

Some Puzzles About the Evil of Death

Fred Feldman

I. THE PUZZLES

Death is nothing to Epicureans. They do not fear or hate death. They do not view death as a misfortune for the deceased. They think death is no worse for the deceased than is not yet being born for the as yet unborn. They say that ordinary people, who look forward to their deaths with dismay, are in this irrational. Why do they hold these odd views?

In his central argument for these conclusions, Epicurus says:

So death, the most terrifying of ills, is nothing to us, since so long as we exist, death is not with us; but when death comes, then we do not exist. It does not then concern either the living or the dead, since for the former it is not, and the latter are no more.¹

The argument seems to turn on what has been called “The Existence Condition”—nothing bad can happen to a person at a time unless he exists at that time.² If we agree that the dead don’t exist, we seem driven to the conclusion that nothing bad can happen to us once we are dead. It is just a small step then to the conclusion that death itself is not bad for those who die.

Although some may find reassurance in this ancient bit of reasoning, most of us cannot help but view it as sophistry. Except in cases in which continued life would be unbearable, death is taken to be a misfortune for the one who dies. We cry at funerals; we

¹Epicurus, “Letter to Menoeceus,” trans. C. Bailey, *The Stoic and Epicurean Philosophers*, edited and with an introduction by Whitney J. Oates (New York, N.Y.: The Modern Library, 1940), pp. 30–31. Lucretius presents essentially the same argument. See *On the Nature of Things*, trans. H. A. J. Munro, and *The Stoic and Epicurean Philosophers*, p. 131.

²Jeff McMahan, “The Evil of Death,” *Ethics* 99 (1988), pp. 32–61, at p. 33. He calls it “The Existence Requirement.”

grieve for the deceased. Especially when a young person dies, we feel that she has suffered a great misfortune. And it apparently seems to most of us that our attitude is perfectly rational. So we have our first puzzle: how can being dead be a misfortune for a person, if she doesn't exist during the time when it takes place?

According to the most popular anti-Epicurean view, death is bad for a person primarily because it deprives him of certain goods—the goods he would have enjoyed if he had not died.³ This so-called “Deprivation Approach” thus seems to require that we make a certain comparison—a comparison between (a) how well off a person would be if he were to go on living and (b) how well off he would be if he were to die. The claim is that when death is bad for a person, it is bad for him because he will be worse off dead than he would have been if he had lived. The second puzzle arises because it appears that any such comparison is incoherent. It seems to be, after all, a comparison between (a) the benefits and harms that would come to a person if he were to live and (b) those that would come to him if he were to die. However, if he doesn't exist after his death, he cannot enjoy or suffer any benefits or harms after death. So there apparently is no second term for the comparison. Thus, the Deprivation Approach seems in a covert way to violate the Existence Condition, too.⁴

Suppose we find some coherent way to formulate the view that a person's death is a misfortune for him because it deprives him of goods. Then we face another Epicurean question: *when* is it a misfortune for him? It seems wrong to say that it is a misfortune for

³I am by no means the first to defend this sort of answer. Similar views are defended (or at least discussed with some enthusiasm) by a number of philosophers. See, for example, Jeff McMahan, “The Evil of Death”; Thomas Nagel, “Death,” *Noûs* 4 (1970), pp. 73–80, revised and reprinted in *Moral Problems*, ed. James Rachels (New York, N.Y.: Harper and Row, 1975), pp. 401–409; Roy Perrett, *Death and Immortality* (Dordrecht, The Netherlands: Martinus Nijhof Publishers, 1987); L. S. Sumner, “A Matter of Life and Death,” *Noûs* 10 (1976), pp. 145–171; Douglas Walton, *On Defining Death* (Montreal, Quebec: McGill-Queen's University Press, 1979); and Bernard Williams, “The Makropulos Case: Reflections on the Tedium of Immortality,” in B. Williams, *Problems of the Self* (New York, N.Y.: Cambridge University Press, 1973).

⁴For a vigorous defense of the claim that the standard view involves an illegitimate comparison, see Harry Silverstein, “The Evil of Death,” *The Journal of Philosophy* 77 (1980), pp. 401–424.

him while he is still alive—for at such times he is not yet dead and death has not yet deprived him of anything. It seems equally wrong to say that it is a misfortune for him after he is dead—for at such times he does not exist. How can he suffer misfortunes then?

Another problem confronts the anti-Epicurean. If we can find a coherent way to say that early death is bad for us because it deprives us of certain goods, then we probably will have found a coherent way to say that late birth also deprives us of certain goods—the goods we would have enjoyed if only we had been born earlier. Yet virtually nobody laments his late birth, or thinks it a misfortune that he wasn't born years or decades earlier. Lucretius presented a forceful statement of this puzzle. He said:

Think too how the bygone antiquity of everlasting time before our birth was nothing to us. Nature therefore holds this up to us as a mirror of the time yet to come after our death. Is there aught in this that looks appalling, aught that wears an aspect of gloom? Is it not more untroubled than any sleep?⁵

So another puzzle that must be confronted is this: if early death is bad for us because it deprives us of the goods we would have enjoyed if we had died later, then why isn't late birth just as bad for us? After all, it seems to deprive us of the goods we would have enjoyed if we had been born earlier.

There are other puzzles about the evil of death. Some of these will be addressed as we go along. But these are the main questions I mean to discuss here.

II. METAPHYSICAL AND AXIOLOGICAL ASSUMPTIONS

Before I propose my answers to these questions, I should mention some of my metaphysical and axiological assumptions. First among these, perhaps, is the assumption that there are possible worlds. I am inclined to think that a possible world is a huge proposition fully describing some total way the world might have been, including all facts about the past, present and future. Nothing I say here depends on this particular view about possible worlds. So long as it countenances an appropriate number of ap-

⁵*On the Nature of Things*, p. 134.

propriately detailed possible worlds, any other coherent view will do as well.

I write as if a given individual may exist at several different possible worlds. This may seem controversial, but I think it is really not. Suppose Myron is an actual person. Suppose he actually smokes. I may ask you to consider some possible world in which Myron does not smoke. This may seem to commit me to the view that there are other worlds relevantly like our (concrete) world, and that in addition to being here in our (concrete) world, the actual concrete Myron (or perhaps a counterpart) is also located at these other places. That, it seems to me, would be strange.

In fact, however, I hold no such view. When I ask you to consider some world in which Myron does not smoke, I am just asking you to consider a huge proposition that fully describes some total way the world might have been, and which entails the proposition that Myron exists but does not smoke. Since it is more convenient to do so, I write in a “realistic” way about other possible worlds—as if they were giant, concrete planets far from Earth, but populated by many earthlings.

I assume that it makes sense to speak of the degree of similarity between possible worlds. Indeed, it seems to me that there are many similarity relations among possible worlds. Later I will have more to say about the details of the similarity relations that are most important for present purposes. However, if we have some particular similarity relation in mind, then it will make sense to speak of some world as being “most similar” in that way to a given world. Sometimes instead of speaking of similarity I speak of “nearness.” It is just another way of expressing the same idea.⁶

Now let us briefly turn to axiology. Possible worlds can be evaluated in various ways. One sort of evaluation is “objective” and “non-relational.” Suppose that the very simplest form of hedonism is true. According to this view, pleasure is intrinsically good and pain is intrinsically bad. Nothing else has any (basic) intrinsic value. Let’s suppose that there is a way to measure the amount of pleasure contained in an episode of pleasure; let’s suppose simi-

⁶The *locus classicus* of many of these ideas is David Lewis, *Counterfactuals* (Cambridge, Mass.: Harvard University Press, 1973).

larly that there is a way to measure the amount of pain contained in an episode of pain. Suppose further that the pleasure-measure and the pain-measure are commensurate, so that it makes sense to subtract amounts of pain from amounts of pleasure.⁷ We can then say that the intrinsic value of a possible world is determined as follows: consider how much pleasure is experienced throughout the history of that world; consider how much pain is experienced throughout the history of that world; subtract the latter value from the former; the result is the hedonic value of the world. The simplest form of hedonism says that the intrinsic value of a world is equal to the hedonic value of that world.⁸

Another way to evaluate worlds is equally “objective,” but is “person-relative.” That is, instead of asking how good a world is, we ask how good it is *for a certain person*. When I speak of how good a world is for a certain person, I mean to indicate the portion of that world’s goods and evils that the individual in question enjoys and suffers at that world. Suppose again that the simplest form of hedonism is true. Then the value of a world, w , for a person, s , is determined in this way: consider how much pleasure s enjoys throughout his lifetime at w ; consider how much pain s suffers throughout his lifetime at w ; subtract the value of the latter from the value of the former. The result is the value of w for s , or $V(s,w)$.

I assume that these values can be expressed with numbers in such a way that higher numbers indicate greater value for the person; zero indicates neutrality for the person; negative numbers indicate badness for the person. Since $V(s,w)$ is a measure of how well s fares at w , I sometimes refer to this as s ’s “welfare level” at w .⁹

⁷I attempted to present a clear formulation of this view about axiology in my *Doing the Best We Can: An Essay in Informal Deontic Logic* (Dordrecht, The Netherlands: D. Reidel, 1986). See especially Section 2.2.

⁸I doubt that many moral philosophers would endorse anything like this simplest form of hedonism. Indeed, I wouldn’t endorse it either. My point here is primarily to indicate something about the structure of an axiological view—it should yield an ordering of worlds in terms of value. In an attempt to make this conception most obvious, I have assumed that there is a value function taking us from worlds to numbers. This structural approach is consistent with a wide variety of substantive axiological theories.

⁹It should be obvious that in interesting cases, no one could possibly

There is a question concerning a person's welfare level at worlds at which he does not exist. The proposed account leaves this value undetermined. Although it plays no role in my argument, I stipulate that if s fails to exist at w , then $V(s,w) = 0$. This thesis is suggested by the proposed account of relativized value, since if a certain person does not exist at a world then he enjoys no pleasure there and suffers no pain there.

In fact, I do not think that a person's real welfare level is determined in the simple-minded hedonistic way I have sketched. I am inclined to think that several other factors may contribute to determining how good a world is for a person. Among other things, I suspect that the amounts of knowledge and freedom that a person enjoys, as well as the extent to which he is forced to suffer injustice are also important. However, I prefer to proceed here on the pretense that hedonism is true. I have several reasons.

First and foremost, there is the historical reason. I am engaged in a debate with Epicurus about the evil of death. Epicurus was a hedonist. Some commentators have suggested that in order to answer Epicurus, we must reject his axiology—that his view about the evil of death is inextricably tied to his hedonism. I think this is a mistake. I want to show that, even if we accept the Epicurean axiology, we can still reject the Epicurean conclusion about the evil of death.

A second reason for assuming hedonism is strategic. The central intrinsic value-bearing properties associated with hedonism are ones that a person can have at a time only if he is alive and conscious then. I want to show how death can be an evil for the deceased even if this hedonistic axiology is assumed. Thus, I take myself to be trying to show that death may be an evil for a person even according to an axiology maximally hostile to this notion. If I succeed, it will be pretty easy to see how to extend the solution in the direction of more plausible axiologies.

It should be clear, then, that certain sorts of solution are ruled out by my axiological assumptions. I will not be able to say (as Thomas Nagel and others have suggested)¹⁰ that death is bad in

calculate the value of a world for a person. On the other hand, we could have reason to believe that worlds of a certain specified sort would be uniformly *worse* for someone than worlds of some other specified sort.

¹⁰Nagel discusses this idea in his now classic paper, "Death," cited above

something like the way in which being the subject of nasty rumors is bad. Clearly enough, one can be the subject of nasty rumors even after one has died. If we think this is bad for a person, then we will want to say that one's welfare level at a world can be adversely affected by things that happen after one ceases to exist at that world. Another sort of example involves the failure of one's life projects. One's life projects may come unraveled after one has died. If we think this is bad for a person, then we can cite another way in which one's welfare level at a world may be reduced by things that occur after one's death.

These claims about welfare levels are controversial, and strike me as being implausible. I would rather stick to a much more hard-nosed axiology—an axiology according to which one's welfare level at a world is determined entirely by things that happen during one's life there. Thus (for purposes of illustration) I have adopted a form of simple hedonism. According to this view, if a person never learns of nasty rumors, and never suffers from them, then they don't affect his welfare level. If a person never learns that his life project has come to naught, and never suffers from this frustration, then it doesn't affect his welfare level. Only pains and pleasures can affect a person's welfare level at a world—and these he must experience during his life.

A final advantage of the hedonistic axiology is its simplicity. If we assume that the fundamental bearers of intrinsic value are experiences of pleasure and experiences of pain, and we assume that these are in principle subject to unproblematic quantification, then the determination of a person's welfare level at a possible world becomes quite straightforwardly a matter of simple arithmetic. To find s 's welfare level at w , just subtract the amount of pain s suffers at w from the amount of pleasure s enjoys at w . Although the axiology is admittedly quite crude, its simplicity makes it especially useful for illustrative purposes.

I assume that any statement to the effect that something is good (or bad) for a person can be paraphrased by a statement to the effect that some *state of affairs* is good (or bad) for the person. Fur-

in fn. 3. A similar approach to the evil of death is suggested by George Pitcher in "The Misfortunes of the Dead," *American Philosophical Quarterly* 21 (1984), pp. 183–188.

thermore, I assume here that a state of affairs (such as the state of affairs of *Myron smoking*) is just a proposition (in this case, the proposition *that Myron smokes*). Thus, for present purposes, it makes no difference whether we say that a certain state of affairs obtains, or whether we say that a certain proposition is true.

In any case, instead of saying that smoking (apparently an activity) would be bad for Myron, we can say instead that *that Myron smokes* (a state of affairs) is bad for Myron. Instead of saying that a bowl of hot soup (apparently a physical object) would be good for me, we can say that what would be good for me is *that I have a bowl of hot soup*, and thus again represent the thing that is good for me as a state of affairs. I prefer to write in this way, since it induces a sort of conceptual tidiness and uniformity.

I am also going to assume that when a person dies, he goes out of existence. In fact, I think this assumption is extremely implausible. No one would dream of saying that when a tree dies, it goes out of existence. Why should we treat people otherwise? My own view is that a person is just a living human body. In typical cases, when the body dies, it continues to exist as a corpse. So the thing that formerly was a person still exists, although it is no longer alive (and perhaps no longer a person). Of course, I recognize that some people go out of existence at the moment of death—for example, those located at Ground Zero at the moment of a nuclear blast. For present purposes, I will assume that everyone does. Once again, I do this in part for historical reasons—Epicurus seems to have accepted this view about death and nonexistence—and in part for strategic reasons. I want to show how death can be bad for the deceased even on the assumptions (a) that things that directly affect a person's welfare level can happen to that person only at times when he exists, and (b) that death marks the end of existence for the deceased.¹¹

III. THINGS THAT ARE BAD FOR PEOPLE

The central question here is how a person's death can be bad for him. The claim that someone's death is bad for him is an instance

¹¹Some commentators suppose that we stop existing when we die, but we don't stop "being." They also suppose that appealing to the existence/being distinction helps solve the problem about the evil of death. For an

of a more general sort of claim—the claim that some state of affairs is bad for some person. It would be surprising if it were to turn out that we need two independent accounts of what's meant by statements to the effect that something is bad for someone: one account of the meaning of such a statement when the relevant object is the person's death, and another account of the meaning of such a statement when the relevant object is something other than the person's death. Surely the statement about death ought to be nothing more than an interesting instance of the general sort of statement. So let's consider the more general question first, and then focus more narrowly on the specific case concerning death. What do we mean when we say that something would be bad for someone?

It seems to me that when we say that something would be bad for someone, we might mean either of two main things. One possibility is that we mean that the thing would be *intrinsically* bad for the person. So if someone says that a state of affairs, *p*, is intrinsically bad for a person, *s*, he presumably means that *p* is intrinsically bad, and *s* is the subject or "recipient" of *p*. Given our assumed hedonistic axiology, the only things that could be intrinsically bad for someone would be his own pains. Thus, *Dolores suffering pain of intensity 10 from t1 to t3* would be intrinsically bad for Dolores.

On the other hand, when we say that something would be bad for someone, we might mean that it would be "all things considered bad" for him. At least in some instances, this seems to mean that he would be all things considered worse off if it were to occur than he would be if it were not to occur. In this case, the thing itself might be intrinsically neutral. The relevant consideration would be the extent to which it would lead to or prevent or otherwise be connected with things that are intrinsically bad for the person. Consider an example. Suppose we are interested in the question whether moving to Bolivia would be bad for Dolores. Intuitively, this question seems to be equivalent to the question whether Dolores would be worse off if she were to move to Bolivia than she would be if she were to refrain from moving to Bolivia.

example of this approach, see Palle Yourgrau's "The Dead," *The Journal of Philosophy* 84 (1987), pp. 84–101. In this paper, I have made no such distinction.

Letting “b” indicate the state of affairs *Dolores moves to Bolivia*, we can say this: b would be all things considered bad for Dolores if and only if she would be worse off if b obtained than she would be if b didn’t obtain. Now, if we employ the standard account of the meaning of subjunctive conditionals, together with the assumptions about values of worlds for individuals, we can rewrite this as follows: b would be all things considered bad for Dolores if and only if the value for Dolores of the nearest possible b-world is less than the value for her of the nearest possible \sim b-world.¹²

Correspondingly, to say that a state of affairs would be all things considered good for a person is to say that she would be better off if it were to obtain than she would be if it were to fail to obtain. More exactly, it is to say that her welfare level at the nearest possible world where it obtains is higher than her welfare level at the nearest possible world where it does not obtain.

If we make use of the abbreviations introduced above, we can restate these claims as follows:

D1: p would be good for s if and only if $(\exists w) (\exists w') (w \text{ is the nearest } p\text{-world} \ \& \ w' \text{ is the nearest } \sim p\text{-world} \ \& \ V(s,w) > V(s,w'))$

¹²I am suppressing consideration of certain complexities. One that should be addressed concerns cases in which there is no unique nearest world in which a certain state of affairs occurs—several worlds are tied for this distinction. What shall we say then?

Suppose that at the real world Dolores does not move to Bolivia. Then the real world is the nearest world in which she does not move to Bolivia. Suppose that among worlds in which she does move to Bolivia, there are two that are equally near and most near the real world. Then I want to say this: if each of these worlds is worse for Dolores than the real world, then moving to Bolivia would be bad for her; if each is better for her than the real world, then moving to Bolivia would be good for her; if one is better and the other is worse, then it’s not the case that moving to Bolivia would be good for her, and it’s not the case that moving to Bolivia would be bad for her; moving to Bolivia might be good for her and might be bad for her.

If all the nearest b-worlds have the same value for Dolores, then we can use this value when we compute the value of b for Dolores. On the other hand, if the nearest b-worlds differ in value for Dolores, then the computations become more problematic. One possibility would be to make use of the average value for Dolores of these nearest b-worlds. Another possibility would be to say that the value of b for her might be the result of subtracting the value for Dolores of the real world from the value for her

D2: p would be bad for s if and only if $(\exists w)(\exists w')$ (w is the nearest p -world & w' is the nearest $\sim p$ -world & $V(s,w) < V(s,w')$)

If we make use of our assumption that worlds have numerical values for individuals, then we can say precisely *how bad* or *how good* something would be for someone. Suppose that if Dolores were to move to Bolivia the rest of her life would be a nightmare. Considering all the pleasures and pains she would ever experience (including the ones she has already experienced), her life would be worth +100 points. Thus, the value for Dolores of the nearest world in which she moves to Bolivia is +100. Suppose on the other hand that the value for her of the nearest world in which she does not move to Bolivia is +1000. Then she would be 900 units worse off if she were to move to Bolivia. That tells us precisely how bad it would be for her to move to Bolivia. The value for her of moving to Bolivia is -900. So the general principle says that to find the value for a person of a state of affairs, subtract the value for him of the nearest world where it does not obtain from the value for him of the nearest world where it does obtain.

Precisely the same thing happens in the case of a state of affairs that would be good for a person. Suppose it would be good for Dolores to move to Boston. To find out how good it would be for her, consider the value for her of the nearest world in which she does move to Boston. Suppose it is +1100. Consider the value for her of the nearest world in which she does not move to Boston. Suppose it is +1000. Subtract the value for her of the latter from the value for her of the former. The result (+100) is the value for Dolores of moving to Boston.

In its most general form, then, the principle may be formulated as a principle about the overall value (good, bad, or neutral) of states of affairs for persons. The overall value of a state of affairs for a person is the result of subtracting the value for him of the

of one of them, and it might be the result of subtracting the value for her of the real world from the value for her of another. In such a case, we would have to say that there is no number, n , such that the value of b for Dolores = n .

In what follows, I shall write as if there is always a unique nearest world. My main points are not affected by this simplifying assumption.

nearest world where it does not occur from the value for him of the nearest world where it does occur. In other words:

D3: The value for s of $p = n$ if and only if $(\exists w)(\exists w')(w \text{ is the nearest } p\text{-world} \ \& \ w' \text{ is the nearest } \sim p\text{-world} \ \& \ V(s,w) \text{ minus } V(s,w') = n)$.

IV. THE EVIL OF DEATH

The application of these ideas to the case of one's own death is straightforward. Suppose we are wondering whether it would be bad for a certain person, s , to die at a certain time, t . Then we must ask about the value for s of the possible world that would exist if s were to die at t ; and we must compare that value to the value for s of the possible world that would exist if s were not to die at t . If the death-world is worse for s than the non-death-world, then s 's death at t would be bad for s ; otherwise, not.

Let's consider a typical example to see how this works. Suppose I am thinking of taking an airplane trip to Europe. Suppose I'm worried about accidents, hijackings, sabotage, etc. I think I might die en route. I think this would be bad for me. D3 directs us to consider the nearest possible world in which I do die en route to Europe on this trip, and to consider my welfare level at that world. I see no reason to suppose that interesting parts of my past are any different at that world from what they are at the actual world. So I assume that all my past pleasures and pains would be unaffected. The main difference (from my perspective) is that in that world I suffer some terminal pain and then a premature death, and never live to enjoy my retirement. Let's suppose that that world is worth +500 to me—+500 is the result of subtracting the pain I there suffer from the pleasure I there enjoy. Next D3 directs us to consider the nearest world in which I do not die en route to Europe on this trip. The relevant feature of this world is that I do not die a painful and premature death in an airplane accident. Suppose I there do live to enjoy many happy years of retirement. Let's suppose my welfare level at that world is +1100. D3 implies that my death on this trip would have a value of -600 for me. It would be a terrible misfortune.

Two points deserve mention here. One is the fact that D3 is a proposal concerning how *good* or *bad* a state of affairs is for a

person, and not a proposal concerning the extent to which a state of affairs *benefits* or *harms* a person. I am inclined to suspect that the concepts of benefit and harm are in certain important ways different from the concepts of being good for and being bad for a person. One such respect might be this: it might be that it is impossible for a person to be harmed or benefitted by things that happen at times when he no longer exists. It is nevertheless still possible that something bad or something good for a person might occur at a time when the person no longer exists. D3 is not intended to have any direct implications concerning harm and benefit. It is intended to be restricted to the concepts of being good for a person and being bad for a person.

The second point is that nothing I have said here implies that death is always bad for the one who dies. Suppose a person is suffering from a painful terminal disease. Suppose he is considering suicide, and is inclined to think that death might be a blessing. He might be right. If his welfare level at the nearest world where he thus commits suicide is higher than his welfare level at the nearest world where he doesn't commit suicide, then committing suicide would be good for this person.¹³ My point in formulating D3 is simply to show how it is possible for a person's death to be bad for him, not that everyone's death must be so.

Perhaps we can now see where Epicurus went wrong in his argument for the conclusion that one's death cannot be bad for him. Perhaps Epicurus was thinking that the only states of affairs that are bad for a person are the ones that are *intrinsically* bad for him. Since (given our axiological assumptions, which are intended to be relevantly like his) death is not intrinsically bad for anyone, it would follow that death is never bad for the one who dies. But even the most fervent hedonist should acknowledge a distinction between things that are intrinsically bad for a person (which he will

¹³I think these remarks provide the basis for a reply to one sort of argument concerning the alleged irrationality of suicide. Some have said that suicide is always irrational since it is impossible to calculate the value of death for the deceased. See, for example, John Donnelly's "Suicide and Rationality," in *Language, Metaphysics and Death*, ed. John Donnelly (New York, N.Y.: Fordham University Press, 1978); and Philip Devine's *The Ethics of Homicide* (Ithaca, N.Y.: Cornell University Press, 1978), esp. p. 25. If what I have said here is right, the calculations are in principle possible, and some suicides are perfectly rational.

take to be pains) and things that are bad for the person in other ways. D3 is designed to calculate an important sort of *non-intrinsic* value. It tells us the degree of *overall* badness for a person of a state of affairs. Even though my death on my imagined European trip would not be intrinsically bad for me, D3 tells us that it would be overall bad for me.

Another possibility is that Epicurus was thinking that if a state of affairs would be bad for a person, then it must at least *cause* something intrinsically bad for him. Since (given our axiological and metaphysical assumptions) nothing intrinsically bad can happen to me after my death, my death cannot cause anything intrinsically bad for me. Thus, Epicurus might have concluded that my death cannot even be extrinsically bad for me. However, D3 does not calculate extrinsic value by focusing exclusively on intrinsic goods and evils that would befall the person *as a result* of the state of affairs. Rather, it calculates the value of a state of affairs for a person by considering what would happen (whether as consequence or not) if the state of affairs were to occur, as compared to what would happen (whether as consequence or not) if it were to fail to occur. Thus, according to D3, my death would be bad for me not because it would cause me to suffer pain, and not because it would itself be intrinsically bad for me. Rather, it would be bad for me because it would deprive me of 600 units of pleasure that I would have had if it had not happened when it did. More precisely, it would be bad for me because my welfare level at the nearest world where it occurs is 600 points lower than my welfare level at the nearest world where it does not occur.

V. SOME PROPOSED ANSWERS

At the outset, I mentioned some questions about the evil of death. These were prompted by the Epicurean challenge. I will now attempt to answer those questions.

The first question was the question how, given that he doesn't exist after he dies, being dead can be a misfortune for a person. The simple answer is this: a state of affairs can be bad for a person whether it occurs before he exists, while he exists, or after he exists. The only requirement is that his welfare level at the nearest world where it occurs is lower than his welfare level at the nearest world where it does not occur. It may be interesting to consider an

example in which something bad for a person occurs *before* the person exists. Suppose my father lost his job shortly before I was conceived. Suppose that as a result of the loss of his job, my parents had to move to another town, and that I was therefore raised in a bad neighborhood and had to attend worse schools. I would have been happier if he had not lost his job when he did. In this case, the fact that my father lost his job was bad for me—even though I didn't exist when it occurred. It was bad for me because the value for me of the nearest world where he didn't lose his job is greater than the value for me of the actual world (which, on the assumption, is the nearest world where he did lose his job). The same may be true of cases involving things that will happen after I cease to exist (although, of course, such cases will illustrate *deprivation* of happiness, rather than *causation* of unhappiness).

It should be clear, then, that the plausibility of the Existence Condition derives from a confusion. Given our hedonistic axiology, it would be correct to say that nothing *intrinsically* bad can happen to a person at a time unless he exists at that time. You cannot suffer pains at a time unless you then exist. However, even on the same axiology, the *overall* value version of the thesis is not true. That is, it would not be correct to say that nothing *overall* bad for a person can happen at a time unless he exists at that time. Perhaps some Epicureans have been induced to accept the Existence Condition because they fail to notice this distinction.¹⁴

The second puzzle concerns an allegedly illegitimate comparison. It may seem that I am maintaining that when a person's death is bad for him, it is bad for him because he's worse off being dead than he would have been if he had stayed alive. Yet this suggests

¹⁴In "How to be Dead and Not Care: A Defense of Epicurus," *American Philosophical Quarterly* 23 (1986), Stephen Rosenbaum proposes an interpretation of the Epicurean argument. He suggests that one crucial premise is "A state of affairs is bad for a person P only if P can experience it at some time" (p. 218). I would say that the premise is ambiguous. If taken to mean that a state of affairs is *intrinsically* bad for a person only if he can experience it, then (assuming hedonism or any other "experience-based" axiology) the premise may be true—but it is not relevant to the claim that death is bad for the one who dies, since it is most reasonable to take this as the claim that death is *extrinsically* bad for the one who dies. If the claim is understood in this more plausible way as the claim that a state of affairs can be *extrinsically* bad for a person only if he can experience it, then, as I have attempted to show, the premise is false.

that there is some degree of “bad-offness” that he endures while dead. However, since he doesn’t exist while he is dead, he can have no degrees of “bad-offness” then. The question, then, is this: doesn’t my answer presuppose an illegitimate comparison?

My answer presupposes no such comparison. I am not proposing that we compare a person’s welfare level during life to his welfare level during death. I have assumed that one’s welfare level at a world is determined entirely by pleasures and pains that one experiences during one’s life at that world. Thus, the comparison is a comparison between one’s welfare level (calculated by appeal to what happens to one during his life) at one possible world with his welfare level (also calculated by appeal to what happens to him during his life) at another possible world. I have provisionally agreed that nothing intrinsically good or bad can happen to a person at times when he does not exist.

In effect, then, my proposal presupposes what Silverstein calls a “life-life comparison.”¹⁵ To see how this works, consider again the example concerning my imagined death en route to Europe. My proposal requires us to compare the values for me of two lives—the life I would lead if I were to die on the plane trip and the life I would lead if I were not to die on the plane trip. Since (according to our assumptions) the shorter life is less good for me, my death on that trip would be correspondingly bad for me.

The third puzzle was a puzzle about dates. I have claimed that a person’s death may be bad for her because it deprives her of the pleasures she would have enjoyed if she had lived. One may be puzzled about just *when* this misfortune occurs. The problem is that we may not want to say that her death is bad for her during her life, for she isn’t yet dead. Equally, we may not want to say that it is bad for her after her death, for she doesn’t exist then.

In order to understand my answer to this question, we must look more closely into the question. Suppose a certain girl died in her youth. We are not concerned here about any puzzle about the date of her death. We may suppose we know that. Thus, in one sense, we know precisely when the misfortune occurred. Nor are we concerned about the dates of any pains she suffered as a result of that death. We assume that there are none. The present question is,

¹⁵Silverstein, “The Evil of Death,” p. 405.

rather, a question about when her death is a misfortune for her. If Lindsay is the girl, and d is the state of affairs of *Lindsay dying on December 7, 1987*, then the question is this: “Precisely when is d bad for Lindsay?” I have proposed an account of the evil of death. According to that account, when we say that d is bad for Lindsay, we mean that the value for her of the nearest world where d occurs is lower than the value for her of the nearest in which d does not occur. So our question comes to this: “Precisely *when* is it the case that the value for Lindsay of the nearest world in which d occurs is lower than the value for her of the nearest world in which d does not occur?”

It seems clear to me that the answer to this question must be “eternally.” For when we say that her death is bad for her, we are really expressing a complex fact about the relative values of two possible worlds. If these worlds stand in a certain value relation, then (given that they stand in this relation at any time) they stand in that relation not only when Lindsay exists, but at times when she doesn’t. If there were a God, and it had been thinking about which world to create, it would have seen prior to creation that d would be bad for Lindsay. In other words, it would have seen that the value for Lindsay of the relevant d -world is significantly lower than the value for Lindsay of the relevant $\sim d$ -world. And it would have seen this even though Lindsay did not yet exist at that pre-creation moment.

A final puzzle concerns the fact that we feel that early death is a greater misfortune for the prematurely deceased than is “late birth” for the late born, even though each may deprive us of as much happiness as the other.

Suppose Claudette was born in 1950 and will die somewhat prematurely in 2000 as a result of an accident. We may want to say that her premature death will be a misfortune for her. Consider the nearest possible world (call it “ w_3 ”) in which she does not die prematurely. Suppose that at w_3 she lives happily until 2035. Since she has 35 extra years of happiness in w_3 , her welfare level there is higher than her welfare level in the actual world. D_3 yields the result that her premature death is bad for her. But now consider the claim that Claudette suffered an equal misfortune in not having been born in 1915. This fact seems to deprive her of 35 happy years too—the years from 1915 to 1950 when she was in fact born. Yet we feel uncomfortable with the idea that her late

birth is as great a misfortune for Claudette as her premature death. Why is this?

Consider the state of affairs of *Claudette being born in 1915*. Call it “b.” In the actual world b is false. Consider the nearest world where b is true.¹⁶ (In other words, consider what would have happened if Claudette had been born 35 years earlier.) Call this world “w4.” I see no reason to suppose that Claudette lives any longer in w4 than she does here in the actual world. Any such change in lifespan strikes me as being superfluous. I am inclined to suppose that Claudette’s welfare level in w4 is slightly lower than her welfare level in the actual world—after all, in w4 she probably endures hard times during the Great Depression, and maybe even catches measles, whooping cough and other diseases that were rampant in those days. If she has just fifty years to live, she’s better off living them in the second half of the twentieth century, rather than thirty-five years earlier. Thus, given my intuitive sense of how to calculate what would have happened if Claudette had been born earlier, it follows that early death is worse for Claudette than late birth. Her late birth deprived her of very little value; her early death would deprive her of a lot.

The proposed reply to Lucretius’ challenge is thus based on an asymmetry between past and future. When I am asked to consider what would happen if Claudette were to die later, I hold her birth-date constant. It has already occurred, and I tend to think that unnecessary differences in past history are big differences between

¹⁶In “Death” (p. 67), Thomas Nagel claims that late birth does not deprive anyone of anything, since no one could have been born much earlier than she was in fact born. This provides the basis for a quick answer to Lucretius. Derek Parfit makes a similar claim in *Reasons and Persons* (Oxford, England: Oxford University Press, 1984), p. 351. The argument might be based on the essentiality of origins. However, with the development of techniques for the cryopreservation of sperm and eggs, the view seems false. Even if we grant the controversial claim that each person has her origins essentially, we have to acknowledge that once the relevant sperm and egg have been frozen, it is in principle possible for her to be conceived at any time in the next thousand years or so. I grant, of course, that the issue of the essentiality of origins deserves independent discussion. I simply assume that it makes sense to speak of what would have happened if Claudette had been born earlier. This makes it possible to look more deeply into the puzzle suggested by Lucretius.

worlds. Thus, it is more natural for me to suppose that if she were to die later, it would be because she lives longer. On the other hand, when I am asked to consider what would have happened if she had been born earlier, I do not hold her deathdate constant. Instead, I hold her lifespan constant, and adjust the deathdate so as to accommodate itself to the earlier birthdate.

Someone might claim that I have made an unfair comparison. They might want to insist on holding lifespans constant. They might say that Claudette would be better off living longer if the extra time were tacked on to the end of her life. They might say that Claudette would not be any better off if the extra time were tacked on to the beginning of her life. (That is, if she were born in 1915 instead of 1950 but lived until 2000 anyway.) The question is vexing, since it is hard to discern Claudette's welfare levels in the appropriate worlds. My own inclination is to say that if she lives 85 happy years in each world, then her welfare level at the one is equal to her welfare level at the other. In this case, I can't see why anyone would think it would be better for her to have the 35 years tacked on at the end of her life rather than at the beginning. When the comparison is fair, D3 generates what seem to me to be the correct results. And the results are that the deprivation of 35 happy years of life is a bad thing, whether these years would have occurred before the date at which Claudette was in fact born, or after the date on which she in fact died.

There are, after all, two ways in which we can rectify the apparently irrational emotional asymmetry. On the one hand, we can follow Lucretius and cease viewing early death as a bad thing for Claudette. On the other hand, we can at least try to start viewing late birth as a bad thing. My suggestion is that in the present case, the latter course would be preferable.

I think it must be granted that our emotional reactions toward pleasures lost by early death are quite different from our emotional reactions toward similar pleasures lost by late birth. If my proposal is right, this emotional asymmetry is irrational. To see this, consider a variant of the case involving Claudette. Suppose (to make the case very "clean") that Claudette never experienced any pleasures or pains, but that if she had died later, she would have enjoyed one especially great pleasure ("the Late Pleasure") in her old age. Suppose similarly that if she had been born earlier, she

would instead have enjoyed an equally great pleasure (“the Early Pleasure”). In either case, her life would have contained exactly one pleasure.

Given natural assumptions, my proposal yields the result that Claudette’s late birth was just as bad for her as was her early death. Yet I suppose that at times near the end of her life, Claudette and her friends would have been more upset about her impending early death than they would have been about her late birth. Perhaps this emotional asymmetry is to be explained by the fact that we tend to think that the past is fixed, whereas the future is still open. Thus, we may feel that there’s no point in lamenting the fact that Claudette missed the Early Pleasure. On the other hand, we may feel that there was a “real chance” that she might have enjoyed the Late Pleasure. Her loss of that seems a greater misfortune.

Another possibility is that we have what Derek Parfit has called “a bias toward the future.” Once they are past, we become indifferent toward our pleasures and pains; while they are still in the future, we care deeply about them.¹⁷ If hedonism is true, this sort of asymmetry is wholly irrational. Nevertheless, it might be a deep-seated feature of human psychology.

I want to emphasize the fact that my central proposal here concerns a value-theoretic question, not a question in psychology. I mean to be discussing the question about the relative evil of early death and late birth. I have not attempted to answer the psychological question about the differences in the ways in which we react to early death and late birth. If my proposal is right, then (to a large extent) our emotional reactions may be irrational.

VI. AN OBJECTION AND A REPLY

In “Death and the Value of Life,” Jeff McMahan considers and rejects an account of the evil of death very much like the one I mean to defend.¹⁸ He cites a number of difficulties for any such

¹⁷Derek Parfit, *Reasons and Persons*. An interesting proposal based on some Parfitian ideas can be found in “Why is Death Bad?” by Anthony Brueckner and John Martin Fischer, *Philosophical Studies* 50 (1986), pp. 213–221.

¹⁸McMahan (in “The Evil of Death,” cited above in n. 3) discusses what he calls “the revised possible goods account.” This is relevantly like my

view. One concerns a young cavalry officer who is shot and killed in the charge of the Light Brigade. According to the story, the officer was shot by someone named “Ivan.” McMahan stipulates that if the officer had not been killed by Ivan’s bullet, he would have died just a few seconds later by a bullet fired by Boris. McMahan says that “. . . our answer to the question of what would have happened had the officer not died when and how he did will be that he would have lived for a few seconds, and then he would have been killed. This leads to the unacceptable conclusion that his actual death was hardly a misfortune at all.”¹⁹

McMahan goes on to offer various revisions of the original proposal, but these seem to me to be changes for the worse (and I explain why below). It seems to me that D3 generates appropriate results.

It is important to distinguish several different things that happen in this example. Let us call the gallant officer “Herbert,” and let us suppose the time of his death was 3:30 p.m., October 25, 1854—or “t.” Here are some states of affairs that we should distinguish:

- P1: Herbert dies at exactly t.
- P2: Herbert dies near Balaclava.
- P3: Herbert dies in the charge of the Light Brigade.
- P4: Herbert dies as a result of being shot by Ivan.
- P5: Herbert dies in his youth.

It should be clear that we have five different states of affairs here. In fact, each is logically independent of each of the others. Furthermore, it should come as no surprise if some of these are worse for Herbert than others. Given the details of the story, it turns out that P1 and P4 are not very bad for Herbert. Neither of these deprived Herbert of much happiness, since if he hadn’t been killed at t by Ivan, he would have been killed seconds later by Boris. It’s hard to see why this calls for any alteration of D3. These states of affairs seem to me not to be very bad for Herbert. The real tragedy here is not that he died exactly at t, or that he died as

proposal. He claims that it runs into the “problem of specifying the antecedent” (p. 43).

¹⁹Ibid., p. 46.

a result of being shot by Ivan; the real tragedy is that he died so young. Thus, P5 should be the focus of our attention.

We must consider the nearest possible world in which P5 does not occur. Let's call it "w5." What sort of life does Herbert live there? Perhaps in w5 Herbert is one of the few survivors of the charge; perhaps he is wounded, but recovers and goes on to live a long and happy life. Of course, I don't know precisely what happens to Herbert in w5—but it is reasonable to suppose that in w5 Herbert's welfare level is significantly higher than it is here in the actual world. After all, in w5 Herbert does not die in his youth, but is otherwise as much as possible like he is here in the actual world. In any case, according to D3, the badness of P5 for Herbert is equal to the difference in value for Herbert between w5 and the actual world. This might be a significant difference. He might have led a long and happy life if he had not died in his youth.

I mentioned earlier that I think that McMahan's view is less plausible than D3. On McMahan's proposal, we are asked to consider what happens in a world far more distant than w5. McMahan asks us to consider the nearest world *in which the whole causal sequence leading up to Herbert's death fails to occur*. As McMahan remarks, in the example cited, this may mean considering a world in which the Crimean War does not occur.²⁰ This strikes me as being implausible. To see how it could go wrong, suppose that Herbert loved excitement. If there had been no Crimean War, he would have sought excitement elsewhere. He would have taken up mountain climbing, and would have been killed in 1853. Given these assumptions, McMahan's proposal yields the surprising result that being killed in the Crimean War was *good* for Herbert. It seems to me to make much more sense to consider a nearer world—a world in which the Crimean War occurs, Herbert participates, but does not die a premature death. w5 is supposed to be such a world, and Herbert is better off in w5 than he is in the real world.

VII. CONCLUSION

I have attempted to formulate a coherent answer to the ancient challenge set by Epicurus. I have claimed that there is nothing

²⁰“ . . . we must presumably imagine that the Crimean War did not occur, in which case the threat from Boris would not have occurred either” (ibid., p. 47).

paradoxical about the idea that death may be bad for the one who dies. My answer is a version of the traditional view that death is bad (when it is bad) primarily because it deprives the deceased of goods—the goods he would have enjoyed if he had lived. I have attempted to provide my answer within a predominantly Epicurean framework. I have assumed that hedonism is true, and I have assumed that when a person dies, he goes out of existence. I have attempted to show that even if we grant these assumptions, we can still maintain that death can be evil for the deceased. I have furthermore attempted to show that if we formulate our account properly, we can provide satisfactory answers to some puzzling questions: “How can death be bad for the deceased if he doesn’t exist when it occurs?” “When is death bad for the deceased?” “Is there an illegitimate comparison between the welfare of the non-existent and the welfare of the existent?” “Why is death worse than prenatal nonexistence?” Along the way, I have also discussed the merits of some other proposed solutions to the puzzles.²¹

University of Massachusetts at Amherst

²¹Many friends provided much-needed criticism and support, for which I am thankful. I am especially grateful to Gary Matthews for his encouragement on this and related projects. Earl Conee, John Fischer, Ed Gettier, Ned Markosian, Neil Schaefer, Harry Silverstein, and the editors of *The Philosophical Review* made helpful suggestions. Earlier versions were subjected to useful criticism at Montclair State College and Drew University.