumption of this argument. My main interest is to examine the plausibility of the second assumption, the one on which the utilitarian and his critics agree.

## II.

One common objection to utilitarianism, one aimed at the second assumption, should be dispensed with immediately. It is often claimed that it is impossible in situations of ordinary life to foresee all the consequences of an act. And in cases in which the consequences of two or more acts have nearly or exactly the same value, the doubt cannot be resolved at the time the decision must be made. The critic claims that in weighing good against good (or if only evil alternatives are open to the agent, evil against evil), the balance is so delicate that only future and unpredictable consequences will resolve the doubt. Since the resolution of these cases of moral conflict depends on future consequences that are not (at least reasonably) foreseeable, the utilitarian has no real solution to moral conflicts.<sup>(5)</sup> So the utilitarian does not have an easy way of resolving apparent moral dilemmas. While this objection is a common one, I believe that it is mistaken. When one asks what the purpose of an ethical theory is, there are at least two different answers.<sup>(6)</sup> One purpose of an ethical theory is to provide an account of obligation-making characteristics. That is, an ethical theory is supposed to provide an account of the characteristics which all and only morally required acts have by virtue of which they are morally required. It might be suggested, however, that another purpose of an ethical theory is to provide a procedure which, if followed, would provide one in practice with helpful or correct answers to particular moral questions. On this latter view an ethical theory is to provide an agent with a decision-making procedure. The criticism being considered here attacks utilitarianism because it does not provide a decision-making procedure for agents in situations of

moral conflict. Formulated in a strong way, this objection purports that utilitarianism makes moral knowledge in many (or all) situations impossible. Let us assume that this is a consequence of utilitarianism. What should be noted is that this criticism does not raise doubts about utilitarianism as a theory of obligation-making characteristics. When utilitarianism is understood as a theory about what one objectively ought to do, it is untouched by this criticism. Utilitarianism does in principle provide a moral answer to every situation the agent faces.

This response, I believe, shows that the common objection to utilitarianism can be answered. But, ironically, this same response can be made to the utilitarian's charge that Ross's theory provides no adequate way of handling cases of moral conflict. We may set out Ross's deontological theory as follows:

(RD) An agent is morally obliged to do A in circumstances c if and only if doing A in c produces a greater balance of *prima facie* rightness over *prima facie* wrongness than any alternative action open to the agent in c.<sup>(7)</sup>

Utilitarians typically claim that Ross's theory (and all other deontological theories) provides no solution at all to the problem of how to handle conflicting ought-claims. Surely, though, Ross can respond, à la the utilitarian, that this is only an objection to his theory as a decision-making procedure. In principle, Ross might argue, we know that an agent ought to do that act which will maximize the balance of *prima facie* rightness over *prima facie* wrongness; thus as a theory of obligation-making characteristics, Ross's account handles the problem of ought-conflicts. Ross might go on to admit, though, that in many situations it is very difficult, and perhaps even impossible, for the agent to know what he ought to do in these conflict cases; that is, Ross might concede that his theory provides no decision-making procedure.<sup>(8)</sup> It may be an important weakness of a theory that it provides one with no *practical* guidance in these difficult situations. However, the utilitarian can hardly object to Ross's view on this account, since his theory is surely plagued with the same difficulty. The utilitarian might, of course, have other objections against Ross's theory. He might, for example, argue that Ross's version of non-naturalism is incoherent or unintelligible. But this is not an objection

to the way the theory handles ought-conflicts. Objections of this sort are more fundamental in nature, and not appropriately discussed in this context. So the first assumption of the utilitarian's argument is at least questionable.

### III.

Let us now examine the second assumption of the argument, the assumption that any theory concerned with the maximization of a single moral value has an easy way of resolving moral conflicts. To reiterate, a genuine moral dilemma is a situation in which an agent ought to do each of two actions, both of which he cannot do. We may represent a dilemmatic situation as follows:

- (1) p ought to do A in c.
- (2) p ought to do B in c.
- (3) p x can not do both A and B in c.

The second assumption of the utilitarian's argument claims that utilitarianism rules out situations characterized by (1), (2), and (3). According to (U) an agent ought to do an act if and only if it produces better consequences than any alternative open to him. If there is a situation in which more than one action is optimific, and the agent can do any one, but only one, of these acts, then he is to appeal to principle (P). He ought to do any one of the optimific acts, but which one he does is morally indifferent. Given this, the utilitarian will never affirm (1), (2), and (3) together; hence, genuine moral dilemmas are ruled out.

But is the situation this simple? If the utilitarian says that an agent ought to do act X, then he is committed to saying that act X possesses the obligation-making characteristics. Given the definition of (U), there are at least two possible utilitarian accounts of obligation-making characteristics. One possibility is this: if an agent ought to do act X in circumstances c, then X produces consequences *at least as good* as the consequences of any alternative open to the agent in c. Let us call this the *first* utilitarian account of obligation-making characteristics. The *second* utilitarian account is this: if an agent ought to do X in circumstances c, then X produces *better* consequences than any alternative open to the agent in c.

It is easy to show that the utilitarian cannot accept the first account. Consider a situation in which there are two maxima, A and B, either of which can be done but both of which cannot be done. By hypothesis, A produces consequences at least as good as the consequences of any alternative open to the agent in *c*; so the utilitarian must assert that the agent, *p*, ought to do A. By the same line of reasoning, the utilitarian must say that *p* ought to do B. Thus if we adopt the first utilitarian account of obligation-making characteristics, (1), (2), and (3) will be jointly satisfied in situations of many maxima; that is, the utilitarian will be committed to saying that there are genuine moral dilemmas. In these situations, no matter what an agent does, some maximum act will go unperformed. There is an additional problem with the first account. Given a commonly accepted distribution principle of deontic logic,

(DP) '*p* ought to do A in *c* and *p* ought to do B in *c*' entails '*p* ought to do both A and B in *c*',

it follows from (1) and (2) that

(4) *p* ought to do both A and B in *c*.

The conjunction of (4) and (3) provides a counterexample to the principle that 'ought' implies 'can'. Yet this is a principle that is central to utilitarianism. In fact, given the definition of (U), a necessary condition for an act's being morally required is that it be an alternative open to the agent. The utilitarian can avoid this latter result, of course, if he rejects (DP). Some, in fact, have argued that (DP) must be given up.<sup>(9)</sup> However, their main defense of this claim appeals to the assumption that there are genuine moral dilemmas; hence, the utilitarian cannot employ their arguments for rejecting (DP). So unless he has some new reason for denying this principle of deontic logic, such a move will seem much too *ad hoc*. The utilitarian must, then, reject the first account of obligation-making characteristics because it commits him to the view that there are genuine moral dilemmas and it forces him to give up the principle that 'ought' implies 'can'.

So the utilitarian must accept the second account of obligation-making characteristics. What makes X obligatory for an agent is that X produces better consequences than any alternative open to him. On this account the utilitarian must affirm neither (1) nor (2) in situations of many maxima. (1) is false because A does not produce better consequences than any alternative open to the agent; by hypothesis, it does not produce better consequences than B. By the same line of argument, (2) is false. In situations of many maxima, then, the agent is not obligated to do A and is not obligated to do B. Here the agent is to appeal to principle (P). What he is obligated to do, presumably, is the disjunctive act,  $A \vee B$ . One can now see why the utilitarian claims that while deontological theories face the *problem* of conflicting obligations, the analogue of these situations for his theory presents the agent with an abundance of riches. So the utilitarian does seem to have an advantage over the deontologist. But to say that in situations of many maxima one is obligated to do the disjunctive act  $A \vee B$  rather than act A and rather than act B commits the utilitarian to two apparently implausible assumptions.<sup>(10)</sup> (i) He must assume that the disjunctive act  $A \vee B$  is an alternative to act A and an alternative to act B (even when A and B cannot be jointly done). (ii) He must assume that the agent's performing the disjunctive act  $A \vee B$  will result in more utility than his performing act A or his performing act B.

There are at least two problems with assumption (i). First, it has been argued convincingly that two actions can be regarded as alternatives in the morally relevant sense only if they are incompatible.<sup>(11)</sup> In fact, if one does not adopt this view, it seems that there is no way to avoid a well known difficulty for utilitarianism raised by Castañeda.<sup>(12)</sup> And clearly the disjunctive act  $A \vee B$  is incompatible with neither A nor B; so it cannot be an alternative to these acts. Secondly, it has been argued that there is no plausible method of assigning values to the consequences of disjunctive actions that is available to the utilitarian. Any available method contravenes the spirit of utilitarianism.<sup>(13)</sup> If this is correct, the utilitarian must deny assumption (i).

Assumption (ii) seems even more implausible. If an act is obligatory on utilitarian grounds, it must produce *better* consequences than any alternative open to the agent. But how can the agent's performing the disjunctive act  $A \vee B$  (when A and B cannot both be done) produce



better consequences than his performing A or his performing B? This seems absurd; surely the consequences are exactly the same. So the utilitarian will want to avoid being committed to assumption (ii). Thus neither utilitarian account of the obligation-making characteristics eliminates ought-conflicts in a satisfactory way.

#### IV.

There are two possible escapes from this problem that must be considered. It is common in ethical theory to distinguish between saying that an act is right (i.e., permissible and justifiable) and saying that an act ought to be done (i.e., one which would be impermissible not to do). If an act is one that ought to be done, then it is right to perform it; but the converse does not hold. The utilitarian may employ this distinction in his definition of (U), and use it to avoid the problem that has been raised here. The amended definition is as follows: an act *A ought to be done* if and only if the consequences of doing A are *better* than those of any alternative open to the agent, and an act A is *right* if and only if the consequences of doing A are *at least as good* as those of any alternative open to the agent.<sup>(14)</sup> Given this, the utilitarian may say that in situations of many maxima there is *no act which the agent ought to perform* because there is no act the doing of which will have *better* consequences than any alternative. This will enable the utilitarian to maintain that there are no genuine moral dilemmas and to avoid the difficulty with disjunctive acts that has been discussed. This response will also allow him to answer another objection frequently raised against his theory, viz., that it requires acts of a trivial sort. The objection is that in any situation, no matter how trivial and seemingly unrelated to morality, there are some acts that will produce more good than the other alternatives, and so we are obligated to perform one of them. But if in situations of many maxima the utilitarian says that there are many rights acts that the agent might perform, but no act that he ought to perform, he will have answered this criticism too.

This suggestion is quite plausible, but it is not without its difficulties. In particular, this revised version of utilitarianism has consequences that are unacceptable. In some situations of many maxima it is

very counterintuitive to say that there is no act that the agent ought to perform. Consider the following case. Two persons (who are equally useful citizens),  $P_1$  and  $P_2$ , are trapped in a burning building, and it is within our power to save one but only one of them. According to this revised version of utilitarianism, saving  $P_1$  is a right act and saving  $P_2$  is a right act, but there is no act that we ought to perform. But surely in this case we want to make the stronger claim that one *ought* to save  $P_1$  or  $P_2$ . In short, this revised version of (U) forces us to treat all cases where there are two or more maxima as being on a moral par. It seems clear, however, that in some of these situations there is something that the agent is required to do, while in others there is no act that one ought to perform. Since this version of (U) cannot account for this, it must be rejected.

One other utilitarian escape route must be considered. The utilitarian may take 'is right' to be the primitive term in his theory, and then define 'is obligatory' in terms of rightness. The account of rightness is the same as before: an act is right if and only if it will have consequences at least as good as any alternative act. The term 'is obligatory' may now be defined as 'that which an agent *must* do in order to do what is right'. When there is only one act that will maximize the good, what is right and what is obligatory will be the same. In situations of many maxima, however, there is more than one right act. If A and B are the two maxima in a given situation, then on the proposed account each of these acts is right. And we can restrict the predicate 'is right' to apply only to act sets where the members are pair wise incompatible and jointly exhaustive. Thus the disjunctive act  $A \vee B$  will not appear as an alternative to A or to B. But what in this situation is obligatory? It is false that A is obligatory because the agent *need not* do A in order to do what is right; for this same reason, B is not obligatory. In this situation what the agent *must* do is some right act or other. Thus what is obligatory is that the agent do act A or act B. This is a modified and more plausible version of the second utilitarian account of obligation-making characteristics. It is more plausible because it enables the utilitarian to avoid talking about disjunctive acts, and hence he is no longer committed to assumptions (i) and (ii).

This account is a natural one for the utilitarian to offer. However, if we take the predicates 'is right', 'is wrong', and 'is obligatory' to be

moral grades that we assign to actions (or disjunctions of actions), several significant problems emerge. First, it is difficult to see how the disjunction 'act A or act B' can, on utilitarian grounds, get a different moral grade (viz., the grade 'is obligatory') than act A alone or B alone (each of which gets the grade 'is right', but not 'is obligatory'). There is *no value produced* by the disjunction of acts A and B which is not also brought about by A alone or B alone. Yet this version of (U) attributes a different moral status to the disjunction 'act A or act B' than it does to act A or to act B. Because of this, what we have is no longer a pure teleological theory.<sup>(15)</sup> Something in addition to the utilities produced determines to what the moral grade 'is obligatory' is assigned. Second, when we are considering which acts are *right*, the question of assessing the disjunction 'act A or act B' does not arise since it is not an alternative to A or to B; indeed, it is not an act at all. But when one asks what is *obligatory*, the answer is act A or act B. On this view, then, when we change moral categories, what is available to be assessed changes too. This seems implausible, or at least suspicious. How can something be available for assessment when we are asking what is obligatory, but not be available when we are inquiring about what is right? Finally, on the proposed definition of (U), the disjunction of acts A and  $\sim A$  (for any A) will be obligatory since, trivially, one *must* do this in order to do what is right. Surely, though, we do not want to say in all cases that an agent is required to do act A or act  $\sim A$ . In many cases we will want to assign the grade 'morally neutral' to such a disjunction of actions. In any case, making such requirements is pointless since the agent cannot help but satisfy them. This version of (U), then, does not seem to be an adequate one.

So the assumption that any theory concerned with the maximization of a single moral value has an easy way of resolving ought-conflicts is a dubious one. Situations of many maxima represent the one case where the utilitarian must say that what is obligatory and what is right are not identical. Any account of obligatoriness and rightness that the utilitarian gives must satisfy the following requirements: it must not allow each of two maxima, A and B, to be obligatory (or else there will be moral dilemmas); it must allow that each of A and B is right; and it must not allow  $A \vee B$  to be an alternative to A and to B. Neither of the first two accounts satisfies all of these requirements. And though the latter two accounts discussed here do fulfill these requirements, there

are other serious problems with them. It does not follow that the utilitarian must allow for genuine dilemmas. What has been shown, though, is that eliminating real moral conflicts poses theoretical difficulties (of an internal sort) for utilitarianism that are just as severe as those that plague deontological theories.<sup>(16)</sup>

University of North Carolina  
at Greensboro

Terrance C. McCONNELL

#### NOTES

<sup>(1)</sup> Included among the many who argue that the utilitarian's handling of ought-conflicts is a strength of the theory are John Stuart MILL, *Utilitarianism* (Indianapolis: Bobbs-Merrill, 1957), 32-33, William K. FRANKENA, *Ethics* (Englewood Cliffs: Prentice-Hall, second edition, 1973), 34; and Rolf E. SARTORIUS, *Individual Conduct and Social Norms* (Encino, Calif.: Dickenson Publishing, 1975), 10-11. Those who argue that one of the weaknesses of utilitarianism is that it rules out the possibility of genuine moral dilemmas are Bernard WILLIAMS, *Morality: An Introduction to Ethics* (New York: Harper & Row, 1972), 93-94; P.H. NOWELL-SMITH, «Some Reflections on Utilitarianism», *Canadian Journal of Philosophy*, II (1973), 430; and Richard BRONAUGH, «Utilitarian Alternatives», *Ethics*, 85 (1975), 177.

<sup>(2)</sup> Such arguments are presented in my «Moral Dilemmas and Consistency in Ethics», *Canadian Journal of Philosophy*, VIII (1978), 269-287.

<sup>(3)</sup> This familiar argument is set out explicitly by SARTORIUS in *Individual Conduct and Social Norms*, 10-11.

<sup>(4)</sup> MILL, *Utilitarianism*, 33.

<sup>(5)</sup> This objection to utilitarianism is presented by Stuart M. BROWN, Jr., «Duty and the Production of Good», *The Philosophical Review*, LXI (1952), 308-309, and by P.H. NOWELL-SMITH, «Some Reflections on Utilitarianism», 430.

<sup>(6)</sup> This distinction is set out by R. Eugene BALES in «Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?» *American Philosophical Quarterly*, 8 (1971), 257-265.

<sup>(7)</sup> Such an account is suggested by Ross himself in *The Right and the Good* (Oxford: Clarendon Press, 1930), 41 and 46, and in *Foundations of Ethics* (Oxford: Clarendon Press, 1939), 84.

<sup>(8)</sup> In fact, ROSS does argue in this way in *The Right and the Good*, 30-32 and 41-42.

<sup>(9)</sup> Two who have presented this argument for rejecting (DP) are Bernard WILLIAMS, «Ethical Consistency», *Aristotelian Society Supplementary Volume*, 39 (1965), 120-122, and Bas C. VAN FRAASSEN, «Values and the Heart's Command», *The Journal of Philosophy*, LXX (1973), 15. For a critical discussion of the latter's argument, see my «Moral Dilemmas and Requiring the Impossible», *Philosophical Studies*, 29 (1976), 409-413.

(<sup>10</sup>) One way for the utilitarian to avoid being committed to these two assumptions is to dispense with the notion of *disjunctive acts* and instead talk about a *disjunction of different acts*. This is a modified version of the second utilitarian account of obligation-making characteristics, and I deal with it in section IV.

(<sup>11</sup>) See Lars BERGSTROM, *The Alternative and Consequences of Actions* (Stockholm: Almqvist & Wiksell, 1966) and his «Utilitarianism and Deontic Logic,» *Analysis*, 29 (1968), 43-44.

(<sup>12</sup>) The problem alluded to here is set out by Hector-Neri CASTAÑEDA in «A Problem for Utilitarianism,» *Analysis*, 28 (1967), 141-142. What he shows is that utilitarianism is inconsistent with the converse of what I call (DP), viz.,

(DP') 'p ought to do A and B in c' entails 'p ought to do A in c and p ought to do B in c'.

CASTAÑEDA himself, in «Ought, Value, and Utilitarianism,» *American Philosophical Quarterly*, 6 (1969), 258, suggests that the most reasonable way to avoid this problem is the one suggested by BERGSTROM. See also, Dan W. BROCK, «Recent Work in Utilitarianism,» *American Philosophical Quarterly*, 10 (1973), 249-250.

(<sup>13</sup>) CASTAÑEDA, «Ought, Value, and Utilitarianism,» section VI. For a more detailed discussion of this problem see J. Howard SOBEL, «Value, Alternatives, and Utilitarianism,» *Nous*, 5 (1971), 373-384.

(<sup>14</sup>) Such a definition of (U) is presented by Lars BERGSTROM in «On the Formulation and Application of Utilitarianism,» *Nous*, 10 (1976), 144.

(<sup>15</sup>) «Pure teleological principles are concerned, ultimately, only with the values of the consequences of actions. Thus, when pure teleological principles are applied, particular actions may be viewed only with respect to their teleologically significant properties, that is, those properties *in virtue of which* actions produce utilities or disutilities.» David LYONS, *Forms and Limits of Utilitarianism* (Oxford: Clarendon Press, 1965), 57.

(<sup>16</sup>) For their many helpful comments on earlier versions of this paper, I would like to thank Richard Brough, Norman Dahl, Barry Hoffmaster, and especially Holly Goldman and Rolf Sartorius, whose suggestions I have incorporated in section IV of this essay.