

better action performed a worse one or of why the agent who could have performed a worse action performed a better one rules out even the explanation that it was a matter of chance.

In order to have control over the moral quality of his actions, an agent must have certain requisite abilities—in particular, the abilities necessary to see and understand the reasons and interests he ought to see and understand and the abilities necessary to direct his actions in accordance with these reasons and interests. And if, furthermore, there is nothing that interferes with the agent's use of these abilities—that is, no determining cause that prevents him from using them and no statistical process that, as it were, takes out of his hands the control over whether or not he uses them—then it seems that these are all the abilities that the agent needs to have. But it is compatible with the agent's having these abilities and with there being no interferences to their use that it is not determined whether the agent will perform a good action or a bad one. The responsible agent who performs a bad action fails to exercise these abilities sufficiently, though there is no complete explanation of why he fails. The responsible agent who performs a good action does exercise these abilities—it may or may not be the case that it is determined that he exercise them.

The freedom required for moral responsibility, then, is the freedom to be good. Only this kind of freedom will be neither too much nor too little. For then the agent is not so free as to be free from moral reasons, nor so unfree as to make these reasons ineffective.

SUSAN WOLF

Harvard University

#### THE PRINCIPLE OF MORAL HARMONY \*

**W**ITH a few exceptions, moral philosophers seem to be agreed that, at the level of the individual, morality doesn't necessarily pay. Hardly anyone who thinks about it seriously would maintain that doing what he morally ought to do invariably benefits the agent more than would some worse alternative. However, when we rise from the level of the individual

\* Earlier versions of this paper were read at several places, including SUNY at Fredonia, DePauw, Vanderbilt, Ohio, and Northern Illinois Universities. I am grateful to many of the participants in the discussions for useful comments. Many other friends read and commented on the paper. I am grateful to all of them, especially Eva Bodanszky, Earl Conee, and Terry Horgan.

to the level of the social group, we find that the reverse is true. Quite a few moral philosophers seem to believe that when all the members of a social group do what they morally ought to do, the group as a whole does benefit more than it would have from the performance of any worse alternative set of actions. I shall say that any such view is a version of the Principle of Moral Harmony (PMH). These doctrines are the subject of this paper.

My procedure here is as follows: first I sketch some of the main forms of PMH, and some of the main uses to which it has been put. I then attempt to show that some of the most plausible forms of PMH are false. After attempting to deal with some possible objections, I conclude by assessing the significance of what I have shown.

I

PMH may be formulated in a wide variety of ways. One version is suggested by Berkeley in *Passive Obedience*.<sup>1</sup> Berkeley held that God has brought it about that certain sorts of action are morally right. God's judgments about the rightness of these actions are expressed in a set of moral rules, or precepts. Berkeley calls the set of these rules "the law of nature." Concerning it, he says:

The law of nature is a system of such rules or precepts as that, if they be all of them, at all times, in all places, and by all men observed, they will necessarily promote the well-being of mankind, so far as it is attainable by human actions (239).<sup>2</sup>

In this form, PMH tells us that if everyone always does his duty, mankind as a whole is as well off as it can possibly be. So Berkeley's version of PMH applies only to the case in which we have universal duty doing. In this case, he seems to be saying, we are best off. It should be noted that this version of PMH does not have any direct bearing on the question whether large-scale, but nonuniversal, duty doing promotes welfare more or less than smaller-scale duty doing.

Bentham's version of PMH seems to make just this point. Bentham apparently held that the principle of utility tells us what our duties are. To "pursue" this principle is to do the acts it tells us to do—in other words, to do our duties. Bentham says:

The principle of utility is capable of being consistently pursued; and it is but tautology to say that the more consistently it is pursued, the better it must ever be for humankind.<sup>3</sup>

<sup>1</sup> George Berkeley, *Passive Obedience*, in Mary W. Calkins, ed., *Berkeley: Selections* (New York: Scribner's, 1929), pp. 427–469.

<sup>2</sup> It may be interesting to compare Berkeley's view with Leibniz's. See especially paragraphs 83–90 of the *Monadology*.

<sup>3</sup> Jeremy Bentham, *Introduction to the Principles of Morals and Legislation*, in E. A. Burtt ed., *The English Philosophers from Bacon to Mill* (New York: Modern Library, 1949), p. 796.

So Bentham apparently wished to maintain that if most of us were to do our duty most of the time, "humankind" would be better off than it would be if fewer of us were to do our duty less of the time. More generally, his position seems to be that the greater the extent to which we do what we should, the better off we'll be.

Some contemporary philosophers have affirmed versions of PMH, too. Stephen Toulmin, for example, claims that the "function of morality" is "to correlate our feelings and behavior in such a way as to make the fulfillment of everyone's aims and desires as far as possible compatible."<sup>4</sup> This statement is open to a variety of interpretations, and the context does not make clear just which one Toulmin meant to affirm. Of these, however, the following seems to me to be reasonable, and it may be the one Toulmin had in mind. First we can say that the "fulfillment value" of a state of affairs for a society is the extent to which that state of affairs leads to the fulfillment of the aims and desires of members of that society. We can understand these aims and desires to be the ones they actually have, although ones they might have had under various counterfactual circumstances could also be considered. In any case, Toulmin's statement suggests the view that the fulfillment value for society *S* of the state of affairs in which everyone in *S* does what is morally right is at least as great as the fulfillment value for *S* of any other state of affairs the members of *S* could bring about.

Kurt Baier, Hector-Neri Castañeda, and others have also suggested versions of PMH.<sup>5</sup> So it isn't hard to find moral philosophers who affirm PMH, in one form or another. Furthermore, many of these philosophers go on to make use of PMH in their work. Let us consider some of the main uses to which this doctrine has been (or might be) put.

Various moral philosophers, including Kurt Baier and Brian Medlin, have appealed to PMH in connection with attempts to refute ethical egoism.<sup>6</sup> An argument against this form of egoism, based on PMH, might go as follows: sometimes interests conflict in such a way that, if each participant in a group attempts to maximize self-interest, the group as a whole will be worse off. Thus, if each mem-

<sup>4</sup> *The Place of Reason in Ethics* (New York: Cambridge, 1950), p. 137; parenthetical page references to Toulmin will be in this book.

<sup>5</sup> Baier, *The Moral Point of View*, abridged ed. (New York: Random House, 1965), see especially pp. 106-109; parenthetical page references to Baier are to this book. Castañeda, *The Structure of Morality* (Springfield, Ill.: Charles C Thomas, 1974). J. L. Mackie, *Ethics* (New York: Penguin, 1977), see esp. ch. 5.

<sup>6</sup> Baier, *op. cit.*, pp. 95/6. Medlin, "Ultimate Principles and Ethical Egoism" in David Gauthier, ed., *Morality and Rational Self-Interest* (Englewood Cliffs, N.J.: Prentice-Hall, 1970), pp. 56-63.

ber does what egoism says to do, the group does not do as well as it would if they performed some alternative combination of actions. However, according to PMH, morally right actions cannot have this feature. Doing them must always maximize group welfare. Hence, egoistic actions are not morally right, and egoism is allegedly refuted.

This argument can be clarified by means of an example. Suppose there is a small community living around a lake. Let the community be called 'S', and suppose the members are Vincent, Margaret, George, and Mary. Suppose each produces a moderate amount of garbage, and, as things stand, each takes his own garbage to the community dump a few miles away. Vincent, we may suppose, recognizes that he has two main alternatives with respect to his garbage: he can either (a) take it to the dump, or (b) secretly throw it into the lake. Suppose he could get away with throwing it into the lake, no harm would come from doing so, and he would thereby avoid a bothersome trip to the dump. Egoism says he morally ought to throw his garbage into the lake.

However, egoism tells Margaret, George, and Mary the same thing. Each of them, according to this view, should take advantage of his or her neighbors, and use the lake as a dump. But, if they all did so, the lake would soon be polluted, they would all be worse off, and the community would suffer. PMH says that morally right actions never have this feature. Their joint performance must be maximally beneficial to the community. Hence, it has been argued, egoism must be wrong.

PMH also figures prominently in some discussions of the question, Why should I be moral? Baier claims that one of the main functions of morality is to provide us with a systematic way of resolving conflicts of interest. "By the moral point of view," he says, "we mean a point of view which furnishes a court of arbitration for conflicts of interest" (96). Baier goes on to claim that the acceptance of a moral system "is in the interest of everyone alike" (106–109). His point seems to be that social groups benefit from having moralities since the moral rules provide an especially good way of resolving conflict. Thus, if each of us does what the rules specify, the group as a whole will be better off than it would have been if each of use had behaved in some other way. In light of this, each member of the social unit has an overridingly good reason to "be moral." Each of us should act in accord with the moral rules of our society, Baier seems to argue, because the group will be best off if we do (see 148–162).

Other philosophers make use of PMH in their attempts to give an account of the nature of morality. This is an interesting puzzle: what is it that makes a set of rules a set of *moral* rules rather than, for example, a set of rules of etiquette or prudence? One suggestion is that moral rules must be designed to serve a certain function. Toulmin, as we have seen, says that the function of morality is "to correlate our feelings and behavior in such a way as to make the fulfillment of everyone's aims and desires as far as possible compatible" (137). He goes on to claim that "what makes us call a judgment 'ethical' is the fact that it is used to harmonize peoples' actions" (137). So Toulmin would apparently want to appeal to a version of PMH in an attempt to explain what morality is. The morality of a society is a set of rules adopted by that society for the purpose of resolving "lower-level" conflicts, and thereby "harmonizing" peoples' actions. Thus, a form of PMH might be said to be true in virtue of the meaning of the term 'ethical'.<sup>7</sup>

PMH might also figure in an argument for act utilitarianism. For purposes of argument, we can assume that the "well-offness" of a society may be considered to be the same as the total amount of utility enjoyed by the members of the society. Act utilitarianism tells each of us to do what maximizes utility. It may seem obvious that if each of us does what act utilitarianism tells us to do, our society will be best off.<sup>8</sup> PMH says that morally right actions have this feature, too. Hence, we may want to conclude that the acts act utilitarianism tells us to do are one and all morally right. In other words, act utilitarianism is true.

So PMH has been formulated in a variety of ways, and has been (or might be) put to a variety of philosophical uses.

## II

Now let us consider some of these doctrines a bit more critically.

Bentham maintained that "it is but tautology to say that the more consistently [the principle of utility] is pursued, the better it must ever be for humankind." I take this to mean that the greater the extent to which we individually do what maximizes utility, the better off we'll be collectively. Let's call this PMH<sub>2</sub>.

Even if we evaluate our collective "well-offness" strictly in terms of how much utility we collectively enjoy, this is surely no tautology. It seems clear enough that two moderately bad wrong actions may have a much less bad total impact on our collective well-offness than one moderately good right action together with an absolutely hor-

<sup>7</sup> It might be instructive to compare this view with that proposed by Castañeda in *The Structure of Morality*, ch. 8.

<sup>8</sup> This view is criticized below, pp. 175 ff.

rendous wrong one. Yet the extent to which we do our duty when we do it not at all is apparently less than the extent to which we do our duty when we do it once. So Bentham's principle is no tautology.

To see this in a more concrete case, consider the following possibility. Suppose Vincent has his choice of three alternatives,  $v_1$ ,  $v_2$ , and  $v_3$ , and Margaret has her choice of two:  $m_1$  and  $m_2$ . Suppose these acts are unconnected, and have the following utilities:

$$\begin{array}{ll} v_1 & +2 \\ v_2 & -5 \\ v_3 & -10,000 \end{array} \qquad \begin{array}{ll} m_1 & +2 \\ m_2 & -5 \end{array}$$

If the group composed of Margaret and Vincent (we can consider it to be the set,  $\{M, V\}$ ) does the group act  $\{v_2, m_2\}$ , it does two wrong actions. The group as a whole suffers a utility deficit of  $-10$  (given some natural assumptions). Yet if the group does  $\{v_3, m_1\}$ , then they at least do one right action. Yet the group is much worse off. It suffers a utility deficit of  $-9,998$ . I can see no reason to suppose that similar results would be impossible in the case of a much larger group (even humanity as a whole) or over a much longer period of time (even eternity). Thus, I have to reject  $PMH_2$ .

A more plausible form of PMH might be developed by focusing on a single moment and a single social group. Consider the set of actions the members of the group would perform if each were to do something that's morally obligatory for him or her then. Won't the group as a whole be better off as a result of that than it would have been as a result of any alternative set of actions for the group including some wrong ones? This view, which we may call  $PMH_3$ , may seem attractive, especially to those who look favorably upon utilitarianism. Nevertheless, it seems to me that this version of PMH is false, too. Let us consider an example.

Suppose a group of adults has taken a group of children out to do some ice skating. The adults have assured the children and their parents that, in case of accident, they will do everything in their power to protect the children. Each adult in the party is a good skater and swimmer. Suppose, finally, that, while they are out skating, it just so happens that all the adults are spread out around the edge of the pond. A lone child is skating in the middle, equidistant from the adults. Suddenly, the ice breaks, and the child falls through. There is no time for consultation or deliberation. Someone must quickly save the child. However, since the ice is very thin, it would be disastrous for more than one of the adults to venture near the place where the child broke through. For if two or more were to go out, they would all fall in and all would be in profound trouble. In

fact, let us suppose, no one goes to the aid of the child, and it is only after hard and painful efforts that the child manages to scramble out of the water.

In my view, every one of the adults in this example has failed to do his or her moral duty. Each of them is such that he or she morally ought to have gone to the aid of the child. This intuition is supported by virtually every plausible moral theory I can think of. On act utilitarianism, for example, it is clear that each adult should have saved the child. For each adult faced roughly two main alternatives: to attempt to save the child, or not. Since no one did attempt to save the child, and no one's attempt would have provoked anyone else to try, each adult is such that if he or she had attempted to save the child, he or she would have done so successfully and alone. Hence, attempting to save the child maximizes utility for each adult. On Ross's theory, it seems plausible to say that each should have attempted to save the child. For each promised to do so, thus creating a *prima facie* duty to do so. I can see no overriding *prima facie* duty that lets anyone off the hook. So each seems to have an all-in duty to save the child. Various forms of Kantianism, it seems to me, may plausibly be thought to imply that each adult ought to have gone to the aid of the endangered child.

I prefer, however, to let the example stand or fall with reflective moral intuition, rather than with some controversial normative theory. My moral intuition tells me that this is a case in which each adult should have gone to the aid of the child. In my view, this is closely related to the stipulated fact that each of the adults does this in all the best possible worlds accessible to him or her at the time.

It is also quite clear, I think, that this example refutes PMH<sub>3</sub>. For here we have a case in which there is a time at which, if each member of the social group were to do what in fact is his or her moral duty, then the group as a whole would be much worse off than it would have been if only one were to do what in fact is his or her duty. For it is clear that if all were to go out on the thin ice, many would get cold and wet. Indeed, we may even suppose that all of them, including the child, would fail to survive. That result is enormously worse than the result of just one doing his or her duty. For in that case only two get cold and wet, and all survive. Thus, PMH<sub>3</sub> is false. It is not the case that a group is inevitably better off if everyone does his or her duty.

The ice-skating example illustrates a sort of case that refutes many forms of PMH. The important aspects of the case may be illustrated in a simple utility matrix. In this matrix, we show a case in

which there are two agents,  $V$  and  $M$ .  $V$  has two choices,  $v_1$  and  $v_2$ , and  $M$  has two choices,  $m_1$  and  $m_2$ . We may think of  $v_1$  and  $v_2$  as representing, respectively,  $V$ 's going out on the ice, and  $V$ 's remaining on the edge of the pond. Similarly for  $m_1$  and  $m_2$ . The numbers in the matrix represent the amounts of utility that would result from the joint performance of the relevant actions.

		$M$	
		$m_1$	$m_2$
$V$	$v_1$	$A_{-100}$	$B_{+100}$
	$v_2$	$C_{+100}$	$D_{-50}$

In fact, let us suppose that  $m_2$  is performed. Then  $V$  may choose between doing  $v_1$  and thereby helping to produce +100 units of utility (square  $B$ ) or doing  $v_2$ , and thereby helping to produce -50 (square  $D$ ). Clearly, act utilitarianism implies that  $v_1$  is the morally right choice.

Let us also suppose that  $v_2$  is in fact performed. Then  $M$  may choose between doing  $m_1$  and thereby helping to produce +100 (square  $C$ ) or else doing  $m_2$  and thereby helping to produce -50 (square  $D$ ). Clearly,  $m_1$  is obligatory, at least on utilitarian grounds.

The upshot is that  $m_1$  and  $v_1$  are reasonably thought to be obligatory, even though  $\{m_1, v_1\}$  is clearly the worst group action of the lot. Given some innocent and natural assumptions (such as that the actions here are timewise identical and that the agents in the example ought to do the acts that have the best results), this shows that  $PMH_3$  is false.

Reflection on this utility matrix may shed some light on the connection between PMH and act utilitarianism. In my view, there are cases in which the implication of these doctrines conflict. We must not assume that if each of a set of actions is utility-maximizing (and so individually right according to act utilitarianism) then the group action composed of them will also be utility-maximizing (and so right according to a version of act utilitarianism applicable to group actions).

When we assign a utility to an action, we consider how much pleasure and pain it would produce, if it were performed. In other words, we consider how much pleasure and pain it does produce in "the nearest possible world" in which it is performed. It may be that in the nearest world where an action  $v_1$  is performed, it produces a lot of pleasure and no pain. It may be that in that world another action



$m_1$  is not performed. Furthermore, it may be that in the nearest world in which  $m_1$  is performed, it produces a lot of pleasure and no pain, and  $v_1$  is not performed. It is consistent with these assumptions to stipulate further that in the nearest world in which  $\{m_1, v_1\}$  is performed, it produces a lot of pain and no pleasure. The utility of a group action bears no interesting relation to the sum of the utilities of the actions composing it.

A utilitarian might want to insist that the utility of a group action has to be equal to the sum of the utilities of its components. This seems to me to be straightforwardly false, as I have tried to show. Something like it, however, may be true. We should note that actions have their utilities contingently. It makes sense to speak of the utility an action would have had, under various counterfactual circumstances. We may want to say that the utility of a group action is equal to the sum of the utilities its components would have had, if they had been performed together. In this way, we focus on the nearest world in which the group action is performed, and maintain that the utility produced by the collective act must be distributed among its parts. This view seems to me to be reasonable, although I am not entirely convinced that it is true.<sup>9</sup>

So if a utilitarian wants to extend his normative theory so as to make it apply to group actions, he or she must be wary of assuming that utility-maximizing group actions are invariably composed of utility-maximizing components. Thus, the utilitarian must be wary of saying that morally right group actions are invariably composed of morally right components. In fact, I can see no straightforward way to connect the normative status of a group action with that of each of its components. An act can be right for a group to perform, and each component thereof wrong for its agent, and vice versa.<sup>10</sup>

### III

Some might claim that my ice-skating example seems to work only because it is a case in which in fact everyone fails to do his or her duty. Furthermore, it might be claimed,  $PMH_3$  is a principle about what happens when everyone does his or her duty. Hence, the example might be thought not to run counter to  $PMH_3$ , properly understood.

I'm not moved by this sort of objection. In the ice-skating example, each of the adults ought to perform a certain action.  $PMH_3$  tells us that, if they were all to perform these actions, the group as a

<sup>9</sup> So far as I know, the best discussion of this topic is still to be found in David Lyons, *Forms and Limits of Utilitarianism* (New York: Oxford, 1965).

<sup>10</sup> For an interesting twist on this point, see Gerald Barnes, "Utilitarianisms," *Ethics*, LXXXII, 1 (October 1971): 56-64.

whole would be best off. I have claimed that, if they were to perform these actions together, the group would be far worse off. I admit that, had some of them behaved differently, some of these actions would not have been obligatory, Nevertheless, as it stands, the example does run counter to  $PMH_3$ .

In any case, it may be instructive to consider a revision of  $PMH_3$ . Let us formulate a version that applies only to cases in which everyone in the example in fact does what's obligatory. This may suffice:

$PMH_4$ : If there is a society  $S$  and a time  $t$  such that every member of  $S$  does something he or she morally ought to do at  $t$ , then  $S$  as a whole is better off than it would have been if some members had done something wrong instead of what they did.

This principle is of less interest than  $PMH_3$ , since it applies only to cases in which all the members of some society do what they should. Nevertheless, I think  $PMH_4$  is false. I shall assume a version of utilitarianism for ease of exposition here, and present and analyze a purported counterexample.

Suppose the residents of the village are plagued by mosquitoes. Suppose there is only one way in which they can get rid of the mosquitoes. Each of the residents would have to fumigate his or her yard with a special insecticide, and they would all have to do it simultaneously at noon. If some, but not all of them were to fumigate, terrible results would ensue. In this case, universal cooperation is essential. Suppose no one else is in fact going to fumigate at noon (although each of them easily could). Suppose nothing Vincent can do will make any of the others cooperate. Vincent has two main choices: he can fumigate at noon, or he can refrain from fumigating at noon. Clearly, if he fumigates, he does so alone and so helps to produce very bad results. If he does not fumigate at noon, nothing especially bad will happen, but the mosquitoes will not be killed. So it seems that Vincent ought to refrain from fumigating.

Similar arguments apply to each of the rest. Each of them ought to refrain from fumigating at noon. Suppose noon comes, and no one fumigates. Each of them does what he or she ought to do.

Now consider the relevant group action they perform at noon. This action is composed of several acts of nonfumigation, and each such act is morally obligatory. Yet it should be clear that the group as a whole had an alternative that would have had enormously better results. They could have got rid of the mosquitoes. So  $PMH_4$  is

refuted. In this example, every member of the group in fact does what he or she ought to do, and yet the group is not best off.<sup>11</sup>

This case has a structure slightly different from that of the ice-skating case. Its essential features can be represented in another utility matrix. Here we simplify again, assuming that the society has only two members, and that each member has only two choices:

		<i>M</i>	
		<i>m</i> <sub>1</sub>	<i>m</i> <sub>2</sub>
<i>V</i>	<i>v</i> <sub>1</sub>	<i>A</i> <sub>+100</sub>	<i>B</i> <sub>-100</sub>
	<i>v</i> <sub>2</sub>	<i>C</i> <sub>-100</sub>	<i>D</i> <sub>+50</sub>

If we assume that *m*<sub>2</sub> is in fact performed and view the matrix from *V*'s point of view, we see that *V* has two choices: he can do *v*<sub>1</sub> (fumigate alone), which will yield -100 (square *B*), or he can do *v*<sub>2</sub> (refrain from fumigating), which will yield +50 (square *D*). Clearly, *v*<sub>2</sub> is by far his better choice.

If we also assume that *v*<sub>2</sub> is in fact performed and view the matrix from *M*'s point of view, we see that she has two choices: she can do *m*<sub>1</sub>, which will yield -100 (square *C*), or she can do *m*<sub>2</sub>, which will yield +50 (square *D*). Clearly, *m*<sub>2</sub> is her better choice.

Yet the group action {*m*<sub>1</sub>, *v*<sub>1</sub>} is preferable from the standpoint of group welfare. The group {*M*, *V*} would be much better off if {*m*<sub>1</sub>, *v*<sub>1</sub>} were performed than it would be if {*m*<sub>2</sub>, *v*<sub>2</sub>} were performed. Yet, as we have seen, {*m*<sub>2</sub>, *v*<sub>2</sub>} contains two obligatory actions, each of which in fact was performed. I think this sort of example shows that PMH<sub>4</sub> goes wrong.

Another sort of objection to my ice-skating example may be based on the view that I have failed to distinguish between conditional obligation and absolute obligation. Thus, some may grant that Vincent has a conditional obligation to save the child if no one else does, but they may want to deny that Vincent has an absolute obligation to save the child. These objectors could go on to claim that the same is true of each of the other adults in the example. Each

<sup>11</sup> As we saw above, Berkeley apparently maintained that if everyone eternally does what's morally right, the "well-being of mankind" will be maximized. Let's call this PMH<sub>1</sub>. One obvious problem with this view is that it lacks practical relevance. It is too late to actualize the only case to which it applied. Beyond that, however, I think we can now see that this principle, too, is no conceptual truth. For there might be a world whose whole history consists of what goes on in our community at noon. There everyone "eternally" does his duty. Yet the group would have been better off if they had all acted otherwise.

conditionally ought to save the child if no one else does, but none absolutely ought to save the child.

If this claim is correct, then my alleged counterexample does not work. For PMH<sub>3</sub> does not say anything about group actions composed of merely conditionally obligatory individual actions. Furthermore, it would appear that no bad results would ensue if each adult in the example were to go to the aid of the child if no one else did. For, if they all did that, then exactly one would go to the aid of the child, and group welfare would be maximized.

This objection obviously raises some serious issues, and I can't attempt to settle them here.<sup>12</sup> Nevertheless, I want to sketch an answer.

In the first place, I grant that Vincent conditionally ought to go to the aid of the child if no one else does. Similarly for the rest. Secondly, I admit that no bad results would be produced by the joint fulfillment of these conditional obligations. Nevertheless, I also believe that each adult absolutely ought to go to the aid of the child. I've given some reasons for thinking this to be true. Each promised. Each is such that nothing else he or she could do would produce better results. Each has a non-overridden *prima facie* duty to do so. I think it is clear that the results of the group action composed of all these acts of going to the aid of the child would be terrible. So I feel that I have presented a legitimate counterexample to PMH<sub>3</sub>.

Moreover, I believe that it can be shown that if each adult has a conditional obligation to go to the aid of the child if no one else does, and no one has it in his power to make any of the others go, then each has an absolute obligation to go to the aid of the child. To see this, let us focus on Vincent, and see what he ought to do, absolutely and conditionally.

It is agreed that Vincent ought to go to the aid of the endangered child if no one else does. It is also agreed that in fact no one else does. However, from these two premises we may not infer that Vincent absolutely ought to go to the aid of the child. We cannot detach an absolute obligation from a conditional obligation and its condition.

Yet in the case at hand I believe we have another premise that enables us to detach our conclusion. That premise is that the condition is "inevitable." More precisely, it is that, no matter what Vincent does, the others will not go to aid of the child. Nothing he can do will make them do it. This fact, together with the conditional

<sup>12</sup> I have discussed these issues at greater length in "Iffy Oughts," currently unpublished.

obligation to go to the aid of the child if no one else does, entails that Vincent absolutely ought to go to the aid of the child.<sup>13</sup>

Corresponding arguments apply to each of the other adults. Therefore, the objection can be rebutted. If the adults have the conditional obligation the objectors allege them to have and it is stipulated that no one can make anyone else go to the aid of the child, then each of them absolutely ought to go to the aid of the child. Hence, I feel that my counterexample stands.

#### IV

By way of conclusion, I want to say a few words about what the failure of PMH may show. In the first place, I think we must view with skepticism any philosophical argument in which PMH appears as a premise. So, for example, if someone attacks egoism, appealing to the fact that egoism conflicts with some form of PMH, we will probably have to reject the argument. For we have no reason to suppose that any interesting form of PMH is true. Equally, if someone attempts to defend act utilitarianism by claiming that, of all normative theories, it alone is consistent with some form of PMH, we must reject this argument as well. For, in the first place, most forms of PMH are false. In the second place, I believe it has been shown that act utilitarianism is not consistent with the interesting forms of PMH. Similar comments apply to all other uses of PMH in arguments or theories. Insofar as any philosophical view depends upon PMH, that view is suspect.

But beyond this, I think that the failure of PMH may show us something of interest about the concept of morality itself. Some moral philosophers think of morality as a sort of social institution, designed to serve certain social purposes. For example, some suggest that morality is designed to harmonize conflicting interests. Others say it is designed to maximize social welfare.

I am convinced that this conception of morality is mistaken. Of course, I recognize that social groups adopt or establish various moral principles, and that frequently the selection of these principles is in some way guided by judgments of social utility. However, it seems clear to me that this sort of adoption or establishment of a moral principle cannot help to make the principle true, if in fact it is false. Beliefs about what's right and wrong, no matter how widespread, may be mistaken.

So let us distinguish between moralities—moral codes adopted by various social groups—and Morality—the truth about what we

<sup>13</sup> See also P. S. Greenspan, "Conditional Oughts and Hypothetical Imperatives," this JOURNAL, LXXII, 10 (May 22, 1975): 259–276.

ought to do. I admit that moralities may be social instruments (although I suspect that no even remotely plausible morality would be consistent with any interesting form of PMH). My main thesis, and what I've attempted to show here, is that Morality is not consistent with PMH. If I am right about this, then we should reconsider the doctrine that Morality is primarily designed to maximize some social goal, such as harmony or welfare. I have tried to show that, if Morality has this purpose, then it fails to do what it is supposed to do. My own view, for what it's worth, is that Morality doesn't have any purpose. Indeed, the suggestion that Morality has a purpose strikes me as being as odd as the suggestion that mathematics, or ornithology, has a purpose.

FRED FELDMAN

University of Massachusetts/Amherst

### BOOK REVIEWS

*The Sense of Order*. E. H. GOMBRICH. Ithaca, N.Y.: Cornell University Press, 1979. 411 p. \$38.50.

This book complements *Art and Illusion*: "one (is) concerned with representation, the other with pure design" (ix). Does any *philosophical* problem of design correspond to the philosophical problem of representation? That is unclear. Consider some comparisons. Unlike a representation, a decoration stands for nothing. Usually a represented figure is in front of a background. Decoration is all background. Seeing representations requires a "search for meaning" (256). Seeing decoration requires global, unfocused perception. Expectation about how a pattern will repeat is analogous to "the beholder's share" in representation. A design motif, like a representational schema, gives a way of organizing the making and seeing of art.

For representations, but not designs, there is an objective test of accuracy. A Constable copies what it depicts more completely than a Giotto. "Given the aim of creating a convincing picture of reality, this is the way the arts will 'evolve' . . . (for decoration) both the aims and the means are more diffuse" (210). When discussing representations, Gombrich shows that we cannot understand the meaning of a work without knowing its historical context. By contrast, little space here is devoted to the historical context of dec-