

WHERE'S THE HARM IN DYING?

MATTHEW HANSER

The University of California, Santa Barbara

It is natural to suppose that at least part of the reason why killing a human being is typically so morally objectionable is that the victim typically suffers such serious harm. Jeff McMahan thus spends the second chapter of his admirable and long awaited book, *The Ethics of Killing*, considering why and to what extent death is bad for, a misfortune for, or comes as a harm to, the one who dies (he uses these formulations interchangeably). The question is not, he stresses, why it is bad for us that we aren't immortal. The question is why and to what extent the actual event of one's ceasing to be alive comes to one as a harm (p. 103). In answering this question McMahan defends a modified version of what might be called the 'standard' account of death's badness. According to the standard account, someone suffers harm in dying if and only if his life as a whole would have gone better had he not died that particular death. If we ignore, for simplicity, the possibility that the value of an earlier portion of one's life might depend upon how that portion fits into the overall pattern of one's life—a pattern potentially affected by how and when one's life comes to an end—we can say that according to the standard account, someone suffers harm in dying if and only if the extra life he thereby loses would have had positive value for him. The better the extra life would have been, the greater the harm suffered in dying; and if the extra life would have been bad, death comes not as a harm but as a benefit. McMahan modifies this account in two ways. First, he argues that the degree to which someone's death is bad for him is a function not just of the value of the extra life he thereby loses, but also of the strength of the 'prudential unity relations'—psychological continuity and the like—that would have linked his present self to the future self living that extra life. The more weakly he would have been related to that self, the less of a misfortune he suffers in losing that extra life (pp. 105–106). Second, McMahan argues that the more fortunate someone has already been in life, the less bad it is for him to lose extra life in dying (pp. 141–145). Suppose that a twenty-year old and a sixty-year old are both killed in an accident, and that both are thereby deprived of twenty extra years of life. (Owing to an incurable condition, the twenty-year old had a maximum life-span of forty.) They suffer equivalent losses in dying, but the sixty-year old suffers a lesser misfortune, having gained more from the life he's already had.

These modifications to the standard account play important roles in McMahan's larger project. The first modification, for example, helps support his claim that fetuses suffer negligible harm when killed. For present purposes, however, I am concerned with what his account has in common with the standard account. According to both, in order to determine whether someone suffers harm in dying we must compare the value of his actual life

with the value of the life he would have lived had he not died the death he actually did. This comparison gives rise to what McMahan calls *the metaphysical problem*: when asking how good someone's life would have been had he not died the death he actually did, just how much of the causal history leading up to his death are we to imagine being different (pp. 107–109)? Consider McMahan's example of a young man who develops acute leukaemia, which disease ultimately results in his death from a severe haemorrhage. Had that particular haemorrhage not occurred, but had everything leading up to it remained the same, the young man would have died anyway, shortly thereafter, from some other complication of his cancer. Had he not developed leukaemia in the first place, however, he would have lived for many more years. To which alternative history, if either, should we compare his actual history, when determining whether his death was a misfortune? The moral of the example might seem to be that there *is* no single correct comparison. Different comparisons are appropriate in different contexts and for different purposes. This would account for the ambivalence most of us feel when a young person dies from a protracted and painful illness: in such cases death can seem both a great tragedy and a welcome release. McMahan acknowledges that this is a plausible analysis of the leukaemia example but argues that with most deaths such ambivalence would be out of place. In most cases we think there is a single answer to the question whether the person who dies thereby suffers a misfortune. McMahan describes the criteria he thinks we tacitly employ in choosing the appropriate comparison, but he admits that he sees no way to justify the claim that these are the uniquely correct criteria (pp. 110–117).

McMahan's account of death's badness faces another serious problem: it apparently yields incorrect results when death is 'over-determined'. In the present context, death counts as over-determined not when it results from the joint operation of two or more causal factors, each sufficient to produce death on its own, but when death would have resulted from another cause shortly thereafter had the actual cause not operated first. Suppose, for example, that A kills B, and that a group of hit men would have killed him later in the day if A had not done so first. McMahan's account entails that B suffers negligible harm in dying, since he would have lived but a few hours longer had his actual death not occurred. Now consider an alternative scenario in which it is the hit men who kill B, A having abandoned his murderous plan. McMahan's account entails that B suffers great harm in this scenario, since we may suppose he would have lived for many more years had the hit men not killed him. But this pair of results seems unacceptable. Surely B suffers just as serious a harm in the first scenario as he does in the second. The fact that in the first scenario the hit men would have caused B to suffer this harm if A had not done so first does not alter the fact that as events actually unfolded, A caused the harm.

Or so one might think. But McMahan insists that a person *doesn't* suffer serious harm in dying if he would have died soon anyway of some other cause. McMahan's response to those who disagree has two prongs. First, he argues that the claim that death is a serious harm in such cases is indefensible.

In order to get the result that B suffers serious harm in the first scenario, argues McMahan, we must compare B's actual history with an alternative history in which he is killed by neither A nor the hit men. But now suppose that if neither A nor the hit men had killed him, he would have been hit by a bus the following week. If we are still to maintain that B suffers serious harm when A kills him, we must imagine away this potential cause of death as well. We must compare his actual history with a history in which he is killed by neither A nor the hit men nor the bus. But at what point do we stop imagining away alternative causes of death? McMahan calls this *the problem of the terminus*, and he argues that neither of the seemingly plausible answers—that we stop when the alternative life reaches a normal length or that we stop when it reaches the maximum possible length—yields acceptable results in all cases (pp. 120–127). Furthermore, insofar as we *can* make sense of such a 'cumulative' loss (McMahan dubs it the *overall* loss the victim suffers in dying), this loss is not attributable solely to the victim's actual death. The overall loss B suffers in dying is attributable to his actual death *taken together with* the fact that hit men would have killed him if A hadn't *plus* the fact that a bus would have run him over if neither A nor the hit men had killed him *plus* the fact that . . . and so on. The loss properly attributable to his death *alone* is still negligible (pp. 127–128). Death is consequently not a serious harm when it is over-determined.

The second prong of McMahan's response is to explain away any lingering sense we may have that death is a harm in such cases. Perhaps, he suggests, we are mistaking the victim's overall loss for the loss attributable to death alone (pp. 127–128). Or perhaps our desire to say that a person can suffer harm in dying an over-determined death arises from our conviction that acts which produce such deaths are often seriously objectionable (p. 122). In Chapter 3, however, McMahan argues that when the victim is a person (a being with a special sort of worth), the primary objection to killing him (against his will) is not that we thereby harm him but that we thereby fail to *respect* him as we ought (p. 242). The badness of death consequently cannot be inferred from the wrongness of killing.

Despite McMahan's vigorous defence of the view, I remain reluctant to accept that over-determination renders death benign. Yet this result seems unavoidable if we accept some version of the standard account. The question thus arises whether we must accept such an account. McMahan thinks we must. He writes:

If it is right that when we die we cease to exist, it seems to follow that death cannot be bad because of its intrinsic features in the way that, for example, suffering is. For nonexistence has no intrinsic properties, positive or negative. Since the badness of death cannot be intrinsic, it must instead be comparative. Death must be bad by comparison with what it excludes. Thus the central problem in the evaluation of death is understanding what exactly is excluded by death in any particular case. Obviously the alternative to death must be continued life, but what sort of life and of what duration? (p. 98)

In order to understand McMahan's reasoning we must first understand his distinction between intrinsic and comparative badness. I take it that in the present context 'intrinsic' means the same as 'non-comparative', so I shall use the latter expression. Considered in itself, being in pain is bad for one—it is a *non-comparatively* bad state to be in. Having a million dollars, by contrast, is not, in itself, bad; but nor is it as good as having two million. Having one million dollars is thus bad *as compared to having two*—it is a *comparatively* bad state to be in, relative to that alternative. Now let's consider events. One way an event can be bad for someone is in virtue of putting him into a bad state. Being kicked hard in the shin, for example, is bad for one because it puts one into a non-comparatively bad state: it causes one pain. Being cheated out of an extra million dollars is bad for one because it puts one into a comparatively bad state: it causes one to have (let us suppose) just one million dollars instead of two. Using the expressions in *derivative* senses, then, we can say that being kicked hard in the shin is non-comparatively bad for one, insofar as it puts one into a non-comparatively bad state, and that being cheated out of an extra million dollars is comparatively bad for one, insofar as it puts one into a comparatively bad state. But we must remember that these are derivative senses. Strictly speaking it is not the events that are comparatively or non-comparatively bad, but the states that they produce.

Now let us turn our attention to death. The word 'death' can be used to refer either to the *state* of being dead or to the *event* of ceasing to be alive. Unfortunately, these two uses are not clearly distinguished in the passage quoted above. McMahan begins by observing that since non-existence lacks intrinsic properties, death, now in the sense of being dead, cannot be an intrinsically (non-comparatively) bad state; he concludes that death's badness must be comparative. But how is the word 'death' to be understood in this conclusion? McMahan does not, I think, mean to conclude that being dead is a comparatively bad state to be in. If, as he supposes, we cease to exist when we die, then being dead has no value for us at all, positive, negative or neutral. And if being dead has no value for us, then being dead cannot be *worse* for us than being alive. Being dead is neither an intrinsically (non-comparatively) nor a comparatively bad state to be in. McMahan's intended conclusion, then, is presumably that if the *event* of death is bad for one, its badness must be comparative. Does his argument support this conclusion? From the fact that being dead is not intrinsically (non-comparatively) bad, it follows that if the event of death is bad for one, this is not owing to its putting one into a non-comparatively bad state. The event of death cannot be non-comparatively bad for one, in the derivative sense introduced above. But from this it does not follow that death must be *comparatively* bad for one, in the derivative sense introduced above. For perhaps death's badness does not derive from its putting one into a bad *state* at all.

Indeed, the event of death clearly *isn't* comparatively bad, in the sense introduced above, since being dead is not a comparatively bad state to be in. Of course McMahan never claims otherwise. When he concludes that death's badness must be comparative, he seems to be employing the phrase in a somewhat broader, *though still derivative*, sense. An event is comparatively bad,

in this broader sense, if and only if its badness derives from its resulting in something else's being, or having been, comparatively bad. This is the sense in which the standard account takes death to be bad: death is bad for the victim if and only if it results in his having had a comparatively bad *life*—a life worse than the one he would have had if that event had not occurred. But even if we use the phrase in this broader sense, McMahan's argument does not show that death's badness must be comparative. For again, the argument tacitly assumes, without justification, that death's badness must be derivative. It assumes that if death's badness does not derive from its resulting in something else's being intrinsically (non-comparatively) bad, then its badness must derive from its resulting in something else's being comparatively bad. The argument ignores the possibility of death's being bad in its own right, so to speak, quite independently of whether it results in anything else's being bad. To illustrate the sort of possibility I have in mind, I shall, in my remaining space, briefly sketch a 'non-derivative' account of death's badness.

Consider first, as a potential model for death, the harm of losing one's power of sight. The standard account would hold that one suffers harm in losing one's sight if and only if one's life would have gone better had that particular event of sight-loss not occurred. But an alternative analysis is available: loss of sight is a harm because the power of sight is an important human good. (I cannot here pursue the question what makes sight a human good—I can imagine several possible answers.) In one respect, this account is quite similar to the standard account: both claim that one suffers harm in losing one's sight only insofar as one thereby loses something good. According to the standard account, however, the good one loses is a certain quantity of well-being, namely the difference between the higher level of well-being one would have enjoyed had one not been blinded and the lower level of well-being one actually enjoys given that one has been blinded. According to the account I am sketching, by contrast, the good one loses when one loses the power of sight is just that: the power of sight. According to this account, even if losing the power makes one's life go *better* than it would otherwise have gone, the loss, considered in itself, still constitutes a harm.

Now if the power of sight is an important human good, I think it follows that it is bad for a human to lack it. Even if one's life ends up going better than it would have gone had one possessed the power, lacking the power is still, in itself, bad for one. Unlike dying, then, losing one's sight puts one into a bad state. But the badness of lacking the power does not figure into the present account of why one suffers harm in *losing* it. Rather, the power's status as a good explains *both* why it's bad to lack it and why one suffers harm in losing it. Our account of why death is a harm can thus parallel this account of why losing sight is a harm. Death consists in the loss of an assortment of quite basic powers which we may, for convenience, call *vital* powers. (Possession of these powers is in turn a precondition for the possession of all the other human powers.) According to the account I am suggesting, one suffers harm in dying because the vital powers one thereby loses are important human goods.

I leave it an open question whether the vital powers are goods in all possible circumstances, and hence whether death, considered in itself, is always a

harm. (Even if death were always a harm, it could sometimes also bring benefits. I shall say more about this possibility shortly.) Perhaps illness and senility can so far destroy one's 'higher' powers that life no longer has value; and perhaps the vital powers cease being goods when the kind of life they can support permanently ceases to have value. But the vital powers' status as goods is not undermined simply by the fact that there is a contract out on one's life. Over-determination creates no difficulties for this account. Nor is the account vulnerable to McMahan's 'metaphysical' problem. Whether one suffers harm in dying depends not upon whether one would have lived a longer, better life had one's actual death not occurred, but only upon whether the powers in whose loss one's death consists were goods.

I am inclined to say that a harm's seriousness is determined primarily by the *type* of good in whose loss it consists. If this is right, then except for those cases, if any, in which the vital powers no longer constitute goods, everyone suffers pretty much the same harm in dying. The account thus allows for considerably less variation in the seriousness of the harm people suffer in dying than does McMahan's. McMahan would presumably see this as a point in favour of his account: one of the facts he thinks an adequate account of death's badness should explain is that death comes as a greater misfortune to some than to others. I grant him the point with respect to misfortune, but I think it a mistake to link the notion of harm to that of misfortune. Death may come as a greater misfortune to those whose prospects are brighter, but I think that other things being equal, a person whose future promises much happiness and a person whose future promises much less suffer equally serious harms in dying. One advantage of this position is that it goes a long way towards vindicating what McMahan calls *the equal wrongness thesis*: the thesis (roughly) that the wrongness of killing a person does not vary with the amount of well-being he loses (p. 235). McMahan too wishes to accommodate the equal wrongness thesis, but he succeeds in doing so only by divorcing the wrongness of killing persons from the seriousness of the harm they suffer; and although I agree that the seriousness of the harm suffered is not the whole explanation why killing persons is wrong, I think McMahan goes too far in denying that it is even part of the explanation.

Certain moral phenomena, however, might initially suggest that McMahan is right to take the seriousness of the harm suffered in dying to vary with the amount of well-being thereby lost. Consider, for example, the duty to give aid. Other things being equal, if we can prevent either of two people from suffering harm, but not both, we should come to the aid of the one in danger of suffering the more serious harm. (Or so it is often said.) Now suppose that two people are in danger of dying, and that one would live for ten more years if saved, the other for twenty. Shouldn't we save the latter person? And doesn't this suggest that he is in danger of suffering more harm? But an alternative analysis is available: we should save the latter person not because he'd suffer more harm in dying, but because he'd receive a greater benefit in being saved. He'd receive a greater benefit because the day of his permanently losing his vital powers would be longer delayed. Next, consider the duty not to harm people. If we can avoid harming either of two people, but not both, then

other things being equal, we should avoid harming the person who would thereby suffer the greater harm. (Or so it is often said.) Now suppose that the harm at issue is death, and that, as before, one of the potential victims stands to lose ten extra years of life, the other twenty. Shouldn't we avoid killing the person whom death would deprive of more extra years? But even if this is right, it does not follow that this person would suffer more harm in dying. After all, if, as I suppose, the two would suffer equal harm in dying, we'd still have to employ *some* criterion in deciding whom to spare, and the relative magnitudes of their prospective misfortunes would seem to provide as reasonable a criterion as any. In any case, other reasonable criteria would also be available. There would be much to be said for flipping a coin, for example.

Finally, consider the case of someone trapped in a burning car who begs to be killed painlessly before the flames reach him. If we do as he asks, surely we bestow upon him a benefit; and it might be thought that in order to get this result we must accept something like the standard account. But this is not so. In fact, I think that the account I've suggested provides a preferable analysis. According to the standard account, death from a merciful bullet comes to this person as a benefit, not a harm, because his life would have gone worse, not better, had he burned to death instead. I think, however, that killing him painlessly should be seen not as a case of pure benefiting, but as a case of *justified* harming. The person suffers serious harm in dying painlessly, but we are justified in inflicting this harm upon him because it is the only means by which we can bestow upon him the even greater benefit of preventing him from dying a horribly painful death. Where the standard account sees only benefit, the account I am suggesting sees both benefit and harm.

The way of thinking about the harm of death that I've sketched obviously requires further elaboration and defence. Perhaps it will turn out to have serious drawbacks of its own; perhaps, on balance, McMahan's modified version of the standard account will prove more acceptable. But the standard account's drawbacks should not be minimised. Indeed, no one has confronted them more seriously, insightfully or honestly than McMahan. Those of us unsatisfied with his account would do well to take his work as our model when investigating alternatives.