Nous 38 (2004): 1-28; penultimate draft; please use published version for citation.

When is Death Bad for the One Who Dies?[1]

Ben Bradley

1. INTRODUCTION

Just about everyone believes that death is an evil for the person who dies. Yet there are well-known difficulties for this view. Perhaps the most important difficulty was raised by Epicurus in the following famous passage from his "Letter to Menoeceus":

> 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.

Epicurus might be interpreted as endorsing the following premises: (1) Anything that is bad for someone must be bad for that person at a particular time.[2] (2) There is no time at which death is bad for the one who dies. (Death is not bad for someone before she dies; it is not bad for her once she dies, because from that point on she no longer exists.) Therefore, (3) death is not bad for the one who dies.[3]

While some still find this argument attractive,[4] most find its conclusion to be obviously false; thus much effort has been expended trying to refute its premises. In what follows I will examine the premises of Epicurus' argument in more detail, and draw out the axiological and metaphysical assumptions behind those premises. I will argue

that there are good reasons to reject the second premise of Epicurus' argument. I will then provide a positive account of the time at which death is bad for the one who dies, and I will argue that my account is superior to competing accounts given by Thomas Nagel, Fred Feldman and Neil Feit. Roughly, I will argue that death is bad for the person who dies at all and only those times when the person would have been living well, or living a life worth living, had she not died when she did.[5]

2. EXAMINATION OF THE EPICUREAN ARGUMENT

Let us examine the premises of Epicurus' argument individually.

Premise 1: Anything that is bad for someone must be bad for that person at a particular time.

Thomas Nagel rejects Premise 1, and claims that death is bad "timelessly" for the one who dies (Nagel 1970, 78).[6] Nagel is a deprivation theorist; he believes that death, when bad for someone, is bad in virtue of the goodness it keeps that person from having. Death deprives us of all the things that make life valuable, such as pleasant experiences, friendships, and meaningful projects; that is why death is bad. However, according to Nagel, while the subject of a misfortune "can be exactly located in a sequence of places and times, the same is not necessarily true of the goods and ills that befall him" (Nagel 1970, 77). Nagel presents an example designed to prove his point. The example involves a man who has suffered severe brain damage, reducing him to the status of a "contented infant" (Nagel 1970, 77). Such a man would be an appropriate object of pity. But according to Nagel, Epicurus' questions about the time when death is bad for the one who

dies can also be raised about this brain-damaged man and the time when the brain damage is bad for him. There is no reason to pity this man now, because he is content; the intelligent man he used to be no longer exists (or so says Nagel), so we cannot pity him either.[7] According to Nagel, the correct way to see this as a misfortune is to "consider the person he was, and the person he could be now."[8] When we think just about the person who now exists, and consider whether that person is undergoing a misfortune, we make "a mistaken assumption about the temporal relation between the subject of a misfortune and the circumstances which constitute it" (Nagel 1970, 77).

But what is this mistaken assumption? Nagel says it is "to restrict the goods and evils that can befall a man to nonrelational properties ascribable to him at particular times" (Nagel 1970, 77). It would indeed be a mistake to limit goods and evils to ones involving only nonrelational properties; this would be to deny the existence of any kind of non-intrinsic value. Deprivation theorists believe that death is bad because of what it prevents the dead from having, namely valuable lives. So deprivation theorists believe that death is extrinsically bad for the one who dies, not intrinsically bad. Therefore all deprivation theorists must agree with Nagel that it would be a mistake to restrict the goods and evils that can befall someone to ones involving only intrinsic properties. But it is not clear that rejecting this assumption solves the problem. We might still wonder, even of an extrinsic evil, when it is evil for someone. In Nagel's example, when is the brain damage bad for the man whose brain has been damaged? Isn't it possible to provide an answer to this question? Nagel seems to think it is not.

> If we apply to death the account suggested for the case of dementia, we shall say that although the spatial and temporal locations of the individual who suffered the loss are clear enough, the misfortune itself cannot be so easily located. One must be content just to state that his life is over and there will never be any more of it. That fact, rather than his past or present condition, constitutes his misfortune, if it is one.[9]

Why must we be content with this? Why cannot the event of one's death, or the condition of being dead, or the condition of being brain damaged and helpless as an infant, constitute a misfortune? And why can we not locate the time of this misfortune? (I leave aside questions about its spatial location.) To put the question more accurately: why can we not locate the time or times at which death is bad for the one who dies?

This raises an interesting question about relations generally: must they always hold at a particular time in order to hold at all? Or are there some relations that can hold "timelessly"? Feit has argued that in order for a relation to hold at all, there must be some time at which it holds.[10] If this principle were true, it would provide solid justification for Premise 1. However, this principle seems too strong, as relations between times, such as the earlier-than and later-than relations, seem to hold without holding at any particular time.[11] But when we restrict ourselves to talking about evils, it seems plausible to suppose that they must be bad for people at particular times in order to qualify as evils. If I get sick, and my getting sick results in my being uncomfortable for a few days, it is natural to say that my getting sick was bad for me during those three days. It is also natural to say that my getting sick was not bad for me at any times before I got

sick, and – assuming that the sickness has no lingering or delayed effects – that my getting sick stopped being bad for me a few days afterwards. Furthermore, if my getting sick had not caused me any discomfort at any time, I would likely deny that it was bad for me at all.

This is just one example, so of course it does not prove that all evils must be bad at particular times. There are various possible counterexamples to Premise 1. Suppose that Andy sustains a head injury at time t, and that this injury results in Andy's living out the remaining twenty years of his life with no good or bad experiences – perhaps in a comatose state. Now suppose that there are two ways Andy's life could have gone after t had he not been injured; call them L1 and L2. In L1, Andy experiences ten years of extreme happiness after t, followed by ten years of moderate unhappiness. In L2, Andy experiences ten years of moderate unhappiness after t, followed by ten years of extreme happiness. In both L1 and L2, the degree of unhappiness felt during the unhappy period would have been severe enough that a life consisting merely of experiences like that would not be worth living, and that Andy would slightly prefer being comatose to having those experiences; but the degree of happiness felt during the extremely happy period would have far outweighed the unhappiness contained in the unhappy period. Finally, suppose that there is no fact of the matter concerning which of L1 or L2 would have been Andy's life had he not been injured; facts about the actual world and the laws of nature are insufficient to decide the matter.[12] It would seem that given these assumptions, there is no time such that we can say that Andy's injury is bad for him at that time. At any given time, there are two ways his life could have been going, one good and one bad. One life is better for Andy than his actual life at that time, and one life is worse. Yet, I

would still want to say that Andy's injury was bad for him. Both L1 and L2 are overall worthwhile existences; hence, Andy's injury at t deprived him of an existence that would have been much better on the whole than his actual comatose existence. There might be other counterexamples to Premise One. Perhaps the evil of never seeing one's beloved again is an example of a timeless evil; never getting what one deserves might be another such example.[13]

These are difficult cases. Let us grant, provisionally, that such evils are timeless; what should we conclude about the evil of death? If we follow Nagel, and deny that death must be bad at a particular time in order to be bad at all, then we have two options. First, we could deny that any evils are bad for people at particular times. This seems implausible; it prevents us from saying very natural things about the times of other evils, as in the case of sickness discussed above. On the other hand, we could claim that some evils are bad for people at particular times, but death is not one of those evils. This would enable us to hold Nagel's view and also say the things we want to say about the time of the evil of sickness, but it raises important questions. What distinguishes the evils that are bad for people at times from those that are "timelessly bad" for people? In which category does death belong? Could some deaths be timelessly bad while others are bad in the way sickness is bad? I think it is important to remember why death is supposed to be bad for us. Death prevents us from doing things we want to do, and experiencing things we want to experience. In this way it is like sicknesses and injuries, except that the latter are not permanent, and typically cause us pain in addition to preventing us from doing things. If this is so, it seems preferable to have an account of the evil of death that makes its evilness similar to that of the evils of ordinary sicknesses

and injuries, so long as such a view can be made coherent.[14] It would be better still to have a theory that enables us to count all evils as "timeful." If we maintain that the evil of death is a special, timeless kind of evil, we only give comfort to the Epicureans, who will view their doubts about the evil of death as having been justified by our need to make distinctions between ordinary evil and the evil of death. They will rightly view the making of this distinction as a desperate maneuver.

Thus I do not accept Nagel's reasons for rejecting Premise 1. Contrary to Nagel, I think that (in most cases at least) there are particular times after a person dies at which that person's death can be said to be bad for him or her. Furthermore, I will argue below, after introducing some necessary machinery, that a good deprivation theory can account for apparently timeless evils such as the ones just presented. Thus, I do not wish to reject Premise 1 of Epicurus' argument, which means we must turn our attention to the second premise.

Premise 2: There is no time at which death is bad for the one who dies.

Epicurus' defense of Premise 2 is based on a number of controversial principles. These principles include hedonism, or (roughly) the view that pleasures and pains are the only goods and evils, and what Feldman calls the "termination thesis," or the view that people cease to exist at the moment of death.[15] If we cease to exist at death, we can no longer experience any pain; thus if hedonism is true, death cannot be bad for us. Feldman gives this evocative description of Epicurus' reasoning: "At the very moment when the Reaper clutches us in his bony embrace, we go out of existence. Since the nonexistent cannot be harmed, death cannot harm us" (Feldman 1992, 4).

Some might reject Premise 2 by rejecting hedonism, and denying that one must be experiencing anything at all, or even existing, in order to be the subject of intrinsic evil; Nagel seems to take this approach (1970, 78), and perhaps Aristotle held a similar view in his discussion of posthumous goods and evils in Book I of Nicomachean Ethics. But this view is unnecessarily contentious; surely there is a way to account for the evil of death without claiming that the dead can be subjects of intrinsic evil. Others might reject Premise 2 by claiming that those who die do not go out of existence, but continue existing as corpses; some materialists hold such a view.[16] While such a view has some plausibility, it fails to account for cases of annihilation; surely when determining whether someone's death was bad for her, it matters little whether she (i.e., for such materialists, her body) remained intact after death.

Another strategy for rejecting Premise 2 may be to point out again the distinction between intrinsic and extrinsic evils. Epicurus saw a problem in attributing disvalue to death because he assumed that the only things that can be good or bad for a person are sensations; since nonexistent people do not have sensations, nothing good or bad can happen to them. But while it is at least arguable that the only things that can be intrinsically good or bad for someone are sensations, sensations are not the only things that can be extrinsically good or bad for someone. In particular, the *causes* of our sensations may be extrinsically good or bad for us. Deprivation theorists attribute extrinsic value, not intrinsic value, to death; death causes us not to have any sensations, which is worse for us than having good sensations.[17]

Still, some may find that the move to extrinsic value does not satisfy their worry. Perhaps the worry can be expressed in the following argument. Badness-for is a relation

that relates something that happens or obtains (an event or state of affairs) and a person. In order for a relation to relate things, it must relate them at a time (putting aside exceptions already discussed).[18] In order for a relation to relate things at a time, both relata must exist at that time. For example, my computer cannot be under my desk right now unless both my computer and my desk exist right now. Therefore, the badness-for relation can hold between a state of affairs and a person only at times when the person exists; so nothing can be bad for a person at times when she does not exist. Thus, there is no time at which death is bad for the one who dies. This is what has been called the "no subject" problem, and I think it is at the heart of Epicurus' argument.

On the face of it, this argument seems to prove too much. Many garden-variety, non-death-related, particular events are good or bad for a person long after the *events* are over, and thus seem to relate an existing thing to a non-existing thing. Suppose Bill falls and breaks his arm on January 14, 2001, and this event costs him his job. That particular event might remain bad for Bill long after it has stopped existing.[19] We might then ask: how can the particular event consisting of Bill breaking his arm still be bad for Bill, when that event no longer exists? The badness-for relation would be relating something that exists – Bill – with something that doesn't exist anymore – a particular event that occurred on January 14, 2001. There seem to be many other examples of relations that hold between an existing thing and a non-existing thing, many having nothing to do with goodness or badness. The is-a-great-great-grandson-of relation often relates an existing person and a non-existent person. Causal relations between events somehow relate events that do not overlap in time. Semantic relations sometimes relate words, or mental events, to things that existed long ago in the past but no longer exist (like Socrates).

Some relations even seem to relate things at times when neither thing exists; for example, the is-a-more-popular-former-President-than relation currently relates Abraham Lincoln and James K. Polk, neither of whom currently exists.

But not everyone believes that the claims just made are literally true. Those who deny it are called Presentists. Presentism, in the sense relevant here, is the view that only present objects exist – putative past and future objects are unreal.[20] I think Premise 2 of Epicurus' argument presupposes the truth of Presentism; thus we may escape the argument by rejecting Presentism. Rejecting Presentism is a good idea for other reasons, including the need to preserve our ability to refer to past objects like Socrates and to account for the truth of our history books. If attributions of badness to death are in the same boat as references to Socrates, I think they are in good company. Rejecting Presentism solves the no-subject problem. If the past is real, there is no difficulty in locating a subject of the evil of death; it is a past person.[21] Presentism-based objections to deprivation accounts arise in different forms; I will return to discussion of issues relating to Presentism once I have presented my positive account of the time of the evil of death.

3. AN IMPROVED DEPRIVATION ACCOUNT

If we are right to reject Nagel's claim that death is bad without being bad at any particular time, and if we have successfully rejected Epicurus' argument by rejecting Presentism, the question still remains: <u>When</u> is death bad for the one who dies? This question may be approached by first asking the question: <u>Why</u> is death bad for the one

who dies? According to Feldman, someone's death is bad for that person just in case it makes the intrinsic value of that person's life lower than it would have been had the person not died then. So Feldman's view involves a comparison between the way a person's life actually goes and the way it would have gone had things gone differently. Feldman analyzes the "bad-for" relation in terms of the subject-relative values of possible worlds; when we say that s's death at t in the actual world is bad (good) for s, we are saying that the value of the actual world for s is lower (higher) than the value for s of the "closest" world where s does not die at t.[22]

So how does Feldman answer our original question: <u>when</u> is s's death bad for S? "It seems clear to me that the answer to this question must be 'eternally'" (Feldman 1991, 221). When we say that s's death is bad for s, we are comparing the values for s of two possible worlds; according to Feldman, those two worlds stand in their value-relations at all times. It is always true that the nearest world where s does not die at t is better for s than the actual world; therefore, Feldman says, s's death at t is bad for s eternally.[23]

Feldman's view avoids the dubious claim that the "bad-for" relation holds but does not hold at any time. I think Feldman's view does express something true – or at least, it correctly answers one possible question we might be considering – but it also leaves out part of the truth. Even if determining the badness of a death for a person involves a comparison of possible worlds, it still seems wrong to say that if a world is good for a person, it is good for that person at all times; I would rather not say that someone's death remains bad for him thousands of years after it occurs, or that it was bad for him before he died, or that it was bad for him before he even existed. The problem is

that there are at least two different things we might be asking when we ask when someone's death is bad for that person. Feit clearly describes this difference:

> Consider again the question: When is Abe Lincoln's death bad for him? Feldman takes the question to be equivalent to this: When is it true that his death is bad for him? Hence Feldman is led to his version of eternalism about the evil of death. On the other hand, I take the question to be asking this: At which times t is it true that his death is bad for him at t? (Feit forthcoming, 23)[24]

We may understand this second question to be asking a question about the "badness-at-a-time-for" relation: At which times t is it true that s's death is bad-at-t for S? Feit attempts to analyze badness-at-a-time in a way that avoids Feldman's eternalism. Before I get to Feit's view, let me explain my own conception of badness-at-a-time.

Let us examine the question of when any event is bad for a person, and apply our results to the case of death. Suppose I stub my toe at time t1, resulting in some discomfort that lasts for three days, and preventing me from hiking or playing tennis during that time. Suppose the toe-stubbing does not have any lingering effects; it does not prevent me from completing important projects, does not cause me to lose my job, and is soon forgotten. Intuitively, we should say that my stubbing my toe at t1 is (extrinsically) bad for me during my three days of discomfort, but once I recover it is no longer bad for me.[25]

It seems to me that the reason we say that my stubbing my toe at t1 is bad for me during the three days following t1 is that my life would have been going better for me during that time had I not stubbed my toe. That is, the intrinsic value of the portion of my life that takes place during the three days after t1 is lower than it would have been had I not stubbed my toe. We also want to say that my stubbing my toe is no longer bad for me once the pain goes away and I am able to resume my normal activities. Its effect on me, and on the value of my life, is over. Note that it continues to be the case that my stubbing my toe had a bad effect on the value of my life, and this will eternally remain so; but it is no longer the case that it is currently having bad effects.

It is not hard to see how to say something basically similar in the case of death. The death of s at t1 is bad for s at t2, where t2 is later than t1, provided that at t2, the life s would have been living had s not died at t1 would have been better for s than no life at all. Thus, Socrates' death has long ceased to be bad for Socrates, since had he not died when he did, he would still have been dead for many years; but JFK's death might still be bad for him today, since he might well be living happily had he not been assassinated. Just as, on Feldman's view, we can talk of a world being good or bad for a person, I suggest we can also talk of a time being good or bad for a person. The time when I felt pain from the toe-stubbing was a bad time for me; some of the times when Socrates was dead, but would have been alive had he not drunk the hemlock, were bad times for Socrates. Talk of times being good or bad for people is ubiquitous; some nostalgic types say that their college years were "the best years of their lives"; we often talk of someone who is "going through a bad time right now." I think determining the values of times for people is crucial to answering questions about the time of the evil of death.

This raises an interesting question: How do we determine the value of a time for a person? I think that the intrinsic value of a time for a person s ought to be determined in a way similar to the way that the intrinsic value of a world for s is determined. On a simple view, to determine the intrinsic value of a world w for s, we sum the basic intrinsic values of the states occurring in w that have s as a subject.[26] We might determine the intrinsic value of a time t for s at world w in a similar way as follows:

> IVT: The intrinsic value of a time, t, for s, at world w = the sum of the basic intrinsic values of the states obtaining at w during t that have s as a subject.[27]

Some features of IVT are worth pointing out, as they will become relevant later. First, IVT entails a more general principle that we might call "internalism" about the intrinsic values of times for people: the intrinsic value of a time t for a person s is entirely determined by the intrinsic values of states obtaining at t – the intrinsic values of states obtaining at times other than t are irrelevant except insofar as they affect the determination of the intrinsic values of states obtaining at t. This internalist principle needs to be distinguished from a stronger principle one might also hold, according to which the intrinsic value of a state of affairs is determined independently of the values of any other states.[28] While I think the stronger principle is justified, defending it is not necessary for my purposes here. IVT is compatible with views about intrinsic value according to which the intrinsic value of something can depend on its relations to other things.

Second, IVT allows us to determine the subject-relative intrinsic values of not only instants of time, such as precisely 12 noon on Jan. 26, 2001, but also durations of time, such as the period from 12 noon to 1pm on Jan. 26, 2001. To determine the intrinsic value of a duration of time for a person, we simply look at the basic intrinsic values of all the states obtaining during that time. I think these are both desirable features of the view.

We may now determine the overall value of a state of affairs for a person at a time by comparing the intrinsic value of the time for the person with what the intrinsic value of the time would have been for the person had that state of affairs not obtained. This involves comparing the intrinsic value of the time for the person in the actual world with the intrinsic value of the time for the person in the nearest possible world at which the state does not obtain. We may state this view in the following way:

> OVT: The overall value of p for s at <w,t> = the intrinsic value of t for s at w minus the intrinsic value of t for s at the nearest world to w at which p does not obtain.

In the case of death, we may assume for now – perhaps contrary to Aristotle – that all times after a person's death have intrinsic values of zero for that person (I will return to this assumption below). Thus, according to OVT, s's death at t1 is bad for s at time t2 iff the intrinsic value of t2 for s, at the nearest world in which s does not die at t1, is greater than zero; s's death at t1 is good for s at t2 iff the intrinsic value of t2 for s, at the nearest world in which s does not die at t1, is less than zero.

It is worth pointing out that OVT inherits the internalism of IVT; it entails that the value of any state p for subject s at time t is wholly determined by the s-relative intrinsic values of states obtaining at t, in actual and counterfactual circumstances. The intrinsic value of a state obtaining at any other time is irrelevant, except insofar as it affects the determination of the intrinsic value of a state that obtains at t.

One consequence of considering the value of a life at both durations and instants is that there may be different yet equally true answers to the question "is s's death bad for s now?" Suppose we know that if s were alive, s would have been in great pain at this very moment (perhaps due to a planned operation), but that s's life generally would have been happy. One person might say that it is good for s, now, that s is not alive; another might say that it is bad for s, now, that s is not alive. If the first person is talking about the present moment, what he says may be right. If the second person is talking about a larger time, that includes the duration of the planned operation but also includes many other times, what he says may be right.

OVT also enables us to avoid the apparent counterexample to Premise 1 of the Epicurean argument presented in Section Two above. In the counterexample I presented, it seemed that there was no time at which Andy's injury was bad for him. But this is true only if we are talking about moments; there are many durations at which Andy's injury can be said to be bad for him, including the twenty-year period after his death. No matter whether L1 or L2 would have been Andy's life, that twenty-year period would have been better for him than a life spent comatose. We might use a similar strategy to account for other apparent cases of timeless evils, such as never seeing one's beloved again or never getting what one deserves. Take the latter case; suppose Fred deserves a teaching award

but never gets it. The closest possible world in which Fred gets the teaching award might be a world in which he gets the award in 2002, or it might be a world in which he gets the award in 2003. Suppose getting the award would make Fred's life better for some period of time afterwards. Then his never getting the award could be said to be bad for him during a period of time starting sometime in 2002, and ending sometime after 2003. This does not seem implausible. There might be no way, even in principle, to determine a smaller period of time during which his never getting the award is bad for him; but that isn't enough to show that it is a timeless evil. Other apparently timeless evils could get the same treatment.[29]

In formulating OVT, I have bypassed certain controversies concerning counterfactuals. In particular, in using OVT to evaluate s's death, I have been supposing that the closest possible world in which s does not die at a particular time is a world in which s's life went exactly the same up until (or shortly before) the time of s's death in the actual world. That is, I have ignored the possibility of "backtracking" counterfactuals.[30] I have ignored them not because I think they are always false; in fact, following Lewis, I think some backtracking counterfactuals are true. However, also following Lewis, I think they are true only in "special" or "non-standard" contexts. In most contexts in which we are evaluating someone's death, backtrackers are ruled out. In an unusual context in which backtrackers are permitted, it might actually turn out to be true that s's death was bad for s before s died. Suppose, for example, that I know that Jane is doomed to die from a congenital disease that always causes great pain before death. Suppose I have been out of touch with Jane for a while, and then come to find out that she died two months ago. Knowing already how she must have died, I might say that

it is bad for Jane that she died, not just because of what her death deprived her of, but because the fact that she died indicates that she must have been in great pain before then. In such a context, I am supposing that the closest possible world in which Jane did not die two months ago is a world in which Jane was not in great pain a little more than two months ago. Such special contexts are required to account for what has been called "signatory" value.[31] Belief in backtracking counterfactuals and signatory value is not, however, required for belief in OVT; in fact, many people are inclined to reject backtrackers and are skeptical of the existence of signatory value, which is fine for my purposes here. But belief in backtrackers might help OVT account for certain problematic cases, as I will show in Section Five below.

I think OVT has more plausible consequences than other views. OVT is similar to Feldman's view, except that whereas Feldman's view allows us to compare only the subject-relative intrinsic values of entire worlds, OVT would have us compare the subject-relative intrinsic values of temporal segments of worlds. Thus OVT does not entail that death is eternally bad for the one who dies; it enables us to make different judgments about different times. I view this as an advantage of OVT over Feldman's view. OVT is also superior to Nagel's view. Unlike Nagel's view, OVT allows for uniform treatment of the evils of death and sickness. Nagel's treatment of death as a timeless evil is to be avoided if possible, and OVT allows us to avoid it.

4. FEIT'S VIEW

My view of the time of the evil of death is importantly similar to a view recently presented by Neil Feit. Feit shares my dissatisfaction with the views of deprivation theorists like Nagel and Feldman. Surely Abe Lincoln's death stopped being bad for Lincoln long ago; likewise, his death hadn't yet started being bad for him before he died. If we could determine the time when his death started to be bad, and the time when it stopped being bad, we could say that Lincoln's death was bad for him at times between those two times, but at no other times. This is Feit's strategy; thus, Feit must give an account of the starts-being-bad relation and the stops-being-bad relation. He explains these two concepts as follows:

> D3: p starts being bad for s at t if and only if (1) the overall value of p for s is negative, (2) the value that accrues to s before t equals the value that accrues to s before t in the nearest ~p-world, and (3) there is no time t' later than t, such that the value that accrues to s before t' equals the value that accrues to s before t' in the nearest ~p-world. (Feit 17)
>
> D4: p stops being bad for s at t if and only if (1) the overall value of p for s is negative, (2) the value that accrues to s after t equals the value that accrues to s after t in the nearest ~p-world, and (3) there is no time t' earlier than t, such that the value that accrues to s after t' equals the value that accrues to s after t' in the nearest ~p-world. (Feit 19)

Feit understands overall value, as I do, to be a function of intrinsic and extrinsic value. By the "nearest ~p-world," Feit means the nearest world to the actual world in which p does not obtain. He then defines overall badness-at-a-time in the following way:

> OBAT: State p is overall bad for s at t if and only if (1) p starts being bad for s at or before t, and (2) p does not stop being bad for s before t. (Feit 23)

Feit's view is complicated, but a simple example can illustrate how the view is supposed to work. Let p be the state consisting of my getting sick; suppose p occurs at t1. Suppose that at t1, I start feeling bad, and I feel bad intermittently for three days as a result of p. Then at t2, I stop feeling bad; p has no other effects on my life. Suppose that had I not gotten sick at t1, I would have felt fine between t1 and t2. D3 entails that p starts being bad for me at t1, because (i) p is overall bad for me, (ii) the intrinsic value of the portion of my life that takes place before t1 is the same as it is in the nearest world where p does not obtain (since p has no effects on my life before it occurs), and (iii) at all times t after t1, the intrinsic value of the portion of my life that takes place before t is lower than it is in the nearest world where p does not obtain. D4 entails that p stops being bad for me at t2, because (i) p is overall bad for me, (ii) the intrinsic value of the portion of my life that takes place after t2 is the same as it is in the nearest world where p does not obtain (since p has no effects on my life after t2), and (iii) at all times t before t2, the intrinsic value of the portion of my life that takes place after t is lower than it is in the nearest world where p does not obtain. OBAT entails that p is bad for me at all times between t1 and t2, and at no other times. This seems plausible. It is easy to see how to

apply OBAT to the case of the badness of death. In virtue of D3, OBAT entails that someone's death cannot be bad for her before it occurs.[32] In virtue of D4, OBAT entails that someone's death cannot be bad for her hundreds of years after it has occurred. These results also seem plausible.

However, I think there are other cases where OBAT yields less plausible results. OBAT entails that in order for a state of affairs to be bad for a person at a particular time, it must be bad for that person overall. Feit discusses one possible counterexample to this claim, involving the taking of an experimental drug that causes immediate pain but cures a disease. Let p be the taking of the drug by John; suppose p occurs at t1, and suppose the following chart determines the value of p for John: [33]

Time	Value of w for John	Value of nearest ~p-world for John
Before t1	+100	+100
From t1 to t2	-50	0
From t2 to t3	+300	-200
After t3	0	0

According to Feit, p is not overall bad for John; about this I think he is right. But if p is not overall bad for John, then OBAT entails that p cannot be bad for him at any time – not even from t1 to t2, when p is causing him pain. This is because D3 entails that in order for something to start being bad for someone, it must be overall bad for that person; and OBAT entails that nothing can be bad for a person at a time unless it has started being bad for that person by that time. Since p is not overall bad for John, it never starts

being bad for John; so according to OBAT, there is no time at which p is bad for John. Intuitively, it seems better to me to say that p is bad for John from t1 to t2, but is good for him from t2 to t3, and is good for him overall.

According to Feit, "the experimental treatment may be the best thing that ever happens to [John], and so it is hard to see why we would want to say that it is bad for [John] at certain times" (Feit 11). I think there are good reasons to want to say this. The most important reason is that, as noted above, if we are asking whether p is good for s at t, we shouldn't need to look at the values of states obtaining at times other than t in order to figure this out. We need to look at the values of states obtaining at t in the actual world, and at t in the nearest ~p-world, and that is all. That is, Feit's view violates the internalism requirement. It seems odd to argue that p is not bad for s at t on the basis of the fact that p is not bad for s at some time other than t, or the fact that it is not bad for s overall; to argue in this way is to change the subject. To further illustrate, suppose the values in the example were changed in the following way:

Time	Value of w for John	Value of nearest ~p-world for John
Before t1	+100	+100
From t1 to t2	-50	0
From t2 to t3	+300	-200
From t3 to t4	-500	+700
After t4	0	0

By adding the period from t3 to t4, we get OBAT to give the result that p is bad for John from t1 to t2.[34] OBAT determines that p is bad for John from t1 to t2 on the basis of things that happen to John long after t2. Note that OVT, the view I presented in Section Three, avoids this problem. In this example, OVT entails that p is bad for John from t1 to t2, good for John from t2 to t3, bad for John from t3 to t4, and neither good nor bad for John before t1 or after t4.

In requiring us to determine the values of temporal segments of worlds for people – for example, the value accruing to John during the pre-t1 period of w – OBAT requires an account of the intrinsic value of a duration of time for a person. OVT also requires such an account. Given such an account, OVT makes the determination of the overall value of a state of affairs for a person at a time fairly simple: we compare the intrinsic value for the person of that time at one world with its intrinsic value for the person at another world. In the first example of the experimental drug, OVT entails that p's overall value for John from t1 to t2 is –50; p's overall value for John from t2 to t3 is +500; p's overall value for John from t1 to t3 is +450. Feit's view is relatively complex; to determine when a state of affairs is bad for a person, given the subject-relative intrinsic values of durations of time in the actual world and another possible world, we must first determine when the state of affairs starts being bad and when it stops being bad. This seems like a superfluous step, and if my criticism is correct, the extra step leads to problems. Furthermore, Feit's view leaves out the possibility of determining an <u>extent</u> to which a particular time is overall good or bad for a person. OBAT tells us that a time is good or bad for a person, but it does not tell us <u>how</u> good or bad it is. Since OVT does, it is more complete than OBAT; I view this as another advantage of OVT over OBAT.

5. OBJECTIONS

Despite the advantages OVT has over other views, it still faces a number of possible objections, of which it is possible to discuss only a few.[35] The four objections that follow are, I think, the most important ones, and I hope that my responses to them will help to clarify my view.

Objection 1. OVT makes true certain claims that some might think should turn out false. For example, consider the following sentence:

The Aunt Alice sentence: Aunt Alice's death was good for her during the week after she died, but bad for her during the next month.

If Aunt Alice would have been miserable for a week, then happy for a month had she not died when she did, OVT entails that the Aunt Alice sentence is true. Some would consider this to be a counterexample to OVT.[36]

While the Aunt Alice sentence might sound odd, I believe sentences of the sort are often true. Let us add some detail to the story of Aunt Alice. Suppose she was suffering from a very painful illness; she underwent a risky procedure that, if successful, would have had her feeling well in a week; instead the procedure failed, and she died. Shortly after her death, we might say things like "Aunt Alice was in so much pain; it's good that she is no longer suffering." Later on, we might say something like "it's too bad Aunt Alice is no longer with us; she would have enjoyed seeing little Timmy's recital."

The Aunt Alice sentence is really nothing more than a way of summing up both of these sentiments; understood in this way, the Aunt Alice sentence loses its oddness and seems true to me. Thus I do not consider the Aunt Alice sentence to be a counterexample.

Objection 2. In formulating OVT, I have made some explicit assumptions that may seem questionable, including the assumption that, for a person who does not exist at a given time, the intrinsic value of that time for that person is zero. OVT requires us to make what Harry Silverstein calls a "life-death" comparison (Silverstein 403). That is, OVT requires us to compare two things: (A) the s-relative intrinsic value of a time t at a world where s is alive, and thus exists, at t; and (B) the s-relative intrinsic value of t at a world where s is dead, and thus does not exist, at t. Some, perhaps including Silverstein, will object that while we can sensibly calculate the first term (A) of the comparison, we cannot calculate the second.

> Suppose, for instance, that, using a scale from +10 to –10, A calculates correctly that his average A-relative value level (e.g., his average happiness level), if he continues living, will be +2. What the Epicurean view claims… is that A cannot coherently use such a calculation as the basis for a prudential choice between life and death, since, as no value on the scale (including 0, the "midpoint" value of neutrality or indifference) can intelligibly be assigned to A's death, there is nothing against which the value +2, however "rationally" derived, can intelligibly be weighed. (Silverstein 1980, 410)[37]

But why can't we assign an s-relative intrinsic value of zero to a time when s does not exist? According to IVT, this value will be determined by adding the basic intrinsic values of the states obtaining during t that have s as a subject. So the question becomes this: are there any states that obtain during t that have s as a subject, or are "directly about" s, if s does not exist? If there are, I presume that their basic intrinsic values will all be zero; summing their basic intrinsic values will give us a total of zero s-relative intrinsic value, as I have been assuming all along. However, some might claim that if s does not exist at t, no states with s as a subject can possibly obtain at t. To put this claim another way, there can be no singular propositions about s that are true at a time when s does not exist.[38] The reason is simple: if s does not exist, then a proposition with s as a constituent cannot exist either; and a nonexistent proposition can't be true. If this is so, then there may be a problem for IVT. How do you calculate the sum of zero addends?

This objection amounts to a restatement of the "no subject" problem, and like the previous formulation, it is based on Presentism. Again, I feel the best way to respond to the objection is to reject Presentism, in virtue of the fact that Presentism faces objections involving relations between things that do not overlap in time, such as the ones discussed above in Section Two.[39] If we reject Presentism and maintain that the past is real, we will have no difficulty locating some singular propositions about past objects and persons.[40]

But suppose we want to accept Presentism. What are our options? (1) We could find a way to make Presentism compatible with the existence of singular propositions about nonexistent objects. Perhaps this would involve a commitment to "haecceities." Alvin Plantinga explains the notion of haecceity as follows: "An haecceity is one kind of

individual essence: a property that is essential to its owner, and essentially unique to is owner, in the sense that it is impossible that there be something else that has it" (Plantinga 1995, 199). The haecceity of JFK would simply be the property of being identical to JFK; this property is essential to JFK, and nothing but JFK could have that property. According to this proposal, propositions involving this haecceity would qualify as singular propositions about JFK. But the primary difficulty for this approach is that it is hard to see how the property of being identical to JFK could exist at times when JFK no longer exists.[41] If JFK's haecceity cannot exist without JFK, this approach will not help us.

(2) There is a more complicated solution for those who wish to be Presentists. We could revise our view about the bearers of basic intrinsic value. To this point, I have followed Fred Feldman in supposing that the bearers of basic intrinsic value are singular propositions: propositions that are directly about a person.[42] But we could hold instead that the bearers of basic intrinsic value are general propositions. That is, where we had been supposing that the bearers of basic intrinsic value were propositions (states) such as

S1 Socrates feels pleasure to degree 10 at time t,

we could instead say – supposing that Socrates has certain distinctive properties P1-Pn – that the bearers of basic intrinsic value are propositions such as

S2 Someone with P1-Pn feels pleasure to degree 10 at time t.

Presentists should not have any objection to the existence of propositions like S2. If such propositions can acceptably serve as basics, there would be no difficulty with locating some propositions and adding up their basic intrinsic values to give us a number for comparison. However, a serious difficulty confronts this view: what does a proposition like S2 have to do with Socrates? If we are determining how well Socrates' life went for him, it is hard to see why we should look at states like S2.[43]

(3) There is at least one more possible avenue. I originally assumed that to find the s-relative value of a time at a world, we must first find all the states of affairs with s as a subject that obtain at that time, and then add up their basic intrinsic values. This led to a problem involving times at which s does not exist, because at such times there are no such states. But we could take a different route. We could introduce the idea of subject-relative basic intrinsic value. For any subject s, every state of affairs – not just those about s – would be assigned an s-relative basic intrinsic value. Any state of affairs that does not have s as a subject would be assigned an s-relative basic intrinsic value of zero, since such states do nothing directly to improve or worsen the world for s. For example, if eudaimonism were true, the only states of affairs with non-zero s-relative basic intrinsic value would be states of the form s is happy/unhappy to degree n at time t. All other states would have an s-relative basic intrinsic value of zero.

If we take this route, there would be no problem in obtaining our second term of comparison in a life-death comparison. If s does not exist at t, we could say that t has an s-relative value of zero, since all the states that obtain at t would have zero s-relative basic intrinsic value. Introducing a notion of subject-relative basic intrinsic value is not an ad hoc solution; it makes sense if we follow Temkin in making a distinction in the

theory of value between theories about self-interest and theories about outcomes.[44] Regular old basic intrinsic value would be the basic intrinsic value concept used in a theory about outcomes; subject-relative basic intrinsic value would be the basic intrinsic value concept used in a theory about self-interest. A total theory of value would encompass both types of theory.

However, this option is in essentially the same boat as Presentism-plus-haecceities. It is just as problematic to claim that the property of being good for Socrates can exist when Socrates doesn't exist as it is to say that the property of being identical to Socrates can exist when Socrates doesn't exist. If that is so, then making this distinction, while perhaps necessary anyway, will not help to reconcile my view with Presentism.

Thus we are left with a dilemma. Presentism really does seem to be incompatible with any plausible interpretation of OVT, or at least any interpretation that gives us the result that death is bad for the one who dies. So we must either reject Presentism or reject OVT. I suspect that deprivationist alternatives to OVT, when combined with Presentism, will be equally unable to account for the evil of death. All deprivation theorists claim that a person's death is bad because of the goods of which it deprives that person; so all deprivation theorists require there to be a subject who is being deprived of some goods, and Presentism appears incompatible with the existence of such a subject. If so, then we must either reject Presentism or reject the view that death is bad for the one who dies. The view that death is bad for the one who dies seems to me to be what McMahan calls a "fixed point" or "starting point" in ethics – a conviction that would require extremely convincing reasons to overturn, if it could be overturned at all.[45] In

light of the difficulties already noted, Presentism cannot provide the basis for overturning such a fixed point; it seems more reasonable to reject Presentism.

Objection 3. OVT has us compare the way someone's life actually goes, with her death occurring at some time t1, with the way it goes in the nearest possible world in which she does not die at t1. As a result, it may be thought that OVT faces problems involving overdetermination and preemption. Jeff McMahan provides a good example:

> The Young Pedestrian. A young man, aged twenty, absentmindedly steps off the curb into the path of a bus and is instantly and painlessly killed. During the autopsy, it is discovered that he had a hitherto silent cerebral aneurysm that would inevitably have burst within a week if he had not been hit by the bus. And the bursting of the aneurysm would certainly have been fatal. (McMahan 2002, 117)

OVT seems to entail that the Young Pedestrian's death was not very bad for him, since it deprived him of very little life. This seems like an implausible result.[46]

One diagnosis of the alleged problem with OVT involves the notion of "inherited evils." The idea of inherited evil is easy to understand. Let t1 be the time when YP got hit by the bus, and let t2 be the time when YP's aneurysm would have burst. If YP had not died at t1, he would have died at the later time t2. OVT entails that YP's death at t1 is neither good nor bad for YP at any times after t2. But now move to the point of view of the world in which YP dies at t2 instead. At that world, YP's death at t2 is bad for him at times after t2, since his death at t2 deprives him of a still longer life. According to

what McMahan calls the "Inheritance Strategy," YP's actual death at t1 inherits the evil of his counterfactual death at t2. Thus, his actual death at t1 would turn out to be very bad for him, as we intuitively think it is. We may be led to the Inheritance Strategy by the following line of reasoning: YP's death at t1 could not be less bad for him than his later death at t2 (McMahan 2002, 121). Whatever his possible death at t2 would deprive him of, his actual death at t1 deprives him of that much and more.

But the Inheritance Strategy is fatally flawed. As McMahan points out, "for each successive death a person might have died, we can always ask what would have happened if that death would not have occurred" (McMahan 2002, 126). If YP had not died at t1, he would have died at t2; if he hadn't died at t2, he would have died at t3; if he hadn't died at t3, he would have died at t4, and so on. Surely YP's death does not inherit the evils of all these possible deaths. The Inheritance Strategy leaves us with no obvious way to determine how bad YP's death is for him. McMahan calls this the "Problem of the Terminus," and rightly sees it as a good reason to reject the Inheritance Strategy (McMahan 2002, 127).

The failure of the Inheritance Strategy is not sufficient to establish the truth of a view like OVT; we still need a way to provide the correct evaluation of YP's death. There are a number of strategies we may employ in order to account for examples of overdetermination and preemption.[47] Insofar as we want to compare YP's actual life with one that lasts much longer, we might revise OVT in such a way that we can compare his actual life with his life in some other salient world – not necessarily his life in the closest world where he does not get hit by the bus. The salient world might be one in which YP does not have an aneurysm, even if that world is not the closest one where he doesn't get

hit by the bus. Such a view would retain some of the spirit of OVT, since it would still involve a comparison between subject-relative intrinsic values of possible worlds; it would merely remove the apparently problematic insistence on always comparing the actual world with the closest world where some death does not occur. The closest world would usually be the most salient one, and therefore the one we would choose for comparison; but in apparently unusual cases such as ones involving overdetermination or preemption, more distant worlds might be more salient.

This strategy might seem ad hoc in the absence of further explanation of what makes a world salient for comparison. Perhaps the inherent vagueness of the notion of closeness of possible worlds could provide us with all the flexibility we need to account for these cases in such a way that the evaluation of a death rides piggy-back on a determination of the closest possible world. What is the closest possible world in which YP does not die at t1? Typically, it is a world in which the past as of t1 (or shortly before t1) is completely similar to the actual past as of t1, so that if YP had not died at t1, YP would still have died of an aneurysm shortly after t1 anyway. To suppose that YP would not have had an aneurysm in the first place would be to make a gratuitous historical change. But such a change is not always gratuitous. To illustrate, consider two obvious answers we might give to a question about YP's death:

"Was it bad for YP that he died at t1?"

Answer #1: Yes, very bad. YP was a young man at t1. He would have been better off dying at a much later time.

Answer #2: Yes, but not very bad, because he would have died within a week anyway from a burst aneurysm.

Either of these answers (and probably many others) could be correct, depending on the context. We would typically give Answer #1. When we find out about YP's aneurysm, our answer might change to Answer #2; but even when we know about the aneurysm, we still feel a strong pull to Answer #1. I think the reason is that we are pulled in different directions by different resolutions of the vagueness of the question, and we have different comparisons in mind at the same time.

What counts as the nearest world where YP does not die at t1? Sometimes it is permissible to count a world where YP dies only a week later; but sometimes the nearest world must be one in which he dies at a time long after t1, even if that makes certain backtracking counterfactuals come out true (e.g., if YP hadn't died at t1, he wouldn't have had an aneurysm before t1). Answer #2 seems right if we are asking "was it bad for YP that he died (by being hit by a bus) at t1 (rather than dying of the aneurysm shortly thereafter)?" On the other hand, Answer #1 seems right if we are asking "was it bad for YP that he died at t1 (rather than in old age)?" Different ways of asking the question make different answers turn out to be correct; the parenthetical insertions help to clarify exactly what comparison we want to make. The way of asking the question that favors Answer #2 emphasizes the <u>cause</u> of YP's death; the way that favors Answer #1 emphasizes its <u>timing</u>. We are led to think there is a problem with the YP because the example emphasizes details concerning the cause of YP's death. But the cause of the death is typically much less important in the evaluation of the death than its timing; the timing is what determines (in conjunction with other obvious facts) how bad someone's death is. [48]

Of course, even once we have focused on the timing of YP's death, different answers may be correct depending on whether we want to know whether death at precisely t1 was bad for YP or whether death at roughly t1 was bad for YP. Consider the following answer to the question:

Answer #3: It's not bad that YP died at t1. If he hadn't died at precisely t1, he would have died a millisecond later or earlier. That wouldn't have made a difference. There is nothing particularly special about t1 that makes it a bad time to die.

Usually it is not permissible to count a world where YP dies a millisecond after t1 as the closest world where he does not die at t1, as illustrated by the oddness of Answer #3. Though Answer #3 seems strange, it should be clear that a world in which YP gets hit by a bus at t1-plus-one-millisecond is more similar to the actual world in certain respects than one in which the bus does not hit him at all. This seems to favor Answer #3 over #1 or #2. A context in which Answer #3 makes sense is one in which it is asked "was it bad for YP that he died at t1 (rather than any other time, such as a millisecond earlier or later)?" The reason Answer #3 seems odd is that we do not usually interpret the question in this way. While we are usually interested in the timing of a death for the purposes of evaluating it, we are not interested in the time because of anything having to do with the time itself; we are interested in the time only because we want to know how much of a good life has been taken from the deceased.

Even if I have shown that OVT can account for the claim that was initially supposed to be correct – that YP's death was very bad for him – some will be unsatisfied. OVT accounts for not only that claim, but also the apparently implausible claim that YP's death was not very bad for him. The commitment to multiple answers to the question of

how bad YP's death was for him might seem problematic, but I do not think so. It provides the best explanation of the fact that we are pulled in different directions by the example. Any theory that entails that YP's death is very bad for him absolutely, and requires us to compare his death with a much later possible death, does violence to a conflicting intuition: namely, that it wasn't so bad that the YP died when he did, given that he had an aneurysm. When the aneurysm is emphasized, I do feel a strong inclination to say that his actual death by bus was not all that bad for him. But I also feel a pull to say the opposite, because it is natural to focus on the timing of the death rather than its cause. Incorporating context-sensitivity provides an explanation for the conflicting pulls.

Those who find this approach unsatisfactory might wish to abandon talk of possible worlds altogether. I suspect that such objectors might wish to divorce causation from counterfactuals, and say that YP's actual death at t1 causes him to miss out on many years of life even though the closest world where he doesn't die at t1 is a world where he dies shortly afterwards. According to this line of reasoning, the existence of the aneurysm does not prevent us from saying that the bus accident caused YP to miss out on years of life. If the bus had not hit him, the aneurysm would have caused him to miss out on years of life; but that doesn't prevent us from saying that the bus accident actually caused him to miss out on those years. We can easily reformulate OVT in a way that is compatible with this approach. Instead of talking about possible worlds, we can stick to talk about causation, and distinguish those states of affairs that a death causes to obtain from those that it causes not to obtain. In the example of YP, YP's death at t1 causes certain states of affairs to obtain (such as the state that YP has no experiences at t1-plus-

ten-years), and causes other states not to obtain (such as that YP is happy at t1-plus-ten-years). Call the conjunction of all states of affairs that p causes to obtain at time t the total consequence of p at t; call the conjunction of all states of affairs that p causes not to obtain at t the total prevention of p at t. We could then formulate a possible-worlds-free version of OVT roughly as follows:

> OVT (revised): The overall value of p for s at t = the intrinsic value for s of the total consequence of p at t, minus the intrinsic value for s of the total prevention of p at t.

This revised version of OVT can account for the badness of YP's death, provided we have a theory of causation that entails that the bus accident caused him to miss out on many years of valuable life. I do not know whether this revised version of OVT is an improvement over the version that employs possible worlds apparatus, but perhaps some will prefer it. In any case, I think we now have plausible versions of OVT that can account for a number of different intuitions concerning the Young Pedestrian.

Objection 4. The final objection I wish to discuss is perhaps the most serious.[49] It might be objected that while my proposal gives a sense to the locution 'p is bad for s at t,' the sense I have given to it does not correspond to any ordinary way of using that locution. The following analogy could be given. Someone might offer a definition of the locution 'x kicks y at t' that makes it possible for the following statement to turn out true:

The Nixon sentence: Nixon kicked JFK at 7pm on January 23, 1979.

Perhaps the definition would entail that the Nixon sentence is true provided that Nixon moved his foot in a certain area, and that JFK would have been there had he existed.[50] We would not think that the definition shows it is possible to kick someone who doesn't exist, any more than we would think it possible to show that dogs have five legs by calling a tail a leg.

Similarly, it may be argued, OVT fails to provide a view that is relevant to traditional questions about the evil of death. Epicurus wanted to know whether death can harm the one who dies; OVT provides a way to understand claims about harming the dead in a way that makes them come out to be true, but Epicurus and other ancient philosophers did not understand the claims in the way I suggest. Using words like 'harm' in their conventional senses, it is just as impossible to harm those who don't exist as it is to kick them.

I think this is a mistake. First, there is an important disanalogy between kicking and harming. As I pointed out above, there are many cases in which an event causes harm to someone at times when the event no longer exists. Consider the explosion at Chernobyl. The explosion caused harm to many people; but for most of the victims, at the time they were harmed by the explosion, the explosion no longer existed. Furthermore, some of those harmed did not even exist yet at the time of the explosion. So it is clear that the harmful event and the harmed person need not both exist at the time of the harm, while in the case of kicking, both the kicker and the kicked must exist at the time of the kicking. If the harmful event need not exist at the time of the harm, why must the harmed person exist at that time?

I think OVT does capture at least one important way of thinking about harm that we might have in mind when thinking about the harm death does to those who die. There may be another sense of the word 'harm' according to which the person harmed must exist at the time of the harm. For example, many things that harm someone do so by causing something intrinsically bad to happen to that person – for example, causing them to feel pain. You can't harm the dead in this way. But this is not the only way to think about harm. We can harm people by preventing them from having good experiences. One way to do this is to cause them to stop existing. That is what death does, and that is how it harms people. If this does not seem like a legitimate use of the word 'harm,' then I suggest there may be ways for events to be bad for people without harming them; if we can show that death is bad for the one who dies, we will have done enough to answer the Epicureans even if we can't show that death harms the one who dies.

But even if we have done enough to answer one Epicurean argument, we may not have satisfied everyone. Kai Draper argues that at least some things that are bad for us – in the sense of being comparatively bad for us – are not genuine evils (Draper 1999). They are not evils worth caring about. For example, not finding Aladdin's lamp is comparatively bad for me, but is not a genuine evil, since it does not merit any negative feelings (Draper 1999, 389). OVT entails that my not finding Aladdin's lamp yesterday is very bad for me today, since I would be much better off today if I had found Aladdin's lamp yesterday. Thus, Draper would argue, I have not shown that the fact that something is bad according to OVT entails that it is anything we should care about; so I have not shown that we should worry at all about the evil of death. Perhaps death, though worse for us than continued life, is nevertheless "nothing to us," as Epicurus says.[51]

My intention here is not to draw any conclusions about how people should feel about death. I take my primary burden to be the burden of showing that a famous line of argument does not establish that death is not bad for the one who dies. According to that line of argument, death cannot even be said to be worse for the one who dies than continued life would be. I have tried to show that if we accept OVT and reject Presentism, we can successfully avoid that argument. If I have shown this, I would not have thereby shown that death merits any particular emotional reaction. Further argument is needed to establish that point; Draper is right about that.

I happen to think that despite the great evil of death, it might often be appropriate for me not to worry very much about my own death, except insofar as the worrying causes me to take precautions to prevent my premature death. The reason is not original: worrying too much about death, like worrying too much about anything, makes me unhappy. About this I at least partially agree with Epicurus, though I have attempted to show that Epicurus' reasoning (or a common interpretation of his reasoning) for this conclusion is flawed. I think that a more interesting question than "how should I feel about my eventual death?" is "what should I do about my (or others') eventual death(s)?" When we focus on this question, we notice an important difference between the case of death and the case of Aladdin's lamp: I can often do things to make it more likely that my death will occur later, but I can't do anything to make it more likely that I will find Aladdin's lamp. In this respect, the evil of death is more similar to evils that clearly count as "genuine evils," such as the evils of injury and sickness, than to evils that don't, such as the evil of not finding the lamp. The question of what I ought to do about my eventual death is a question to be answered by a theory in the normative ethics of

behavior; but the axiological and metaphysical considerations presented here can provide some help in answering it, by establishing that it is not incoherent to suppose that death is worthy of doing something about.

[1] This paper was presented at Virginia Tech, North Carolina State, Agnes Scott College, Idaho State, and the Second Annual Bellingham Summer Philosophy Conference; I thank all those present, including (but not limited to) John Carroll, Greg Fitch, Bill FitzPatrick, Valerie Hardcastle, Frances Howard-Snyder, Ned Markosian, Michael Rea, Russell Wahl, and especially my commentator in Bellingham, Elizabeth Harman, for their helpful comments. I thank Neil Feit, Fred Feldman and three anonymous referees for their helpful comments on earlier drafts. This paper grew out of ideas generated from discussion of Feit's forthcoming manuscript.

[2] When I say that death is "bad for" the one who dies, I mean that death is worse for the one who dies than continued life would have been (it negatively affects the well-being of the one who dies; it is "comparatively bad"). When death has this feature, it might also be worthy of a negative emotional response, and therefore be what Kai Draper calls a misfortune or a "genuine evil" (Draper 1999, 390). My main concern is with establishing the coherence of the claim that death is typically worse than continued life, but I will discuss the appropriate attitude to take toward death at the end of the paper.

[3] My primary concern here is not Epicurus exegesis; I am concerned about the argument. No doubt there are other Epicurean arguments in the vicinity of the quoted passage.

[4] See, for example, Mary Mothersill's comments in her recent APA presidential address (Mothersill 1999, 20).

[5] The view I defend is similar to one suggested by William Grey: "The temporal location of the harm of Ramsey's untimely death, I suggest, is the time when Ramsey might have lived" (Grey 1999, 364).

[6] Harry Silverstein also seems to hold a version of this view (see Silverstein 1980, 422, where he attempts to clarify Nagel's view). Silverstein holds that there can be atemporal facts but not atemporal events; so the fact that A died – an atemporal fact – is bad atemporally for A.

[7] It is interesting that Nagel's example involves a case where someone apparently goes out of existence. In virtue of this fact, we may wonder just what is supposed to be gained by considering this example. Nagel wants us to take what we learn from this case and apply it to the case of death; but given Nagel's assumption that the subject goes out of existence due to the injury, the case is too similar to a case of death to provide any independent support for his view.

[8] Nagel 1970, 77, his emphasis. It is probably a mistake for Nagel to think that what is relevant is how the person could be now; there are an infinite number of ways the person could have been, some good and some bad. What is really relevant is how the person would be now (McMahan 1998, 41).

[9] Nagel 1970, 78. Grey (1999) points out that Nagel does not explicitly say that the evil of death is timeless; he merely says it "cannot be so easily located" in time. Perhaps he did not mean to say that the evil of death cannot be located at all, but merely that it cannot be easily or precisely located. If that is Nagel's view, I do not disagree with it so far as it goes; but it does not go far enough to answer the question at hand.

[10] See Feit forthcoming, 4. (Page references are to the manuscript version of Feit's forthcoming paper.)

[11] I thank Fred Feldman for this example.

[12] This supposition is probably the most controversial of the assumptions involved in the example, as it entails that there is not always a unique closest possible world to a given world. Avoiding the counterexample involves denying the assumption, which would be equally controversial.

[13] I thank an anonymous referee for these examples.

[14] This argument is similar to one given by Feit (forthcoming, 2).

[15] For an extended discussion of the termination thesis, see Feldman 2000b.

[16] Feldman is such a materialist (Feldman 1991, 212), but this plays no role in his own response to Epicurus.

[17] It might be argued that since death does not cause any intrinsically bad sensations, death cannot be extrinsically bad. But this line of reasoning depends on a common misunderstanding about extrinsic value. Many things are extrinsically bad without causing anything intrinsically bad to occur; they are bad because they prevent intrinsically good things from occurring. Likewise, many things are extrinsically good without causing anything intrinsically good to occur; vaccinations are one kind of example. For more on this issue, see Bradley 1998.

[18] I am granting this assumption for now, even though it may be false, since I think it holds true when talking about evils (see Section Two).

[19] Of course, the event-type might continue to exist. However, I presume that if we have an ontology of events, particular event tokens are often good or bad for us. (I leave

open the possibility that we might also want to say that event types are good or bad for us.)

[20] See Markosian (forthcoming).

[21] Kai Draper suggests a response to the Epicurean that may amount to the same thing: "I suspect that it makes sense to speak of the dead as occupying a level of well-being because it is possible to refer to the living person who was, and to assign this past-existent a level of well-being in the present based on the no doubt limited extent to which he is now being benefited or harmed" (Draper 1999, 404; his emphasis).

[22] For more on "closeness" of worlds see Lewis 1986. I make no assumptions about the nature of possible worlds.

[23] Feldman has since clarified his view (personal correspondence). Feldman's current view is that death is bad for the one who dies, not at all specific times within a world (as seems to be his view in his 1991), but from the perspective of eternity – a point of view that stands outside all possible worlds. Thus Feldman can deny that someone's death is bad at a specific time within a world before that person exists. This makes his view more plausible, but makes it essentially similar to Nagel's view. I reject Feldman's clarified view for the same reason I reject Nagel's: namely, I think we may also want to make judgments about the badness of death, and other evils, at particular times within a world.

[24] Julian Lamont and Harry Silverstein give similar criticisms of Feldman's view (Lamont 1998, 199-200; Silverstein 2000, 120-121).

[25] I assume that injuries and sicknesses are merely extrinsically bad for people, in virtue of the fact that they cause people pain and – at least temporarily – prevent them from doing the things they would otherwise like to be doing. Naturally, the analogy with death

holds only insofar as death prevents people from doing and experiencing things, since being dead does not cause pain to the one who dies.

[26] For more on basic intrinsic value, see Harman 2000, 110-111, and Feldman 2000a.

[27] Such summative principles are not uncontroversial; objections to such principles have been around for some time, and have appeared recently in Velleman 1991. Those who object to summative principles may substitute their favorite replacement principle.

[28] For arguments against the stronger principle, see Hurka 1998 and Kagan 1998; for a defense of the stronger principle, see Bradley 2002.

[29] Grey employs a similar strategy for dealing with such evils: "To say that there is no precise or locatable time at which harms occur is not to say there is no time at which they occur" (Grey 1999, 364; his emphasis).

[30] See Lewis 1979 for discussion of backtrackers.

[31] See Feldman 1986, 26, and Bradley 1998, 118-119.

[32] That is, unless we allow "backtracking" counterfactuals. See the discussion of backtrackers in Section Three.

[33] Throughout the paper I have been pretending that we can represent the intrinsic values of things using precise cardinal numbers. Some will find it unrealistic to suppose that this can be done; it may seem absurd, for example, to suppose that the intrinsic value of one experience is precisely 2.37 times greater than the intrinsic value of another experience. Whether this is realistic or not, I think it is still helpful. For a defense of the use of numbers in value theory, see Hurka 1993, 86-88.

[34] Strictly speaking, OBAT does not entail that p is bad for John from t1 to t2. This is because there will be some moment t between t2 and t3 such that the value that has

accrued to John as of t will be identical in the two worlds. I think this only serves to make things worse for OBAT, as in this case it seems clear that p is bad for John from t1 to t2. If we use OBAT to determine whether p is bad for John from t1 to t2, the answer we get depends on what is going on at other times – whether or not p is bad for John overall.

[35] One objection that I do not have room to consider here is an objection that McMahan might raise: that OVT does not provide any way to discount the badness of a death based on the value of the life lived by the deceased. See McMahan 2002, 136-145.

[36] Elizabeth Harman gave the Aunt Alice example in her comments at the Bellingham Summer Philosophy Conference.

[37] Silverstein has recently reaffirmed this line of argument (Silverstein 2000, 119).

[38] Putting the claim in this way presupposes that propositions are states of affairs. I assume this for simplicity's sake. Even if propositions and states of affairs are distinct kinds of things, it won't make a difference to the problem noted here, because if no singular proposition about s exists at t, no state of affairs about s exists at t either; nor does any event involving s.

[39] For a discussion of these objections, see Markosian (forthcoming).

[40] An anonymous referee suggests that the objection could be restated to claim that although we may sometimes be able to calculate the value of the second term of the comparison, we cannot always do so. The idea is that dead people are not always the subjects of states of affairs after their deaths. But according to the conception of states of affairs that I have in mind, dead people are always the subjects of such states. The conception of states I have in mind is roughly the one used by Roderick Chisholm (1976,

114). I identify states with propositions; they are expressed by 'that'-clauses; they are bearers of truth and falsity (and, on my view, intrinsic goodness and badness); they are fine-grained entities (that JFK died and that JFK was assassinated by Oswald are distinct states of affairs); perhaps they are "abstract," in some sense. At every time after JFK's death, the proposition that JFK died is true; so JFK is the subject of a true state of affairs at every time after his death.

Assigning s-relative intrinsic value to a time when s no longer exists (but existed in the past) might not be a problem. But what do we say about the s-relative intrinsic value of a time if s never exists? This will depend on whether there can be any singular propositions about a never-existing person. I will not attempt to answer this question here; but it does not seem problematic to say, concerning a never-existing person, that all times have an intrinsic value of zero for him.

[41] See Adams 1986 and Markosian (forthcoming) for further discussion of Presentism and singular propositions.

[42] See Feldman 2000a, 328. The distinction between intrinsic value and basic intrinsic value is important here; Feldman believes that possible worlds are among the things that are bearers of intrinsic value, but states of affairs are the only bearers of basic intrinsic value.

[43] I thank Ned Markosian for discussion of this point.

[44] See Temkin 1993, 258.

[45] See McMahan 1988, 39, and McMahan 2002, 104.

[46] Elizabeth Harman raised a similar objection at the BSPC.

[47] See Feldman 1991, 225, for one way to account for overdetermination; see McMahan 2002, 119-120, for a critique of Feldman's strategy. The view I will defend is similar in certain ways to Feldman's view.

[48] Feit discusses, but does not endorse or reject, a strategy like this (forthcoming, 15). He views it as an alternative to Feldman's view, but I think it is consistent with Feldman's view, and may help to defend Feldman's view against McMahan's charge that it is ad hoc (McMahan 2002, 119-120).

[49] This objection has been raised, in different forms, by Fred Feldman and an anonymous referee.

[50] I thank Fred Feldman for the example.

[51] See Feit (forthcoming) for a more thorough discussion of Draper's examples.

REFERENCES

Adams, Robert. (1986) "Time and Thisness," in Midwest Studies in Philosophy, Volume XI , ed. Peter A. French, Theodore E. Uehling, and Howard Wettstein, 315-29. Minneapolis: University of Minnesota Press.

Aristotle. Nicomachean Ethics. Many versions.

Bradley, Ben. (1998) "Extrinsic Value," Philosophical Studies 91: 109-126.

Bradley, Ben. (2002) "Is Intrinsic Value Conditional?" Philosophical Studies 107: 23-44.

Chisholm, Roderick. (1976) Person and Object. LaSalle, IL: Open Court Publishing Co.

Draper, Kai. (1999) "Disappointment, Sadness, and Death." The Philosophical Review 108(3): 387-414.

Epicurus. "Letter to Menoeceus." Many versions.

Feit, Neil. "The Time of Death's Misfortune." Nous, forthcoming.

Feldman, Fred. (1986) Doing the Best We Can: An Essay in Informal Deontic Logic. Dordrecht: D. Reidel Publishing Co.

Feldman, Fred. (1991) "Some Puzzles About the Evil of Death." The Philosophical Review 100(2): 205-227.

Feldman, Fred. (1992) Confrontations with the Reaper: A Philosophical Study of the Nature and Value of Death. New York: Oxford University Press.

Feldman, Fred. (2000a) "Basic Intrinsic Value." Philosophical Studies 99: 319-345.

Feldman, Fred. (2000b) "The Termination Thesis," in Midwest Studies in Philosophy, Volume XXIV , ed. Peter A. French and Howard Wettstein, 98-115. Malden, MA: Blackwell Publishers.

Grey, William. (1999) "Epicurus and the Harm of Death." Australasian Journal of Philosophy 77(3): 358-64.

Harman, Gilbert. (2000) Explaining Value. Oxford: Clarendon Press.

Hurka, Thomas. (1993) Perfectionism. New York: Oxford University Press.

Hurka, Thomas. (1998) ""Two Kinds of Organic Unity." The Journal of Ethics 2: 299-320.

Kagan, Shelly. (1998) "Rethinking Intrinsic Value." The Journal of Ethics 2: 277-297.

Lamont, Julian. (1998) "A Solution to the Puzzle of When Death Harms Its Victims." Australasian Journal of Philosophy 76(2): 198-212.

Lewis, David. (1979) "Counterfactual Dependence and Time's Arrow." Nous 13: 455-76.

Lewis, David. (1986) On the Plurality of Worlds. New York: Basil Blackwell.

Markosian, Ned. "A Defense of Presentism," in Oxford Studies in Metaphysics, vol. 1, ed. Dean Zimmerman (forthcoming).

McMahan, Jeff. (1988) "Death and the Value of Life." Ethics 99: 32-61.

McMahan, Jeff. (2002) The Ethics of Killing: Problems at the Margins of Life. New York: Oxford University Press.

Mothersill, Mary. (1999) "Old Age." Proceedings and Addresses of the APA, 73(2): 9-23.

Nagel, Thomas. (1970) "Death." Nous 4: 73-80.

Plantinga, Alvin. (1995) "Haecceity," in A Companion to Metaphysics, ed. Jaegwon Kim and Ernest Sosa, 199. Malden, MA: Blackwell Publishers.

Silverstein, Harry. (1980) "The Evil of Death." The Journal of Philosophy 77: 401-424.

Silverstein, Harry. (2000) "The Evil of Death Revisited," in Midwest Studies in Philosophy, Volume XXIV , ed. Peter A. French and Howard Wettstein, 116-134. Malden, MA: Blackwell Publishers.

Temkin, Larry. (1993) Inequality. New York: Oxford University Press.

Velleman, David. (1991) "Well-being and Time." Pacific Philosophical Quarterly 72: 48-77.