

Killing and Starving to Death

JAMES RACHELS

Although we do not know exactly how many people die each year of malnutrition or related health problems, the number is very high, in the millions.¹ By giving money to support famine relief efforts, each of us could save at least some of them. By not giving, we let them die.

Some philosophers have argued that letting people die is not as bad as killing them, because in general our 'positive duty' to give aid is weaker than our 'negative duty' not to do harm.² I maintain the opposite: letting die is just as bad as killing.³ At first this may seem wildly implausible. When reminded that people are dying of starvation while we spend money on trivial things, we may feel a bit guilty, but certainly we do not feel like murderers. Philippa Foot writes:

Most of us allow people to die of starvation in India and Africa, and there is surely something wrong with us that we do; it would be nonsense, however, to pretend that it is only in law that we make a distinction between allowing people in the underdeveloped countries to die of starvation and sending them poisoned food. There is worked into our moral system a distinction between what we owe people in the form of aid and what we owe them in the way of non-interference.⁴

No doubt this would be correct if it were intended only as a description of what most people believe. Whether this feature of 'our moral system' is rationally defensible is, however, another matter. I shall argue that we are wrong to take comfort in the fact that we 'only' let these people die, because our duty not to let them die is equally as strong as our duty not to kill them, which, of course, is very strong indeed.

¹ For an account of the difficulties of getting reliable information in this area, see Nick Eberstadt, 'Myths of the Food Crisis', *New York Review of Books* (19 February 1976), 32–37.

² Richard L. Trammell, 'Saving Life and Taking Life', *Journal of Philosophy* 72 (1975), 131–137, is the best defence of this view of which I am aware.

³ This article is a companion to an earlier one, 'Active and Passive Euthanasia', *New England Journal of Medicine* 292 (9 January 1975), 78–80, in which I discuss the (mis)use of the killing/letting die distinction in medical contexts. But nothing in this article depends on the earlier one.

⁴ Philippa Foot, 'The Problem of Abortion and the Doctrine of the Double Effect', *Oxford Review* No. 5 (1967); reprinted in J. Rachels (ed.), *Moral Problems*, 2nd edn (New York: Harper and Row, 1975), 66.

Obviously, this Equivalence Thesis is not morally neutral, as philosophical claims about ethics often are. It is a radical idea which, if true, would mean that some of our 'intuitions' (our prereflective beliefs about what is right and wrong in particular cases) are mistaken and must be rejected. Neither is the view I oppose morally neutral. The idea that killing is worse than letting die is a relatively conservative thesis which would allow those same intuitions to be preserved. However, the Equivalence Thesis should not be dismissed merely because it does not conform to all our prereflective intuitions. Rather than being perceptions of the truth, our 'intuitions' might sometimes signify nothing more than our prejudices or selfishness or cultural conditioning. Philosophers often admit that, in theory at least, some intuitions might be unreliable—but usually this possibility is not taken seriously, and conformity to prereflective intuition is used uncritically as a test of the acceptability of moral theory. In what follows I shall argue that many of our intuitions concerning killing and letting die *are* mistaken, and should not be trusted.

I

We think that killing is worse than letting die, not because we overestimate how bad it is to kill, but because we underestimate how bad it is to let die. The following chain of reasoning is intended to show that letting people in foreign countries die of starvation is very much worse than we commonly assume.

Suppose there were a starving child in the room where you are now—hollow eyed, belly bloated, and so on—and you have a sandwich at your elbow that you don't need. Of course you would be horrified; you would stop reading and give her the sandwich, or better, take her to a hospital. And you would not think this an act of supererogation: you would not expect any special praise for it, and you would expect criticism if you did not do it. Imagine what you would think of someone who simply ignored the child and continued reading, allowing her to die of starvation. Let us call the person who would do this Jack Palance, after the very nice man who plays such vile characters in the movies. Jack Palance indifferently watches the starving child die; he cannot be bothered even to hand her the sandwich. There is ample reason for judging him very harshly; without putting too fine a point on it, he shows himself to be a moral monster.

When we allow people in far-away countries to die of starvation, we may think, as Mrs Foot puts it, that 'there is surely something wrong with us'. But we most emphatically do not consider ourselves moral monsters. We think this, in spite of the striking similarity between Jack Palance's behaviour and our own. He could easily save the child; he does not; and the child dies. We could easily save some of those starving people; we do not;

and they die. If we are not monsters, there must be some important difference between him and us. But what is it?

One obvious difference between Jack Palance's position and ours is that the person he lets die is in the same room with him, while the people we let die are mostly far away. Yet the spatial location of the dying people hardly seems a relevant consideration.⁵ It is absurd to suppose that being located at a certain map co-ordinate entitles one to treatment which one would not merit if situated at a different longitude or latitude. Of course, if a dying person's location meant that we *could not* help, that would excuse us. But, since there are efficient famine relief agencies willing to carry our aid to the far-away countries, this excuse is not available. It would be almost as easy for us to send these agencies the price of the sandwich as for Palance to hand the sandwich to the child.

The location of the starving people does make a difference, psychologically, in how we feel. If there were a starving child in the same room with us, we could not avoid realizing, in a vivid and disturbing way, how it is suffering and that it is about to die. Faced with this realization our consciences probably would not allow us to ignore the child. But if the dying are far away, it is easy to think of them only abstractly, or to put them out of our thoughts altogether. This might explain why our conduct would be different if we were in Jack Palance's position, even though, from a moral point of view, the location of the dying is not relevant.

There are other differences between Jack Palance and us, which may seem important, having to do with the sheer numbers of people, both affluent and starving, that surround us. In our fictitious example Jack Palance is one person, confronted by the need of one other person. This makes his position relatively simple. In the real world our position is more complicated, in two ways: first, in that there are millions of people who need feeding, and none of us has the resources to care for all of them; and second, in that for any starving person we *could* help there are millions of other affluent people who could help as easily as we.

On the first point, not much needs to be said. We may feel, in a vague sort of way, that we are not monsters because no one of us could possibly save *all* the starving people—there are just too many of them, and none of us has the resources. This is fair enough, but all that follows is that, individually, none of us is responsible for saving everyone. We may still be responsible for saving someone, or as many as we can. This is so obvious that it hardly bears mentioning; yet it is easy to lose sight of, and philosophers have actually lost sight of it. In his article 'Saving Life and Taking Life',⁶ Richard Trammell says that one morally important difference

⁵ On this point, and more generally on the whole subject of our duty to contribute for famine relief, see Peter Singer, 'Famine, Affluence, and Morality', *Philosophy and Public Affairs* 1 (Spring 1972), 232.

⁶ Trammell, 133.

between killing and letting die is 'dischargeability'. By this he means that, while each of us can discharge completely a duty not to kill anyone, no one among us can discharge completely a duty to save everyone who needs it. Again, fair enough; but all that follows is that, since we are only bound to save those we can, the class of people we have an obligation to save is much smaller than the class of people we have an obligation not to kill. It does *not* follow that our duty with respect to those we can save is any less stringent. Suppose Jack Palance were to say: 'I needn't give this starving child the sandwich because, after all, I can't save everyone in the world who needs it'. If this excuse will not work for him, neither will it work for us with respect to the children we could save in India or Africa.

The second point about numbers was that, for any starving person we *could* help, there are millions of other affluent people who could help as easily as we. Some are in an even better position to help since they are richer. But by and large these people are doing nothing. This also helps to explain why we do not feel especially guilty for letting people starve. How guilty we feel about something depends, to some extent, on how we compare with those around us. If we were surrounded by people who regularly sacrificed to feed the starving, and we did not, we would probably feel ashamed. But because our neighbours do not do any better than we, we are not so ashamed.

But again, this does not imply that we should not feel more guilty or ashamed than we do. A psychological explanation of our feelings is not a moral justification of our conduct. Suppose Jack Palance were only one of twenty people who watched the child die; would that decrease his guilt? Curiously, I think many people assume it would. Many people seem to feel that if twenty people do nothing to prevent a tragedy, each of them is only one-twentieth as guilty as he would have been if he had watched the tragedy alone. It is as though there is only a fixed amount of guilt, which divides. I suggest, rather, that guilt multiplies, so that each passive viewer is fully guilty, if he could have prevented the tragedy but did not. Jack Palance watching the girl die alone would be a moral monster; but if he calls in a group of his friends to watch with him, he does not diminish his guilt by dividing it among them. Instead, they are all moral monsters. Once the point is made explicit, it seems obvious.

The fact that most other affluent people do nothing to relieve hunger may very well have implications for one's own obligations. But the implication may be that one's own obligations *increase* rather than decrease. Suppose Palance and a friend were faced with two starving children, so that, if each did his 'fair share', Palance would only have to feed one of them. But the friend will do nothing. Because he is well-off, Palance could feed both of them. Should he not? What if he fed one and then watched the other die, announcing that he has done *his* duty and that the one who died was his friend's responsibility? This shows the fallacy of

supposing that one's duty is only to do one's fair share, where this is determined by what would be sufficient *if* everyone else did likewise.

To summarize: Jack Palance, who refuses to hand a sandwich to a starving child, is a moral monster. But we feel intuitively that we are not so monstrous, even though we also let starving children die when we could feed them almost as easily. If this intuition is correct, there must be some important difference between him and us. But when we examine the most obvious differences between his conduct and ours—the location of the dying, the differences in numbers—we find no real basis for judging ourselves less harshly than we judge him. Perhaps there are some other grounds on which we might distinguish our moral position, with respect to actual starving people, from Jack Palance's position with respect to the child in my story. But I cannot think of what they might be. Therefore, I conclude that if he is a monster, then so are we—or at least, so are we after our rationalizations and thoughtlessness have been exposed.

This last qualification is important. We judge people, at least in part, according to whether they can be expected to realize how well or how badly they behave. We judge Palance harshly because the consequences of his indifference are so immediately apparent. By contrast, it requires an unusual effort for us to realize the consequences of our indifference. It is normal behaviour for people in the affluent countries not to give to famine relief, or if they do give, to give very little. Decent people may go along with this normal behaviour pattern unthinkingly, without realizing, or without comprehending in a clear way, just what this means for the starving. Thus, even though those decent people may act monstrously, we do not judge them monsters. There is a curious sense, then, in which moral reflection can transform decent people into indecent ones: for if a person thinks things through, and realizes that he is, morally speaking, in Jack Palance's position, his continued indifference is more blameworthy than before.

The preceding is not intended to prove that letting people die of starvation is as bad as killing them. But it does provide strong evidence that letting die is much worse than we normally assume, and so that letting die is much *closer* to killing than we normally assume. These reflections also go some way towards showing just how fragile and unreliable our intuitions are in this area. They suggest that, if we want to discover the truth, we are better off looking at arguments that do not rely on unexamined intuitions.

II

Before arguing that the Equivalence Thesis is true, let me explain more precisely what I mean by it. I take it to be a claim about what does, or

does not, count as a morally good reason in support of a value judgment: the bare fact that one act is an act of killing, while another act is an act of 'merely' letting someone die, is not a morally good reason in support of the judgment that the former is worse than the latter. Of course there may be *other* differences between such acts that are morally significant. For example, the family of an irreversibly comatose hospital patient may want their loved one to be allowed to die, but not killed. Perhaps the reason for their preference is religious. So we have at least one reason to let the patient die rather than to kill him—the reason is that the family prefers it that way. This does not mean, however, that the distinction between killing and letting die *itself* is important. What is important is respecting the family's wishes. (It is often right to respect people's wishes even if we think those wishes are based on false beliefs.) In another sort of case, a patient with a painful terminal illness may want to be killed rather than allowed to die because a slow, lingering death would be agonizing. Here we have a reason to kill and not let die, but once again the reason is not that one course is intrinsically preferable to the other. The reason is, rather, that the latter course would lead to more suffering.

It should be clear, then, that I will *not* be arguing that every act of letting die is equally as bad as every act of killing. There are lots of reasons why a particular act of killing may be morally worse than a particular act of letting die, or vice versa. If a healthy person is murdered, from a malicious motive, while a person in irreversible coma is allowed to die upon a calm judgment that maintaining him alive is pointless, certainly this killing is very much worse than this letting die. Similarly, if an ill person who could be saved is maliciously allowed to die, while a terminal patient is killed, upon his request, as an act of kindness, we have good reason to judge the letting die worse than the killing. All that I want to argue is that, whatever reasons there may be for judging one act worse than another, the simple fact that one is killing, whereas the other is only letting die, is not among them.

The first stage of the argument is concerned with some formal relations between moral judgments and the reasons that support them. I take it to be a point of logic that moral judgments are true only if good reasons support them; for example, if there is no good reason why you ought to do some action, it cannot be true that you ought to do it. Moreover, when there is a choice to be made from among several possible actions, the preferable alternative is the one that is backed by the strongest reasons.

But when are the reasons for or against one act stronger than those for or against another act? A complete answer would have to include some normative theory explaining why some reasons are intrinsically weightier than others. Suppose you are in a situation in which you can save someone's life only by lying: the normative theory would explain why 'Doing A would save someone's life' is a stronger reason in favour of doing A than 'Doing B would be telling the truth' is in favour of doing B.

However, there are also some purely formal principles that operate here. The simplest and least controversial such principle is this:

- (i) If there are the *same* reasons for or against A as for or against B, then the reasons in favour of A are neither stronger nor weaker than the reasons in favour of B; and so A and B are morally equivalent—neither is preferable to the other.

Now, suppose we ask why killing is morally objectionable. When someone is killed, there may of course be harmful effects for people other than the victim himself. Those who loved him may grieve, and those who were depending on him in one way or another may be caused hardship because, being dead, he will be unable to perform as expected. However, we cannot explain the wrongness of killing purely, or even mainly, in terms of the bad effects for the survivors. The primary reason why killing is wrong is that something very bad is done to the victim himself: he ends up dead; he no longer has a good—his life—which he possessed before. But notice that exactly the same can be said about letting someone die. The primary reason why it is morally objectionable to let someone die, when we could save him, is that he ends up dead; he no longer has a good—his life—which he possessed before. Secondary reasons again have to do with harmful effects on those who survive. Thus, the explanation of why killing is bad mentions features of killing that are also features of letting die, and vice versa. Since there are no comparably general reasons in favour of either, this suggests that:

- (ii) There are the same reasons for and against letting die as for and against killing.

And if this is true, we get the conclusion:

- (iii) Therefore, killing and letting die are morally equivalent—neither is preferable to the other.

The central idea of this argument is that there is no morally relevant difference between killing and letting die, that is, no difference which may be cited to show that one is worse than the other. The argument therefore contains a premise—(ii)—that is supported only inductively. The fact that the explanation of why killing is wrong applies equally well to letting die, and vice versa, provides strong evidence that the inductive generalization is true. Nevertheless, no matter how carefully we analyse the matter, it will always be possible that there is some subtle, morally relevant difference between the two that we have overlooked. In fact, philosophers who believe that killing is worse than letting die have sometimes tried to identify such differences. I believe that these attempts have failed; here are three examples.

1. The first is one that I have already mentioned. Trammell urges that there is an important difference in the 'dischargeability' of duties not to kill and not to let die. We can completely discharge a duty not to kill anyone; but we cannot completely discharge a duty to save everyone who needs aid. This is obviously correct, but it does not show that the Equivalence Thesis is false, for two reasons. In the first place, the difference in dischargeability only shows that the class of people we have a duty to save is smaller than the class of people we have a duty not to kill. It does not show that our duty with respect to those we *can* save is any less stringent. In the second place, if we *cannot* save someone, and that person dies, then we do not let him die. It is not right to say that I let Josef Stalin die, for example, since there is no way I could have saved him. So if I cannot save everyone, then neither can I let everyone die.

2. It has also been urged that, in killing someone, we are *doing* something—namely, killing him—whereas, in letting someone die, we are not doing anything. In letting people die of starvation, for example, we only *fail* to do certain things, such as sending food. The difference is between action and inaction; and somehow, this is supposed to make a moral difference.⁷

There are also two difficulties with this suggestion. First, it is misleading to say, without further ado, that in letting someone die we do nothing. For there is one very important thing that we do: we let someone die. 'Letting someone die' is different, in some ways, from other sorts of actions, mainly in that it is an action we perform *by way of* not performing other actions. We may let someone die by way of not feeding him, just as we may insult someone by way of not shaking his hand. (If it is said, 'I didn't do anything; I simply refrained from taking his hand when he offered it', it may be replied 'You did do one thing—you insulted him'.) The distinction between action and inaction is relative to a specification of *what* actions are or are not done. In insulting someone, we may *not* smile, speak, shake hands, and so on—but we *do* insult or snub the person. And in letting someone die, the following may be among the things that are not done: we do not feed the person, we do not give medication, and so on. But the following is among the things that are done: we let him die.

Second, even if letting die were only a case of inaction, why should any moral conclusion follow from *that* fact? It may seem that a significant conclusion follows if we assume that we are not responsible for inactions. However, there is no general correlation between the action-inaction distinction and any sort of moral assessment. We ought to do some things, and we ought not do others, and we can certainly be morally blameworthy for not doing things as well as for doing them—Jack Palance was blame-

⁷ This argument is suggested by Paul Ramsey in *The Patient as Person* (New Haven, Conn.: Yale University Press, 1970), 151.

worthy for not feeding the child. (In many circumstances we are even legally liable for not doing things: tax fraud may involve only 'inaction'—failing to report certain things to the Department of Internal Revenue—but what of it?) Moreover, failing to act can be subject to all the other kinds of moral assessment. Not doing something may, depending on the circumstances, be right, wrong, obligatory, wise, foolish, compassionate, sadistic, and so on. Since there is no general correlation between the action-inaction distinction and *any* of these matters, it is hard to see how anything could be made out of this distinction in the present context.

3. My final example is from Trammell again. He argues that 'optionality' is a morally relevant difference between killing and letting die. The point here is that if we fail to save someone, we leave open the option for someone else to save him; whereas if we kill, the victim is dead and that is that. This point, I think, has little significance. For one thing, while 'optionality' may mark a difference between killing and *failing to save*, it does not mark a comparable difference between killing and *letting die*. If X fails to save Y, it does not follow that Y dies; someone else may come along and save him. But if X lets Y die, it does follow that Y dies; Y is dead and that is that.⁸ When Palance watches the child die, he does not merely fail to save the child; he lets her die. And when we fail to send food to the starving, and they die, we let them die—we do not merely fail to save them.

The importance of 'optionality' in any particular case depends on the actual chances of someone else's saving the person we do not save. Perhaps it is not so bad not to save someone if we know that someone else *will* save him. (Although even here, we do not behave as we ought; for we ought not simply to leave what needs doing to others.) And perhaps it even gets us off the hook a little if there is the *strong chance* that someone else will step in. But in the case of the world's starving, we know very well that no person or group of persons is going to come along tomorrow and save all of them. We know that there are at least some people who will *not* be saved, if we do not save them. So, as an excuse for not giving aid to the starving, the 'optionality' argument is clearly in bad faith. To say of those people, after they are dead, that someone else *might* have saved them, in the very weak sense in which that will be true, does not excuse us at all. The others who *might* have saved them, but did not, are as guilty as we, but that does not diminish our guilt—as I have already remarked, guilt in these cases multiplies, not divides.

⁸ This difference between failing to save and letting die was pointed out by David Sanford in a very helpful paper, 'On Killing and Letting Die', read at the Western Division meeting of the American Philosophical Association, in New Orleans, on 30 April 1976.

III

I need now to say a few more things about the counter-intuitive nature of the Equivalence Thesis.

The fact that this view has radical implications for conduct has been cited as a reason for rejecting it. Trammell complains that 'Denial of the distinction between negative and positive duties leads straight to an ethic so strenuous that it might give pause even to a philosophical John the Baptist'.⁹ Suppose John is about to buy a phonograph record, purely for his enjoyment, when he is reminded that with this five dollars a starving person could be fed. On the view I am defending, he ought to give the money to feed the hungry person. This may not seem exceptional until we notice that the reasoning is reiterable. Having given the first five dollars, John is not free to use another five to buy the record. For the poor are always with him: there is always *another* starving person to be fed, and then another, and then another. 'The problem', Trammell says, 'is that, even though fulfilment of one particular act of aid involves only minimal effort, it sets a precedent for millions of such efforts.'¹⁰ So we reach the bizarre conclusion that it is almost always immoral to buy phonograph records! And the same goes for fancy clothes, cars, toys, and so on.

This sort of *reductio* argument is of course familiar in philosophy. Such arguments may be divided into three categories. The strongest sort shows that a theory entails a contradiction, and, since contradictions cannot be tolerated, the theory must be modified or rejected. Such arguments, when valid, are of course devastating. Second, an argument may show that a theory has a consequence which, while not inconsistent, is nevertheless demonstrably false—that is, an independent proof can be given that the offensive consequence is unacceptable. Arguments of this second type, while not quite so impressive as the first, can still be irresistible. The third type of *reductio* is markedly weaker than the others. Here, it is merely urged that some consequence of a theory is counter-intuitive. The supposedly embarrassing consequence is perfectly consistent, and there is no proof that it is false; the complaint is only that it goes against our unreflective, pretheoretical beliefs. Now sometimes even this weak sort of argument can be effective, especially when we have not much confidence in the theory, or when our confidence in the pretheoretical belief is unaffected by the reasoning which supports the theory. However, it may happen that *the same reasoning which leads one to accept a theory also persuades one that the pretheoretical beliefs were wrong*. (If this did not happen, philosophy would always be in the service of what we already think; it could never challenge and change our beliefs, and would be, in an

⁹ Trammell, 133.

¹⁰ Trammell, 134.

important sense, useless.) The present case, it seems to me, is an instance of this type. The same reasoning which leads to the view that we are as wicked as Jack Palance, and that killing is no worse than letting die, also persuades (me, at least) that the prereflective belief in the rightness of our affluent life-style is mistaken.¹¹

So, I want to say about all this what H. P. Grice once said at a conference when someone objected that his theory of meaning had an unacceptable implication. Referring to the supposedly embarrassing consequence, Grice said, 'See here, that's not an *objection* to my theory—*that's* my theory!'¹² Grice not only accepted the implication, he claimed it as an integral part of what he wanted to say. Similarly, the realization that we are morally wrong to spend money on inessentials, when that money could go to feed the starving, is an integral part of the view I am defending. It is not an embarrassing consequence of the view; it is (part of) the view itself.

There is another way in which the counter-intuitive nature of the Equivalence Thesis may be brought out. It follows from that thesis that if the *only* difference between a pair of acts is that one is killing, while the other is letting die, those actions are equally good or bad—neither is preferable to the other. Defenders of the distinction between positive and negative duties have pointed out that in such cases our intuitions often tell us just the opposite: killing seems obviously worse. Here is an example produced by Daniel Dinello:

Jones and Smith are in a hospital. Jones cannot live longer than two hours unless he gets a heart transplant. Smith, who has had one kidney removed, is dying of an infection in the other kidney. If he does not get a kidney transplant, he will die in about four hours. When Jones dies, his one good kidney can be transplanted to Smith, or Smith could be killed and his heart transplanted to Jones . . . it seems clear that it would, in fact, be wrong to kill Smith and save Jones, rather than letting Jones die and saving Smith.¹³

And another from Trammell:

If someone threatened to steal \$1000 from a person if he did not take a gun and shoot a stranger between the eyes, it would be very wrong for

¹¹ There is also some independent evidence that this prereflective belief is mistaken; see Singer, 'Famine, Affluence, and Morality'.

¹² Grice made this remark several years ago at Oberlin. I do not remember the surrounding details of the discussion, but the remark seems to me an important one which applies to lots of 'objections' to various theories. The most famous objections to act-utilitarianism, for example, are little more than descriptions of the theory, with the question-begging addendum, 'Because it says *that*, it can't be right'.

¹³ Daniel Dinello, 'On Killing and Letting Die', *Analysis* 31 No. 3 (January 1971), 85–86.

him to kill the stranger to save his \$1000. But if someone asked from that person \$1000 to save a stranger, it would seem that his obligation to grant this request would not be as great as his obligation to refuse the first demand—even if he has good reason for believing that without his \$1000 the stranger would certainly die. . . . In this particular example, it seems plausible to say that a person has a greater obligation to refrain from killing someone, even though the effort required of him (\$1000) and his motivation toward the stranger be assumed identical in both cases.¹⁴

The conclusion we are invited to draw from these examples is that, contrary to what I have been arguing, the bare difference between killing and letting die *must be* morally significant.

Now Dinello's example is badly flawed, since the choice before the doctor is not a choice between killing and letting die at all. If the doctor kills Smith in order to transplant his heart to Jones, he will have killed Smith. But if he waits until Jones dies, and then transfers the kidney to Smith, he will *not* have 'let Jones die'. The reason is connected with the fact that not every case of not saving someone is a case of letting him die. (Josef Stalin died, and I did not save him, but I did not let Stalin die.) Dinello himself points out that, in order for it to be true that X lets Y die, X must be 'in a position' to save Y, but not do so.¹⁵ (I was never in a position to save Stalin.) Now the doctor is in a position to save Jones only if there is a heart available for transplantation. But no such heart is available—Smith's heart, for example, is not available since Smith is still using it. Therefore, since the doctor is not in a position to save Jones, he does not let Jones die.¹⁶

Trammell's example is not quite so easy to dismiss. Initially, I share the intuition that it would be worse to kill someone to prevent \$1000 from being stolen than to refuse to pay \$1000 to save someone. Yet on reflection I have not much confidence in this feeling. What is at stake in the situation described is the person's \$1000 and the stranger's life. But we end up with the *same* combination of lives and money, no matter which option the person chooses: if he shoots the stranger, the stranger dies and he keeps his \$1000; and if he refuses to pay to save the stranger, the stranger dies and he keeps his \$1000. It makes no difference, either to the person's interests or to the stranger's interests, which option is chosen; why, then, do we have the curious intuition that there is a big difference here?

¹⁴ Trammell, 131.

¹⁵ Dinello, 85.

¹⁶ There is another way to meet Dinello's counter-example. A surprisingly strong case can be made that it would *not* be any worse to kill Smith than to 'let Jones die'. I have in mind adapting John Harris's argument in 'The Survival Lottery', *Philosophy* 50 (1975), 81–87.

I conceded at the outset that most of us believe that in letting people die we are not behaving as badly as if we were to kill them. I think I have given good reasons for concluding that this belief is false. Yet giving reasons is often not enough, even in philosophy. For if an intuition is strong enough, we may continue to rely on it and assume that *something* is wrong with the arguments opposing it, even though we are not sure exactly what is wrong. It is a familiar remark: 'X is more certain than any argument that might be given against it'. So in addition to the arguments, we need some account of why people have the allegedly mistaken intuition and why it is so persistent. Why do people believe so firmly that killing is so much worse than letting die, both in fictitious cases such as Trammell's, and in the famine relief cases in the real world? In some ways the explanation of this is best left to the psychologists; the distinctly philosophical job is accomplished when the intuition is shown to be false. However, I shall hazard a hypothesis, since it shows how our intuitions can be explained without assuming that they are perceptions of the truth.

Human beings are to some degree altruistic, but they are also to a great degree selfish, and their attitudes on matters of conduct are largely determined by what is in their own interests, and what is in the interests of the few other people they especially care about. In terms of both the costs and the benefits, it is to their own advantage for people in the affluent countries to regard killing as worse than letting die. First, the *costs* of never killing anyone are not great: we can live very well without ever killing. But the cost of not allowing people to die, when we could save them, would be very great. For any one of us to take seriously a duty to save the starving would require that we give up our affluent life-styles; money could no longer be spent on luxuries while others starve. On the other side, we have much more to *gain* from a strict prohibition on killing than from a like prohibition on letting die. Since we are not in danger of starving, we will not suffer if people do not regard feeding the hungry as so important; but we would be threatened if people did not regard killing as very, very bad. So, both the costs and the benefits encourage us, selfishly, to view killing as worse than letting die. It is to our own advantage to believe this, and so we do.

University of Alabama in Birmingham