

Penultimate draft. For published version see *The Blackwell Companion to Applied Ethics*, 2002.

Killing and Letting Die

Alastair Norcross

Is there a morally significant distinction between killing and letting die? In particular, is killing worse, in itself, than letting die? This question has received much attention in recent years, in large part due to the perceived connection with the moral issues surrounding euthanasia. A legal and medical consensus has gradually emerged in the United States that "termination of life support is legitimate under certain circumstances". (Meisel) The consensus, which is grounded in the right of both competent and incompetent patients to refuse treatment, also recognizes a clear distinction between active euthanasia and assisted suicide on the one hand, and forgoing life-sustaining treatment, sometimes referred to as "passive euthanasia", on the other. This distinction, often identified with the distinction

between killing and letting die, is seen as a vital element in maintaining the consensus. Though the legal consensus is that there is a sharp distinction between active and passive euthanasia, there is by no means a public consensus on this point. The activities of Doctor Jack Kevorkian, who has assisted on numerous suicides in recent years, command as much public admiration as condemnation.

Before I examine the philosophical debate over killing and letting die, I should set aside a possible complication. The distinction between killing and letting die appears to be a specific case of the more general distinction between doing harm and allowing harm. But this is not quite right. Most cases of killing involve doing harm, and most cases of letting die involve allowing harm. However, there are cases in which death is not a harm, and therefore in which killing does not involve doing harm. These are precisely those cases in which the moral case for euthanasia, either active or passive, is

strongest. When continued life involves overwhelming suffering, death may be a benefit to the sufferer. It follows that even if doing harm is morally worse, in itself, than allowing harm, active euthanasia may not differ morally from passive euthanasia. In what follows, however, I shall ignore this complication, and treat killing as a specific instance of doing harm.

Philosophical debate on the supposed moral difference between killing and letting die has flourished in recent years. In one of the classic attacks on the distinction, James Rachels claims that the American Medical Association policy statement on euthanasia endorses the doctrine that there is an important moral difference between killing and letting die. He further argues that there is no such difference, and thus that the AMA policy is seriously flawed. Rachels' main argument consists of the following pair of cases:

In the first, Smith stands to gain a large inheritance if anything should happen to his six-year-old cousin. One evening while the child is taking his bath, Smith sneaks into the bathroom and drowns the child, and then arranges things so that it will look like an accident.

In the second, Jones also stands to gain if anything should happen to his six-year-old cousin. Like Smith, Jones sneaks in planning to drown the child in his bath. However, just as he enters the bathroom Jones sees the child slip and hit his head, and fall face down in the water. Jones is delighted; he stands by, ready to push the child's head back under if it is necessary, but it is not necessary. With only a little thrashing about, the child drowns all by himself, "accidentally," as Jones watches and does nothing. (Rachels, 115)

The two cases differ only in that Smith kills his cousin while Jones "merely" lets his cousin die. If the difference between

killing and letting die were morally significant in itself, claims Rachels, Smith's behavior would be morally worse than Jones' behavior. However, the two are equally reprehensible. So, killing is not worse, in itself, than letting die.

Rachels' argument follows a standard format for arguments about the moral status of certain factors: present two pieces of behavior that differ only with respect to the factor in question. If the behavior differs morally, the factor is morally significant in itself. If not, it is morally insignificant. This methodology is by no means uncontroversial, though. It rests on the assumption that, if a factor is morally significant in itself, it will be significant wherever it appears. Against this assumption, some have argued that a factor, such as the difference between killing and letting die, could sometimes be morally significant, and sometimes not. (The claim is not that other factors could sometimes cancel out the influence of the factor in question, but that the factor's influence could vary in

cases where *all other factors are equal*.) For example, Frances Kamm claims that the following properties are "conceptual components" of letting die, and are true of some, but not all, cases of killing:

(b') Letting die does not create a threat

(c') Victim is already under a threat

(d') Victim loses what he would have had via the agent in question (killer or non-saver)

(f) The agent's efforts would provide the victim with continued life. (Kamm, 1983, 301)

Furthermore, she claims, these definitional properties of letting die might be morally significant, in that "introducing a property conceptually true of letting die into a case of killing might make the particular killing more easily justified than killing in a case which lacked the property." (Kamm, *Ibid.*) So some pairs of actions that differ only in that one is a killing and the other a letting die might differ morally, while other such pairs

might not. Thus, killing and letting die might differ morally, even though certain instances of killing and letting die do not differ, even when all other factors are equal. In a different article, Kamm focuses on a combination of factors (d') and (f')

letting die, by definition, involves someone losing *only* life he would have had by virtue of my assistance... My acting to cause death will be different, in a morally significant way, from my letting die by inaction when the act causes someone to lose life he would have had independently of my assistance (though he may have it by virtue of the assistance of others).
(Kamm, 1992, 370)

So how does it follow that killing is worse, in itself, than letting die?

Suppose letting die had a property true of it by definition which did make cases less morally objectionable, while killing did not have a definitional property with the same

effect. I suggest that, other things equal between them, this would indicate that killing and letting die differ morally *per se*. (Kamm, 1983, 301.)

This is a difficult claim to assess. Consider the following case for comparison: there are two classes of murder, first-degree and second-degree. Second-degree murder is less serious, both legally and morally, than first-degree murder, because it lacks premeditation, which first-degree murder includes. Do murder and second-degree murder differ morally and legally *per se*?

The question appears to be ill formed. The correct comparison with second-degree murder in this case is not simply murder, but first-degree murder. The factor with moral significance is not the difference between second-degree murder and murder, but rather the difference between premeditated murder and unpremeditated murder. It is the presence or absence of premeditation that has significance. Likewise, when we consider Kamm's argument, the appropriate claim should be

that the factors with moral significance are the four (or one) that she claims to be “conceptual components” of letting die, not the difference between killing and letting die itself. It is, furthermore, not at all clear why the factors she cites should be thought to have moral significance.

Another possible objection to Rachels' example is that it may not be possible to generalize from it to the kinds of cases that most concern us. The example involves two pieces of malevolent behavior. Perhaps there is no moral difference between malevolently killing and malevolently letting die. But we can't necessarily generalize from this to cases of euthanasia, which involve benevolent motivations. Perhaps benevolent killing is worse than benevolent letting die. Of course, this kind of objection can always be raised against an argument from a matched pair of examples. Whatever the distinction we are trying to isolate, if the examples differ only with respect to that distinction, there will be any number of

other factors held equal that could have been different. Merely identifying one such factor poses no real threat to the original argument. Rachels' examples are both in bathrooms, but it would be ridiculous to suggest that we couldn't generalize to killings and lettings die in other locations. It might be objected that the difference between malevolent and benevolent motivation is intuitively more morally relevant than the distinction between one location and another. Maybe so, but what is intuitively (though not uncontroversially) relevant is the distinction between a piece of benevolent behavior and a piece of malevolent behavior. It is not intuitively obvious that the distinction between a distinction involving malevolent behavior and a distinction involving benevolent behavior is morally relevant. At the very least, the objector to Rachels' argument owes an explanation of why such a distinction should make a moral difference between distinctions.

The preceding discussion illustrates some important methodological points concerning arguments for the moral relevance or irrelevance of distinctions. It is always possible to object to an argument from a pair of examples that a particular distinction is sometimes morally relevant and sometimes not. The objection by itself, though, has no force. It must be accompanied by an explanation of why the distinction is morally relevant in some cases and not in others. This is much harder to provide. At the very least, then, Rachels' argument puts the burden of proof squarely on those who would invest the distinction between killing and letting die with moral significance.

So far, I have been treating Rachels' two cases as morally equivalent. The vast majority of moral intuitions certainly judge them to be equivalent. But moral intuitions are by no means infallible guides to moral reality. Perhaps Smith's behavior really is worse than Jones's behavior, but, since both

behave so badly, our moral intuitions are not sensitive to the difference. The difference between the two cases is 'drowned out' by the immorality of both of them. While the discussion focuses on the distinction between killing and letting die in particular, this problem is hard to avoid. Examples involving death are apt to be morally charged.

Partly in response to this kind of worry, Jonathan Bennett focuses on a more general distinction, which he illustrates with a morally neutral example. Bennett labels his distinction the "making/allowing" distinction, but he doesn't claim that it always corresponds to the ordinary usage of those terms. He centers his discussion on a pair of cases involving an unoccupied car rolling off a cliff. In one case, 'Push', an agent gives a stationary car a push, 'making' it roll off the cliff to its destruction. In the other case, 'Stayback', the car is already rolling towards the edge of the cliff. The agent could prevent it from falling, by placing a rock in its path, but she doesn't, thus

‘allowing’ it to be destroyed. An agent makes an upshot come about, if her behavior is ‘positively relevant’ to the upshot; she allows it to come about, if her behavior is ‘negatively relevant’ to it. The positive/negative distinction ‘does not distinguish two kinds of action: there are no negative actions’. (Bennett, 232) It is, rather, a distinction between two kinds of propositions, and thus also facts. The propositions in question concern how an agent behaves at a time, specifically how she moves her body. A negative proposition about an agent's behavior at time t is a very uninformative proposition, in that it doesn't tell us very much about her behavior, but merely excludes a few possibilities. That the agent didn't place the rock in the car's path at t tells us very little about what she did do at t . It merely excludes the relatively small proportion of possible bodily movements that would have got the rock in the path of the car. She could have been doing any number of other things instead, such as dancing a jig, picking daisies, or

sitting stock still. A positive proposition about an agent's behavior at t , on the other hand, is a very informative proposition, in that it rules out most of the possible bodily movements she could have been making. That the agent pushed the car at t tells us a great deal about her behavior, ruling out most of her possible bodily movements. We can now see what it is for an agent's conduct to be positively or negatively relevant to a particular upshot. Suppose that an upshot U comes about at t_2 , and furthermore that an agent, called Agent, could have behaved so that it didn't come about. Let t_1 be the latest time at which Agent could have behaved differently so that U would not have occurred. What we look for is the weakest fact A about Agent's conduct at t_1 which, when added to a description of Agent's environment, yields a complete causal explanation of U . If A is negative, Agent's conduct at t_1 is negatively relevant to U , that is, she allows U

to occur. If A is positive, her conduct is positively relevant to U, that is, she makes U occur.

When Bennett first presented his distinction, he intended it as an analysis of the distinction between killing and letting die. In fact, the distinction fits the vast majority of cases of killing and letting die. However, it is a fairly easy matter to concoct outlandish examples that seem to involve killing, but are judged otherwise by Bennett's account. In the light of such examples, Bennett no longer claims to have provided an analysis of our ordinary uses of 'kill' and 'let die', or, more broadly, 'make' and 'allow'. His purpose is to present a distinction that underlies much of our thinking on these matters, and to show that it cannot be of moral relevance. Few, if any, would dispute his claim that *his* distinction cannot bear any moral weight. Those who think that there is a morally significant distinction between making and allowing must come up with an alternative account. Bennett himself presents

what he considers to be the most considerable rival to his account, provided by Alan Donagan, and argues both that the two accounts collaborate in shaping our intuitions, and that alternative accounts, when viable, are best understood as versions of Donagan's account.

Donagan defines an action as 'a deed done in a particular situation or set of circumstances;... [consisting] partly of [the agent's] own bodily and mental states'. (Donagan, 42) He continues:

Should he be deprived of all power of action, the situation, including his bodily and mental states, would change according to the laws of nature. His deeds as an agent are either interventions in that natural process or abstentions from intervention. When he intervenes, he can be described as causing whatever would not have occurred had he abstained; and when he abstains, as allowing to happen whatever would not have happened had he intervened.

Hence, from the point of view of action, the situation is conceived as passive, and the agent, *qua* agent, as external to it. He is like a *deus ex machina* whose interventions make a difference to what otherwise would naturally come about without them. (Donagan, 42-3)

In considering what would have happened if Agent hadn't acted, Donagan asks what would have happened in 'the course of nature' (his phrase). The course of nature can include not only Agent's physical presence, but also changes in her 'bodily and mental states'. It is the exercise of human agency that gives Agent the option to intervene in the course of nature or to allow nature to take its course. All of Agent's deeds are either interventions or abstentions. Those that make a difference to the course of nature, or what would have happened anyway, are interventions; those that leave the course of nature unchanged are abstentions.

Talk of the course of nature, or letting nature take its

course, clearly underlies a good deal of unreflective thinking about morality. It is, no doubt, connected to the equally unreflective notion of 'playing god'. When confronted with the familiar example of a runaway train heading towards five people trapped on the track, some people will refuse to countenance switching the train to a sidetrack with only one person, on the grounds that they shouldn't interfere with the natural course of events, or that they shouldn't 'play god'. (The same people show remarkably little reluctance to interfere with the natural course of a life-threatening illness, by approving the use of antibiotics, for example.) Such an attitude embodies an assumption that we have a well-defined notion of the course of nature, or that we can know what god plans. Presumably, god was somehow instrumental in sending the train hurtling towards the five people trapped on the track. If we switch the train to the other track, we will have usurped god's unique role in making life and death decisions. It rarely

occurs to proponents of this line of reasoning that an omnipotent being might choose to act through the decisions of moral agents. At least investing the course of nature with moral significance doesn't involve the complications inherent in postulating a supernatural being whose wishes are obscure. It does, however, require an account that can stand up to critical scrutiny. Donagan defines the course of nature, with respect to an agent's action at a time, in terms of the exercise of human agency. Something happens in the course of nature, just in case it would happen if the agent didn't exercise her agency. However, it is not clear just how we are to evaluate the antecedent of a counterfactual of the form 'if Agent hadn't exercised her agency, x would have occurred'. Donagan suggests that the course of nature can include changes in the agent's bodily and mental states. But what kind of changes occur in the course of nature, as opposed to as a result of the exercise of human agency? Presumably physical reflexes can

operate in the course of nature. If a fly swoops towards my head, my hand can move to brush it away in the course of nature. I can also form beliefs in the course of nature, for example that a fly is swooping towards my head. Could I save or end someone's life in the course of nature? Consider a bigoted policeman challenging a black suspect in the street. The suspect reaches into an inside pocket to produce his wallet. The policeman shoots him before he can produce the wallet. A bigoted policeman's shooting a black suspect who reaches for an inside pocket could be just as much a reflex action as my brushing away a fly. On Donagan's account, then, if that policeman exercises his agency to shoot the suspect, he merely allows him to die. A different policeman, who would not reflexively shoot, would be judged to have killed the suspect in those circumstances. If harming others is a deeply entrenched part of your character, you merely allow them to be harmed on those occasions when the harm results from an exercise of

agency. Conversely, if helping others is a deeply entrenched part of your character, you actively harm them on those occasions when you decide not to help. If doing harm is worse than allowing harm, this gives a curious result. In certain circumstances, the better your character, the worse your actions, and the worse your character, the better your actions. Any account of the course of nature that is to do serious moral work will have to avoid consequences such as these.

Even if the notion of the course of nature were to be given a satisfactory philosophical grounding, it is not clear how it can carry any moral weight. Bennett admits that Donagan's intervention/abstention distinction is not so obviously devoid of moral significance as his own making/allowing distinction. However, it is hard to see how it *can* carry moral weight. If my behavior leads to a harmful upshot, and I could have behaved so that the upshot did not occur, what difference does it make whether the upshot would have occurred if I hadn't exercised

my agency at all? Bennett therefore places the burden of proof on those who would invest Donagan's distinction with moral weight. If we are to show that the difference between making and allowing is morally significant, we must either explain why Donagan's distinction carries moral weight, or provide a different account of the distinction and an explanation of how *that* account is morally relevant.

Philippa Foot attempts to invest the killing/letting die distinction with moral significance by presenting a version that appears to be a variation of Donagan's. She presents a distinction that focuses on the question of whether someone is 'the agent' of harm to someone else. When the harm in question is death, this distinction corresponds roughly to the killing/letting die distinction. She illustrates this with the following two stories: In Rescue I we can save either five people in danger of drowning in one place or one person drowning somewhere else. In Rescue II, we can save the five

drowning people only by driving over and killing someone who is trapped on the road. In Rescue I we act permissibly if we save the five, even though the one dies as a result. We let the one die. In Rescue II we do not act permissibly if we save the five. The only way to save the five involves killing the one. The distinction, according to Foot, is between originating or sustaining a fatal sequence on the one hand, and allowing such a sequence to run its course on the other. It is often permissible, she claims, to bring about a harm by the latter method that could not permissibly be brought about by the former. What explains this moral difference? The different types of agency receive their moral significance via their connection with different types of right:

For there are rights to noninterference, which form one class of rights; and there are also rights to goods or services, which are different... Typically, it takes more to justify an interference than to justify the withholding

of goods or services. (Foot, 284)

Originating or sustaining a harmful sequence will usually involve the violation of a right to noninterference, whereas allowing such a sequence to run its course will, at most, involve the violation of a right to goods or services. The former type of right is stronger than the latter, so the former type of agency is less likely to be permissible than the latter. Notice that this way of arguing for the moral significance of the distinction does not imply that the distinction always matters morally. There may be circumstances in which a particular harmful result cannot permissibly be brought about either by interference or by withholding aid. Even if a right to noninterference is stricter than a right to be given aid, it will often be the case that neither right can permissibly be violated.

Foot's claim about the relative strictness of positive and negative rights and duties has a good deal of intuitive support. My right not to be poisoned does seem stronger than my right,

if any, to be given the food I need to survive. However, as an explanation of a morally significant difference between killing and letting die, this appeal to different types of rights simply diverts the question. If the moral difference between positive and negative rights is to provide a satisfactory explanation of the moral difference between killing and letting die, we also need an explanation of the former difference. *Why* is my right not to be poisoned stronger than my right to be given the food I need to survive? The answer that springs most readily to mind is that it is worse to kill me than to let me die. But this can be of no help to Foot's approach, since it merely takes us in a circle.

There is a further problem with any approach, like Foot's, that appeals to the notion of sequences. What is it for a sequence to be already in motion? Consider Foot's Rescue I. If we save the five, we supposedly allow a fatal sequence that is already threatening the one to run its course. But in what

sense is there a fatal sequence already threatening the one? To be sure, if I don't save him, he will die. But consider the one in my path in Rescue II. If I don't stop or divert my car before I reach him, he will also die. Is there a fatal sequence already threatening him? Perhaps so. Now consider a pedestrian by the side of the road. If I swerve towards him, I will kill him. So, if I don't avoid him, he will die. Is there therefore a fatal sequence already threatening *him*? It appears that there are fatal sequences already threatening any potential victim of either a killing or a letting die. What determines, then, whether the victim is killed or let die? Recall Foot's answer. I kill someone, if I *initiate* or *sustain* a fatal sequence. I let someone die, if I *allow* a fatal sequence to run its course. What distinguishes between initiating or sustaining, on the one hand, and allowing on the other? If I run over the one in Rescue II, he dies because I hit him with the car. If I don't run over the one, the five die because I don't rescue them. Perhaps, then, I

initiate or sustain a fatal sequence when someone dies because I do something. I allow a fatal sequence to run its course when someone dies because I don't do something. This explanation would clearly render the notion of a sequence redundant, and reduce Foot's account to a variation on Bennett's. However, consider a variation on Rescue II. I am the dispatcher who responds to the call for help from the five. I learn that one of my rescuers, Jones, is already on the way. However, I know that Jones, a utilitarian, will not hesitate to run over someone trapped on the path in order to get to the five. I also know that another rescuer, Smith, is closer. Smith is a strict deontologist who will not run over someone trapped on the path, even to save a hundred. Knowing this, I send Smith, and call off Jones. Smith refuses to run over the one on the path, with the result that the five die. The five died because I called off Jones. The relevance of my behavior to their deaths is, as Bennett would put it, positive. But have I initiated or sustained a fatal

sequence, or rather 'merely' allowed one to run its course? Intuitively, I have allowed the five to die, and perhaps also allowed a fatal sequence to run its course.

One response to this kind of complication is to declare that the ordinary uses of such terms as 'kill', 'let die', and 'allow', are pretty poor guides to any well-grounded distinctions that could do any useful moral work. This is Bennett's response, and I suspect he is right. Others, however, could respond by introducing greater degrees of complication into their accounts. Matthew Hanser, for example, argues that we should supplement the categories of killing and letting die with a third – preventing people from being saved. This, says Hanser, is a species of doing harm, like killing, but is morally on a par with letting die. Hanser's explanation for the moral difference between killing, on the one hand, and both letting die and preventing people from being saved, on the other, appeals to different kinds of rights. Specifically, killing

involves the violation of the 'right to continued life', whereas both letting die and preventing people from being saved involve the violation of the 'right to life-saving aid'. The former right is more stringent than the latter right, because 'other things being equal, an unvalenced right to a "basic" good is more stringent than an unvalenced right to a good generated by that basic good.' (Hanser, 294) Although this is an interesting suggestion, it is difficult to know how to assess it without an explanation for why the former right is more stringent than the latter right. I suspect that such an explanation will be hard to come by. The distinction between different kinds of right may also prove illusory. For example, consider someone dying of thirst in the desert. If I make a hole in her water canteen, have I violated her right to continued life or her right to life-saving aid? Would it make a difference whether I make the hole while the canteen is full of water, or before the desert rescue team shows up to fill it up?

Warren Quinn presents a version of the distinction between 'harmful positive agency' and 'harmful negative agency', that is intended to carry moral weight, but that turns out to rely on something very close to Bennett's distinction. He states his position as follows:

Harmful positive agency is that in which an agent's most direct contribution to the harm is an action, whether his own or that of some object. Harmful negative agency is that in which the most direct contribution is an inaction, a failure to prevent the harm. (Quinn, 367)

Furthermore, he explains the notion of direct contribution as follows:

An agent's *most direct contribution* to a harmful upshot of his agency is the contribution that most directly explains the harm. And one contribution explains harm more directly than another if the explanatory value of

the second is exhausted in the way it explains the first.
(Quinn, 366)

Quinn's account thus relies on the notions of action and inaction. He is not, of course, claiming that an agent has to be inactive in order to allow something to happen. It is, rather, an agent's inactivity in a particular respect that is important. I might be exceedingly active at the time that I allow some harm to befall you, but my inactivity with respect to preventing the harm is my most direct contribution to the harm. It is my failure to prevent the harm that most directly explains the harm. This sounds very much like Bennett's account. That I failed to prevent a harm at a particular time is most likely a very uninformative fact about me. Quinn's notion of more and less direct explanation could be read in terms of Bennett's stress on the weakest fact necessary to explain an upshot. Although such a reading of Quinn seems to make a lot of sense on its own terms, it is clearly not what Quinn intended. He

explicitly rejects Bennett's approach in terms of facts about behavior. On the other hand, it is hard to see how to make Quinn's approach work without recourse to facts. If we imagine that an agent's behavior consists of actions and inactions, we need to be able to explain the difference between the two. Ordinary language won't be of any help. We can just as easily describe my failure to save the one in Rescue I as my consigning him to his fate. Is this piece of behavior an action or an inaction? Likewise, my driving over the one in Rescue II could be described as my failure to avoid him. We might employ a theory of action according to which my failure to save the one and my consigning him to his fate are actually two different pieces of behavior, the one an inaction and the other an action. This won't help Quinn's analysis, though, since the inaction in this case doesn't explain the death of the one any more directly than does the action. If Quinn's account is best understood as a version of Bennett's account, it clearly can't

explain how harmful positive agency is morally worse than harmful negative agency. Even if it is not a version of Bennett's account, it is hard to see how to get moral significance from the distinction. Why should it matter whether my most direct contribution to a harm is an action or an inaction?

The ease with which it is possible to invent counterexamples to suggested criteria for the uses of 'kill' and 'let die' may ring warning bells concerning the whole enterprise of supplying such criteria. It could be that there simply is no systematic way to characterize the ordinary uses of these terms, no matter how much complication we are prepared to endure. Part of the problem is that there is a good deal of disagreement about particular cases, much of which arises from the fact that many people are guided in their application of the terms, at least in part, by their moral judgments of the cases themselves. For example, those who

regard the discontinuation of life-support for a dying patient as wrong are more likely to label it 'killing', while those who regard it as permissible are more likely to label it 'letting die'.

Jeff McMahan has a different explanation for the difficulty of the task of defining the distinction between killing and letting die:

Our intuitions about killing and letting die are indeed based on considerations that are relatively simple... But, because of the unruly complexity of reality, it is often difficult to determine what these considerations imply about the classification of a particular case. (McMahan, 402)

And what are these relatively simple criteria?

In short, the fundamental intuitive difference between killing and letting die is that in cases of killing we assign primary causal responsibility for a person's death to an agent's intervention in the person's life, whereas,

in cases of letting die, primary responsibility for the death is attributed to factors other than any intervention by the agent. (McMahan, 411)

This seems very similar to Donagan's active/passive distinction, which is also couched in terms of 'intervention'. And, like Donagan's distinction, it doesn't clearly bear moral weight. As McMahan says, 'it is difficult to believe that the way in which an agent is instrumental in the occurrence of an outcome could be more important than the nature of the outcome itself.' (McMahan, 413) Thus, a deeper understanding of the distinction between killing and letting die may reveal that intuitions that are 'central to any morality that we could bring ourselves to accept' are 'apparently ungrounded'. (McMahan, 413)

What, then, explains the widespread, but erroneous, belief that killing is worse, in itself, than letting die? One important factor concerns the kind of examples the distinction

often brings to mind. Examples of killing often involve unwilling victims and malevolent, or at least callously self-interested, perpetrators. Examples of letting die often involve agents who would have to expend great efforts or take great risks to save the lives of the victims, or victims who would not be benefited by continued life. In many of these common examples, then, it is easy to avoid killing, but costly, either to the agent or to the victim, to avoid letting die. In many of these examples, the agent of a killing intends the death of the victim, either as an end in itself, or as a means to personal gain, whereas the agent of a letting die either doesn't intend the death at all, or intends it for the benefit of the victim. Factors such as these, at best contingently connected with the distinction, are often relevant to the morality of either actions or character.

Another, and perhaps more significant, reason for the widespread belief that killing is worse than letting die concerns the cost of rejecting that belief. As I said, the alternatives to

the commonly discussed examples of letting die often involve great cost either to the victim or the agent. It is risky to run into a burning building to save someone trapped inside. A terminally ill patient may receive no benefit, and sometimes considerable harm, from strenuous efforts to prolong her life. Such cases of letting die may, therefore, involve either excusable wrongdoing or no wrongdoing at all. However, most actual victims of letting die would benefit considerably from being saved, and could be saved with very little effort or risk on the part of those of us who let them die. Millions die every year in all parts of the world, many of them young children, as a direct or indirect result of extreme poverty. Modern relief agencies, such as CARE and UNICEF, have made it very easy for those of us who are even mildly affluent to save significant numbers of them. Most of us do very little to help. If we reject the belief that killing is worse than letting die, it is hard to see how we can judge our behavior as anything

less than abominable. This clearly constitutes a powerful motivation, though not a respectable reason, to believe that the distinction between killing and letting die carries considerable moral weight.

It seems that McMahan is right to claim that the intuition that there is a morally significant difference between killing and letting die is ungrounded. It is less clear that he is correct in his contention that this intuition is central to any morality that we could bring ourselves to accept. However, as I have said, if we reject the moral significance of the distinction between killing and letting die, we will most likely have to accept more than just the permissibility of active euthanasia. If there is no morally significant difference between killing and letting die, it is that much harder to justify our neglect of the underprivileged, both in our own country and abroad. We might well be forced to conclude that most of us who possess even modest resources are seriously at fault for not doing more

to help others. This conclusion would certainly be painful. The unpleasantness of a moral conclusion is, however, neither evidence for its falsity nor even its unacceptability. If we have to choose between a position that is ungrounded and one with painful implications, we should grit our teeth and choose the latter.

Works Cited.

Bennett, Jonathan, 'Negation and Abstention: Two Theories of Allowing' (1993), repr. in Steinbock and Norcross, *Killing and Letting Die*, 230-256.

Donagan, Alan, *The Theory of Morality*, (Chicago: Chicago University Press, 1977).

Foot, Phillipa, 'Killing and Letting Die' (1984), repr. in Steinbock and Norcross, *Killing and Letting Die*, 280-289.

Hanser, Matthew, 'Killing, Letting Die and Preventing People from Being Saved', *Utilitas* Vol. 11, No. 3 (November 1999), 277-295.

Kamm, Frances Myrna, 'Killing and Letting Die: Methodological and Substantive Issues', *Pacific Philosophical Quarterly*, 64 (1983), 297-312.

——— 'Non-consequentialism, the Person as an End-in-Itself, and the Significance of Status', *Philosophy and Public Affairs* (Fall 1992): 354-389.

Meisel, Alan, 'The Legal Consensus About Forgoing Life-Sustaining Treatment: Its Status and Its Prospects', *Kennedy Institute of Ethics Journal* Vol. 2, No. 4, 1992.

McMahan, Jeff, 'Killing, Letting Die, and Withdrawing Aid' (1993), repr. in Steinbock and Norcross, *Killing and Letting Die*, 383-420.

Quinn, Warren, 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing' (1989), repr. in Steinbock and Norcross, *Killing and Letting Die*, 355-382.

Rachels, James, 'Active and Passive Euthanasia' (1975), repr. in Steinbock and Norcross, *Killing and Letting Die*, 112-119.

Steinbock, Bonnie, and Alastair Norcross (eds.), *Killing and Letting Die*, 2nd edn. (New York: Fordham University Press, 1994).

Suggestions for further reading.

Abrams, Natalie, "Active and Passive Euthanasia," *Philosophy*, 53 (1978), 257-63.

Belliotti, Raymond, "Negative Duties, Positive Duties, and Rights," *The Southern Journal of Philosophy*, 14 (1978), 581-88.

Bennett, Jonathan, "Acting and Refraining," *Analysis*, 28 (1967), 30-31.

—, "Shooting, Killing, Dying," *Canadian Journal of Philosophy*, 2 (1973), 315-23.

—, "Morality and Consequences," *The Tanner Lectures on Human Values*, vol. II, S. McMurrin ed., Salt Lake City, UT: University of Utah Press, 1981, 47-116.

—, "Positive and Negative Relevance," *American Philosophical Quarterly*, 20 (1983), 185-94.

— *The Act Itself*, Oxford: Oxford University Press, 1995

Casey, John, "Killing and Letting Die: A Reply to Bennett," *Killing and Letting Die*, 1st ed., B. Steinbock, ed., Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1980.

Chandler, John, "Killing and Letting Die--Putting the Debate in Context," *Australasian Journal of Philosophy*, 68 (1990), 420-31.

Conway, David, "Is Failing to Save Lives as Bad as Killing?," *Journal of Applied Philosophy*, 5 (1988), 109-112.

Glover, Jonathan, 'It Makes No Difference Whether or Not I Do It,' *Proceedings of the Aristotelian Society*, 49 (1975), 171-90.

—, *Causing Death and Saving Lives*. London: Penguin Books, 1977.

Hanink, J.G., "On the Survival Lottery," *Philosophy*, 51 (1976), 223-25.

Harris, John, 'The Survival Lottery' (1975), repr. in Steinbock and Norcross, *Killing and Letting Die*, 257-265

—, *Violence and Negative Responsibility*. Routledge and Kegan Paul, 1980.

Kagan, Shelly, 'The Additive Fallacy', *Ethics* 99 (1988)

—, *The Limits of Morality*, Oxford: Oxford University Press, 1989.

Malm, H. M., 'Killing, Letting Die, and Simple Conflicts,' *Philosophy and Public Affairs*, 18 (1989), 238-58.

—, 'Directions of Justification in the Negative-Positive Duty Debate,' *American Philosophical Quarterly*, 27 (1990), 315-24.

Morillo, Carolyn, 'As Sure as Shooting,' *Philosophy*, 51 (1976), 80-89.

—, 'Doing, Refraining, and the Strenuousness of Morality,' *American Philosophical Quarterly*, 14 (1977), 29-39.

Norcross, Alastair, 'Should Utilitarianism Accommodate Moral Dilemmas?', *Philosophical Studies*, Vol. 79, No. 1, (July 1995), 59-85

Parfit, Derek, *Reasons and Persons*, Oxford: Oxford University Press, 1984.

Singer, Peter, 'Famine, Affluence and Morality,' *Philosophy and Public Affairs*, 1 (1972), 229-43.

—, 'Utility and the Survival Lottery,' *Philosophy*, 52 (1977),
218-22.

Thomson, Judith Jarvis, 'Killing, Letting Die, and the Trolley Problem', reprinted in *Rights, Restitution, and Risk: Essays in Moral Theory*, W. Parent, ed., Cambridge, MA: Harvard University Press, 1986, 78-93.

—, 'The Trolley Problem,' *ibid.*, 94-116.

—, *The Realm of Rights*, Cambridge, MA: Harvard University Press, 1990.

Tooley, Michael, 'An Irrelevant Consideration: Killing Versus Letting Die', repr. in Steinbock and Norcross, *Killing and Letting Die*, 103-111.

Trammell, Richard, 'Tooley's Moral Symmetry Principle,' *Philosophy and Public Affairs*, 5 (1976), 305-13.

Unger, Peter, *Living High and Letting Die: Our Illusion of Innocence*, Oxford University Press, 1996.