

CHAPTER 9

.....

THE BADNESS OF DEATH AND THE GOODNESS OF LIFE

.....

JOHN BROOME

WHAT harm does death do you? To put the question differently: when you die, what do you lose by dying? To put it differently again: when you do not die, what do you gain by continuing to live? The question of what harm death does you is the same as the question of what good is done you by living. It is the question of the goodness of your life.

Two extreme answers can be given. One is “everything”; we might think that, for you, your life is everything, and by dying you lose everything. Another is “nothing”; we might think that you lose nothing by dying. I shall start by rejecting these extreme answers. Then I shall go on to the moderate, quantitative answer that I favor.

1. DO YOU LOSE NOTHING BY DYING?

.....

I shall take the “nothing” answer first. Epicurus may be read as giving this answer. He says:

Become accustomed to the belief that death is nothing to us. For all good and evil consists in sensation, but death is deprivation of sensation. . . . 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, 1926, pp. 30–31)

Epicurus seems to be saying that death does you no harm. If this is right, it follows that continuing to live does you no good.

Most of us find the “nothing” answer implausible because we take it for granted that dying would be a terrible thing to happen to us. Epicurus himself may not mean to give this answer. When he says “death is nothing to us,” he may not mean that death does us no harm—I shall come to that. Nevertheless, his argument does supply materials that can be used to construct a case for the “nothing” answer. I shall lay out this case and try to make it persuasive despite its initial implausibility. But in the end I shall argue that it fails.

The beginning of the case is to recognize that the goodness of life has two dimensions: its quality and its quantity. It is quite easy to slip into thinking that the quantity of life does not matter at all; only its quality matters. This is exactly what most of us think about the goodness of life in another context. One way of adding to the quantity of life in the world is by having more babies; that way, more life is lived in total. But most of us do not favor increasing the quantity of life this way. We favor increasing the quality of life of the people who live, but we do not favor increasing the number of people who live. When the Chinese government instituted its one-child policy, its aim was to increase the quality of life of the Chinese. The policy also decreases the quantity of Chinese life: there are fewer Chinese now than there would have been without the policy. But the government did not think of this reduction in quantity as a bad thing, to be set against the increase in quality. Most of us would have agreed.

Moreover, this attitude we commonly have to the number of people can be supported by an argument. Suppose a couple are thinking of having a child, but eventually decide not to. As a result of their decision there is less life in the world than there would have been had they decided differently. Is this reduction in quantity a bad thing? Well, no one is harmed by it. No one is harmed by not being brought into existence. It is not as though there is some child who suffers the misfortune of not existing. There is simply no child, so no one is harmed. Consequently, we might plausibly think no harm is done. We might conclude it cannot be a bad thing to reduce the quantity of life in this way.

This argument needs to be qualified. Perhaps some people will be worse off as a result of the child’s nonexistence. Perhaps her potential parents will come to regret having no child, and perhaps the child would have grown up to make a great contribution to civilization. So perhaps some people will be harmed by her nonexistence. But these are indirect effects, and to keep the argument sharp, let us assume them away. Let us assume there are no indirect effects of this sort, even though in practice there will almost certainly be some. Under this assumption, the argument has some force.

Now back to our context, which is extending life rather than creating life. Bringing more people into the world is one way of increasing the quantity of life. Another is to extend the lives of people who are already in the world. Epicurus

shows us that we can take the same attitude to quantity in this context too, and there is a parallel argument for thinking that quantity has no value. We can ask a parallel question about the quantity of a single person's life. Previously we asked who is harmed by not being created; now let us ask at what time a person is harmed by not continuing to live. Suppose you might have lived longer, but you actually die now. Is that a bad thing for you? Well, there is no time when you are harmed by your early death. As Epicurus says, you are not harmed at any time before your death, since so long as you exist "death is not with [you]." And you are not harmed at any time after your death, since at no time after your death do you exist. Since there is no time when you are harmed by your death, we might conclude your death is not a bad thing for you. In the same way, the previous argument concluded that, since there is no one who is harmed when a couple declines to have a child, they do nothing bad in acting as they do.

In response, you may say there is indeed a time when death harms you: the time when you die. In saying this, you could be making either of two points. The first is that the process of dying is often dreadful. That is obviously true, and it does mean that your death harms you in one way. But it is not relevant to the question I am asking. I am asking what is the benefit to you of continuing to live. Conversely, what harm would be done you by not continuing to live? What harm would be done you by having your life cut short? I sometimes express this question in the form: what harm does your death do you? This is a graphic but not entirely accurate way of putting the question of what harm would be done you by having your life cut short. The terribleness of the process of dying is not a part of the answer to this question. Cutting your life short does not necessarily harm you in this way, because your dying may be dreadful whether it occurs at the end of a long life or a short one. So we can set aside this aspect of the badness of death.

The second point you might be making is this. If death harms you, it is obvious when the harm is done. It is done at the time of your death, since your death does the harm. This is true too, but it is also not relevant to the question I am asking. We must distinguish the time when a harm is caused from the time when it is suffered. If I drop a banana skin on the road, and you later slip on it and hurt yourself, we may say your harm is caused when I drop the banana skin. But it is suffered when you fall. Epicurus is interested in the time when the harm of death is suffered, not when it is caused. His conclusion is that it is not suffered at any time. If there were any harm, it would be caused at the time you die, but that is another matter. We can set aside this point too.

Once those two points are set aside, I think we should agree that there is no time when death harms you. That is a truth we should learn from Epicurus. Epicurus apparently draws the conclusion that, because there is no time when death harms you, it does not harm you at all. But to reach that conclusion, we have to make the further assumption that an event cannot harm you unless it harms you at some time. Is that a good assumption?

Once again, Epicurus supplies us with material that at first seems to support it. He says that "all good and evil consists in sensation." He means that the only sort of

good that can come to us is a good sensation, and the only sort of bad is a bad sensation. This is a version of what is nowadays called “hedonism.” It is highly contentious, and one way of responding to Epicurus is to deny it.¹ But denying hedonism is also contentious, and for my purposes I do not need to deny it. Instead, I shall show that, even if we grant Epicurus’s hedonism, it does not truly support the claim that you cannot be harmed unless you are harmed at some particular time.

So let us assume like Epicurus that all good and evil consists in sensation. Since all sensations occur at particular times, we can quickly conclude that all goods and evils occur at particular times. So the goodness or badness of your life is made up of good and bad things, all of which occur at particular times in your life. This is a consequence of hedonism.

But the notions of *benefit* and *harm* are different from the notions of *good* and *bad*, and just because all goods and bads occur at particular times, it does not follow that all benefits and harms do. *Benefit* and *harm* are comparative notions. Normally, if something benefits you, it makes your life better than it would have been, and if something harms you it makes your life worse than it would have been. “Better” and “worse” are the comparatives of “good” and “bad,” respectively. A comparison is between two things. To determine whether some event benefits or harms you, we have to compare the goodness of your life as it is, given the event, with the goodness it would otherwise have had. The comparison is between your whole life as it is and your whole life as it would have been. We do not have to make the comparison time by time, comparing each particular time in one life with the same time in the other life. So even if the goodness of your life is made up of good and bad things that all occur at particular times, there is no need for the comparison between lives to be made up of benefits and harms that can all be tied down to particular times.

Take an analogy. Suppose the text of a book is shortened before it is published: the last chapter is cut out. The book is shortened by six thousand words, but all the earlier chapters are left intact. Then six thousand words are cut from the book; yet no words are cut from any page in the book. This is so even though every word in the book appears on a particular page. Moreover, had the book been published in the longer, uncut version, every word in the longer book would have appeared on a particular page. The number of words cut from the book is determined by comparing the whole book as it is, with the whole book as it would have been had it not been shortened. It is not determined by comparing any particular page with that same page as it would have been.

Similarly, death may harm you by shortening your life, even though there is no time when it harms you. To determine whether it harms you, we compare the goodness of the shorter life you have, taken as a whole, with the goodness of the longer life you would have had, taken as a whole. If we believe Epicurus’s hedonism, the goodness of the shorter life is made up of the good and bad sensations that occur within it. The goodness of the longer life includes all those sensations too, and it also includes all the good and bad sensations you would have had in later life had you not died. If your life is going well, presumably these extra sensations

would have been predominantly good ones. So the longer life would have been better than the shorter one. You are therefore harmed by the shortening of your life. But there is no time when you suffer this harm, just as, when the book is shortened, no page in the book loses any words. Epicurus's hedonism actually implies that death normally harms you. Epicurus thinks it implies the opposite, but he is making a mistake.

Go back briefly to the analogous argument about the world's population. It fails for the same reason. The question is whether a couple's decision not to have a child is a bad thing. To answer, we must compare the goodness of the world without the child with the goodness it would have had if the child had existed. The world might be better without the child, or worse, or equally good. In particular, it might be worse, even though there is no one for whom it is worse. Again: a book can be shorter than it would have been, even though no page has any fewer words.

The argument I took from Epicurus fails. Epicurus is right that there is no time when death harms you. But even granted hedonism, it does not follow that death does not harm you. It may harm you, even though it harms you at no time.

2. SHOULD WE MIND ABOUT DYING?

I took the argument from Epicurus, but Epicurus may not mean to argue that death does not harm you. By "death is nothing to us," he may mean simply that you should not mind about dying. It is possible that you should not mind about dying even though your dying will harm you. Perhaps that is what Epicurus thinks.

How could it be so? If dying will harm you, surely you should mind about it. Not necessarily. It depends on what you should care about. Dying will harm you, but possibly you should not care about what happens to you, yourself. You are a person, with a life that extends from when you come into existence to when you go out of existence. Caring about what happens to you involves caring about the whole of that life. But why should you care about that? For instance, as an alternative, why should you not care just about what happens to you in the present? What you care about may change from time to time. Why should you not, at each particular time, care about just what happens to you at that particular time?

This needs to be put carefully. Probably you anyway care about what happens to other people besides yourself. But you probably care in a different way about what happens to you yourself. Call this sort of care "self-care." The suggestion is that you should attach your self-care, not to what happens to the person you are, with the whole of your life, but just to what happens to you in the present.

Wittgenstein uses the expression "living in the present," and I think this is what he means by it. He points out: "For life in the present there is no death. Death is not an event in life" (1961, p. 75). He is saying that, so long as you care only about what happens to you in the present, rather than about yourself as a whole, you will

never encounter death among the things you care about. Your death does not occur during your life, so for you it is never in the present.

Possibly Epicurus is making a similar point. Since there is no time when death harms you, death does not harm you in the present, whatever time happens to be the present. So if you should care only about what happens to you in the present, you will never have any reason to mind about dying.

I am here not concerned with the correct interpretation of Epicurus. I am interested in how good it is for you to continue living. This is a question about the good of you, the person you are, who has a whole life. It is not about what you should care about at any particular time. The question is whether dying—ceasing to live—harms you. I asked whether we could find in Epicurus’s remarks any reason for thinking it does not. His remarks provide the materials for an argument, but in the end the argument fails.

3. DO YOU LOSE EVERYTHING BY DYING?

Now I come to the opposite extreme answer to the question “what do you lose by dying?”: the answer that you lose everything.

Here is an argument that supports this answer: after you die you will not have anything, so in dying you will lose everything. But this argument is invalid. Its premise is true: after you die, indeed you will not have anything. But it is true only in a peculiar way, and in this peculiar way it does not support the conclusion that in dying you will lose everything. I shall explain why not.

The sentence “you will not have anything” can be true in two different ways. One is when you exist and do not have anything. In this case, the negation contained in the sentence is often called “internal,” because it negates the sentence’s predicate. The sentence may be parsed “you will (not have anything).” The other way is when you do not exist. The negation is then “external” because it negates the sentence as a whole. The sentence can only be understood as meaning “it is not the case that you will have something.” (A little point of English that may be confusing: in our context, “anything” replaces “something” under negation.)

The premise of the argument—that after you die you will not have anything—is true in this second way and not the first. This is what I called the “peculiar” way of being true. In the same peculiar way it is true that Pegasus—a winged horse—has no wings, because he does not exist. In the same way, too, it is true that Nelson now has no left arm; it is true because, being dead, Nelson now does not exist.

An external negation does not support the claim that something is lost. Think some more about Nelson’s arms. Before he attacked Santa Cruz de Tenerife, Nelson had a right arm. Afterward he did not. The negation here is internal: after the attack, Nelson did (not have a right arm). Because the negation is internal, we may correctly draw the conclusion that Nelson lost his right arm in the attack. His

loss consisted in the difference between two states of Nelson: the previous state in which he had a right arm and the subsequent state in which he did not. So an internal negation makes loss possible.

Compare what happened at the Battle of Trafalgar, when Nelson died. Before the battle, Nelson had a left arm. Afterward he did not. But in this case the negation is external. After the battle, it was not the case that Nelson had a left arm, but it is not correct to say that he did (not have a left arm), since he did not exist. Because the negation is external, we cannot correctly conclude that Nelson lost his left arm at Trafalgar. There is no comparison to be made between a previous state of Nelson in which he had a left arm and a subsequent state in which he did not, since subsequently there was no Nelson.

There were indeed two states of Nelson's body: a state previous to the battle and a state subsequent to it. This means the body could have lost its left arm. As it happens, it did not; the arm remained attached to the body. In any case, whatever happened to his body, Nelson himself did not lose an arm at Trafalgar.

In the same way, there is no comparison to be made between your state before you die and your state after you die. From the premise that after you die you will not have anything, with its external negation, we cannot correctly conclude that you lose everything by dying. The argument I have been discussing purports to make a temporal comparison between what you have before your death and what you have after it. The argument fails because there is no real temporal comparison to be made.

To determine what you lose as a result of a particular event, we do have to make a comparison. But there are two sorts of comparison we might make, and two corresponding sorts of loss. One comparison is temporal. We compare what you have after the event with what you had before it. If you have less afterward, you have suffered a temporal loss. We can only make a temporal comparison like this if you exist both before and after the event. When the event is your death, you do not exist afterward. Therefore, your death cannot cause you temporal loss.

I admit we say that, at your death, you lose your life, and this loss is plainly meant to be temporal. But this is a unique idiom; we do not commonly say that, at your death, you lose other things besides your life. When a husband dies, we say his wife loses her husband, but we do not say the husband loses his wife. We recognize that would be false. We should also recognize that, although idiomatic, it is strictly false to say that the husband loses his life.

The other sort of comparison we may make is atemporal, and yields an atemporal sort of loss. This sort of comparison does not require you to exist both before and after the event that causes the loss. When we ask what you lose by dying, we do not have to answer the question by comparing what you have after your death with what you have before it. We can instead compare what you have, given that you die at a particular time, with what you would have had if you had not died then. When we think this way, "what you have" does not refer to what you have at a particular time, but to what you have atemporally, taking your whole life together. What you lose by dying, understood this way, is not everything. It is just a part of the longer

life you would have led, had you not died when you do. What you lose by dying is not your life, but only the rest of your life.

Our question is “what do you lose by dying?” The incorrect answer “everything” encourages the idea that living is infinitely good for you. But no one should believe that. No one’s life is infinitely good. How could it be? Dying shortens your life by only a finite length of time. Our human lives are only finite in length, and during them we can experience and achieve only a finite number of things.

The only way to answer the question correctly is to understand it atemporally. What you lose by dying is the finite difference between a longer life and a shorter one. This answer lies between the extremes of “everything” and “nothing,” or between “infinity” and “zero.” It is “something.”

4. HOW MUCH DO YOU LOSE BY DYING? A PRACTICAL QUESTION

But “something” is not a good enough answer; we need to know how much. Excluding the two extremes puts us in the domain of quantities. I have said that what you lose by dying is the rest of your life. I now turn to assessing what you lose in quantitative terms. More exactly, I turn to assessing the value to you of what you lose. How bad for you is your loss? I said that what you lose is the rest of your life; how bad is that? Put differently, how good for you is the rest of your life?

We should not expect all deaths to be equally bad; some people lose more by their deaths than others do. Indeed, some unfortunate people benefit from their deaths. I shall continue to write of the badness of death, but I mean to allow for the possibility that some deaths are good, which is to say that they have negative badness.

The badness for you of your death is the difference between the goodness of the longer life you would have led had you continued living and the goodness of the life you actually do lead. In general, the badness of death is the difference between the goodness of a longer life and the goodness of a shorter one. (In some cases this difference may be negative.) So, to assess the badness of a particular death, we need first to work out what life is led by the person who dies and what life she would have led had she continued living. Then, second, we need to judge the goodness of those two lives.

The first task is partly empirical and partly a matter of evaluating the counterfactual notion of “the life the person would have led.” Some writers on death give space to evaluating this counterfactual,² but I shall not. One reason is that it is not particularly a problem for the philosophy of death; it is a problem of counterfactuals in general. Another reason is that the counterfactual is not important in practice, as I shall soon explain.

In the rest of this chapter, I shall therefore concentrate on the second task, of judging the goodness of lives. The task of judging the badness of death transmutes into this task of judging the goodness of lives. We have to compare the goodness of a longer life with the goodness of a shorter one.

I have described this difference as the goodness of the rest of your life. This is correct in one sense: it is the amount of good that the rest of your life would bring you if you lived it. But it is not necessarily the amount of good that you would enjoy during the rest of your life. The goodness of your life may be determined holistically, in a way that involves interactions among different parts of it. So the rest of your life may benefit you in some way that is not simply by being good in itself. It may add good to earlier parts of your life, or it may contribute in other ways to the goodness of your life as a whole. There will be examples in section 5.

Why do we need to do this quantitative work? Because in practice important decisions hang on it. Life and death decisions are constantly being made—I mean decisions that affect the lengths of people's lives. Some are on a small, individual scale; others on the scale of the whole world. On a small scale, all of us regularly make decisions that shorten or lengthen our lives. Statistically, each doughnut shortens your life. Is it worth it? That is probably something you do not want to think about. But in other cases, you will want to make the calculation. If you have a terminal illness, you will need to decide at what point to give up aggressive treatment aimed at prolonging your life, and accept only palliative care till you die. You may think carefully about that. Your decision may depend on your judgment of the goodness of extending your life—for instance, on whether you have a work of art to finish and on whether you expect to lose your capacities.

You will be weighing the quantity of your life against its quality. You may need to do this for yourself explicitly only in rare and tragic circumstances. But when the decision is for other people, you will need to be more careful. You can be cavalier about your own doughnuts, but not about other people's lives.

Governments in particular make decisions that lengthen or shorten many people's lives, so they need to judge the goodness of those lives. Governments often have to weigh some people's lives against others. They also often have to weigh the quantity of lives against the quality of lives. Take the provision of health care. Some treatments (such as hip replacements) improve the quality of people's lives without extending them. Some (such as heart replacements) extend lives. Many governments explicitly or implicitly set priorities among different sorts of treatment. To do so properly, they must weigh the quality of life against the quantity of life. They need to assess the goodness of people's lives.

On a much larger scale, we must decide what to do about global warming. One of the greatest harms that global warming will do is to kill huge numbers of people. It will kill them in floods and famines and in heat waves; it will kill them by extending the range of tropical diseases; and it will kill them in marginal areas of the world by making them poorer—poverty is a killer. By reducing our emissions of greenhouse gases, we can reduce the number of people who will be killed. But to do that we shall have to sacrifice some of the quality of our own lives. What sacrifices

should we make? What reduction in the quality of our lives in the present is worthwhile for the sake of extending the quantity of people's lives in the future? Again, we need to assess the goodness of people's lives, and weigh quality against quantity.

So the practical need for judging the goodness of lives is as an input into decision making. I am not suggesting that goodness is the only input. Fairness in the distribution of goods also matters, particularly in making public decisions that influence which people die and which survive. But goodness is one input.

Decision making is a matter of choosing among a number of options. Each option will lead to a particular state of the world. But we never know exactly what state will result from the option we choose; the results are always uncertain to some degree. To judge the goodness of each option we therefore have to take account of its uncertainty. Expected utility theory tells us the correct way of doing so. In principle, we must assess the goodness of each state of the world that might result from the option, and calculate a weighted average of those goodnesses, in the way expected utility theory tells us to. The details of the method do not matter here.

For example, suppose you are deciding whether to reject or accept aggressive treatment for your terminal disease. If you reject it, you will die after some time, which is uncertain, and your life till that time will have some quality, which is also uncertain. If you accept aggressive treatment, you will die after some time that is probably longer, and your life till that time will have some quality that is probably less good. You should assess the goodness of each possible result of each option, and compare the weighted average goodness of the possible results of one option with the weighted average goodness of the possible results of the other.

None of this requires you to evaluate the counterfactual notion of "the life you would have led" had you decided differently. For practical purposes, you do not need to evaluate this counterfactual. You do not need to assess the badness of your death in a way that involves it. So it is not needed for practical purposes. That is the main reason why I do not give space to it.

5. THEORIES OF THE GOODNESS OF LIVES

The goodness of lives determines how bad it is to die; the question of how bad death is transmutes into the question of how good life is. So how good is a life? I am sorry to say this is too difficult a question for me to answer here. It is one of the topics of my book *Weighing Lives*, and even there I was able to offer only a "default theory" of the goodness of lives (2004, chs. 15–17). I meant a theory that it is reasonable to hold so long as there are no good arguments against it. A very large range of theories are available and, in the present state of discussion, choosing among them is generally more a matter of intuition than argument. Here, I shall survey part of the range, and provide a partial taxonomy of it. I shall give just a few examples of the theories that are available. I shall not try to adjudicate between them.

Two clarifications are needed at the start. One is about the tightness of the scale of goodness we should aim at. The simplest aim would be just to put lives in the order of their goodness—to determine which lives are better than which. That would be enough to answer the most basic question about the value of a person's death: is it good or bad? It is good if the shorter life the person leads is better than the longer one she would have led, and bad if the shorter life is worse.

But for practical decision making we need more than this. Because the results of a decision are always uncertain, we need to apply expected utility theory. That requires us to have a cardinal scale for goodness, including the goodness of lives. So in valuing lives, we need to aim for a cardinal scale.

I do not say we can expect to get one. Indeed, it seems unreasonable to expect even a determinate ordering of lives by their goodness. It seems likely that there is sometimes no determinate answer to the question of whether it would be a good or a bad thing to prolong a particular life, at some cost in its quality. We should surely expect a lot of incommensurability in assessing the goodness of lives. But I shall have to leave aside the question of how this incommensurability should be dealt with. The theories I shall mention ignore it. But this does not vitiate their value, because there may well be ways of extending them to take incommensurability into account.³

A second clarification is that I am dealing with the goodness of lives for the person who lives them; I call this their “personal goodness.” The question is how good it is for a person to live a particular life. We could also ask how good it is simpliciter, rather than for the person, that a person lives a particular life. This is a question of “general goodness,” as I call it. It is not the topic of this chapter. As it happens, I think personal goodness generates general goodness:⁴ if one life *A* is better for the person who lives it than another *B*, it is better simpliciter that this person lives life *A* than that she lives life *B*. But that assumption plays no part in this chapter.

Theories of the goodness of lives range from very particular ones to ones that have a lot of formal structure. One example of a very particular theory is that the only thing that makes a life good for a person is the excellence she achieves, so the goodness of a life is the amount of excellence it contains. I shall concentrate on more structured theories.

5.1 Distributed Theories

More structured theories can be developed if we make what I shall call the assumption of “distribution.” It is the assumption that the goodness of a life is distributed across times. Put another way, it is the assumption that the goodness of a life is made up of its goodness at all times. It is really a combination of two assumptions. First, that there is such a thing as the goodness of a life at particular times—I shall call this “temporal goodness” and contrast it with the “lifetime goodness” of a life as a whole. Second, that the goodness of the life supervenes on its goodness at times. In mathematical terms, it is a function of its goodnesses at times.

It will be helpful to express this assumption in symbols. For convenience, I shall imagine that times are discrete. The assumption of distribution is that lifetime goodness G is given by the formula:

$$G = G(g_1, g_2, g_3, \dots, g_n).$$

The indices $1 \dots n$ denote a sequence of times, and $g_1, g_2, g_3, \dots, g_n$ are the temporal goodnesses of the life at those times. I call $G()$ the “goodness function.” Let us call a theory of goodness “distributed” if it satisfies the assumption of distribution. Different distributed theories disagree about the form of the goodness function.

The assumption of distribution does not specify the scale on which temporal goodness is measured. Different views about the goodness function require different scales, as will appear.

You might well reject the assumption of distribution. You might think there are good features of a life that cannot be assigned to particular times. I have already mentioned the theory of goodness as excellence. You might hold this theory and also think that the excellence of a life cannot always be assigned to particular times. For another example, there is the view that the lifetime goodness of a life depends on its degree of internal coherence in some way that is not reflected in its sequence of temporal goodnesses. Perhaps a life is better if it is directed toward one particular broad aim, rather than toward an eclectic mixture of aims. It might not be possible to allocate this good feature of the life to any particular times within the life.

Still, the assumption of distribution is not very restrictive. It leaves room for very many different distributed theories. It even allows a life to have temporal goodness at times outside the boundaries of the life; it does not rule out posthumous or antenatal goods and bads. However, for the sake of convenience in what follows, I shall assume that all the temporal goodness of a life occurs within its temporal boundaries. So in the formula above, I shall assume that the times indexed by $1 \dots n$ constitute the sequence of successive times within the life. It would take only minor adjustments to remove this assumption.

The assumption of distribution leaves it open how temporal goodness is determined. The nature of well-being is a large and hotly debated subject, and much of this debate is about temporal goodness. Does your good at a time consist in your experiences at the time, as Epicurus assumed, or in the satisfaction of the preferences you have at the time, or in something else? All of this is left open by the assumption of distribution, and it is left open by all the distributed theories I shall mention in this section. These are theories about how a person’s good at particular times comes together to determine her lifetime good. They are not theories about the nature of her good at particular times.

In particular, the assumption of distribution does not require temporal goodness at a time to be determined only by events that happen at that time. Indeed, it is implausible that it would be, unless events are construed very broadly. Suppose, say, that you work hard on a project, and the project is later successful. We might think that your later success adds to the goodness of the earlier times when you work hard on it; it might cause a change in the goodness of those earlier times. This would be

a sort of backwards causation of goodness, and it is perfectly compatible with the assumption of distribution. I mentioned in section 4 that, were your life threatened at some time, but you survive, the rest of your life might benefit you by adding value to earlier parts of your life. If that is so, it is another example of backward causation of goodness, and it is compatible with the assumption of distribution.

5.2 Additively Separable Theories

A vast range of forms are possible for the goodness function. The one that may first spring to mind is simple addition, which makes a life's lifetime goodness the total of the goodness of its times. Call this the "total theory":

$$G = g_1 + g_2 + g_3 + \dots + g_n.$$

It demands a tight scale for measuring temporal goodness: the scale must be cardinal. Furthermore, goodness at each time must be comparable with goodness at other times. Furthermore again, the scale must have a fixed zero, which means it is a ratio scale. This is because we need to compare the lifetime goodnesses of lives of different lengths. Take a life of some length, and imagine shortening it by removing the last time in it. Imagine the person dies one time earlier, that is to say. According to the total theory, the shorter life is worse than the longer one if and only if the goodness of that last time is above zero. So the level of the zero makes a difference to the ordering of lives.

The total theory belongs to a class of distributed theories that may be called "additively separable." The characteristic of additively separable theories is that they treat the goodness of a life as the sum of values, each of which is assigned to a particular time and is a function of the temporal goodness of that time. The value assigned to a time must be independent of the goodness of other times, and of the length n of the life. Put roughly, each time can be valued independently of other times. The general formula of an additively separable theory is:

$$G = v_1(g_1) + v_2(g_2) + v_3(g_3) + \dots + v_n(g_n).$$

I shall call $v_1()$, $v_2()$ and so on the "temporal value functions." The form of these functions is independent of n and of temporal goodness at other times.

In the total theory, the temporal value functions are the identity function. Other additively separable theories have other functions. The "weighted total theory" departs from the total theory only by giving different weights to goodnesses at different times:

$$G = a_1g_1 + a_2g_2 + a_3g_3 + \dots + a_ng_n.$$

Here, a_1 , a_2 are constants that specify the weights. If later weights are greater than earlier ones, later times in life count for more than earlier ones. One result is that a life that improves over time is better than one that deteriorates, if they both have the same total of temporal good, well-being. This is a consequence of the weighted total theory.⁵ By contrast, some authors—generally economist or public health analysts—"discount" the goodness of later times in a life (e.g., Murray 1994). This means they give later times less weight than earlier ones.

Delete

Another additively separable formula is:

$$G = v(g_1) + v(g_2) + v(g_3) + \dots + v(g_n).$$

Here all the temporal value functions are the same, $v()$. Take the case where $v()$ is an increasing, strictly concave function, which means its graph slopes upward but curves downward. Then we may call the theory “prioritarian.” It gives priority to improving bad times over improving good ones. This has the indirect effect of assigning more goodness to a life that has an even tenor than to one that has extreme highs and lows, if they both have the same total of temporal good. It gives indirect value to evenness, that is to say.

5.3 Constant-Length Additively Separable Theories

A different type of theory is the “average theory,” that the goodness of a life is the average of its temporal goodnesses. Its goodness function is:

$$G = g_1/n + g_2/n + g_3/n + \dots + g_n/n.$$

This theory does not require temporal good to be measured on a ratio scale; a cardinal scale is enough. It is one formulation of the view, mentioned in section 1, that only the quality of life matters, and not its quantity.

The average theory is not additively separable by the definition I gave. It does treat the goodness of a life as the sum of values, each of which is assigned to a particular time and is a function of the temporal goodness of that time. However, the form of this function depends on the length of the life n .

When lengthening or shortening the life is not in question, the average theory is equivalent to the total theory. Among lives that are all the same length, it orders them just as the total theory does. So the average theory is additively separable among lives with the same length.

But think about extending a life by one time. If the temporal goodness of this time is above the average of the existing times, then the life is improved by extending it. If it is below, the life is made worse by extending it. So the value of adding an extra time depends on the temporal goodness of other times. In this sense the value of this extra time is not independent of other times. That is why I do not count the average theory as truly additively separable. Instead, I say it is “constant-length additively separable.”

Another theory in the same class is:

$$G = g_1 + g_2 + b(g_2 - g_1) + g_3 + b(g_3 - g_2) + \dots + g_n + b(g_n - g_{n-1}).$$

This theory gives value to improvements in temporal good, as the weighted total theory can do, but it does so more directly. b is the weight assigned to improvements. This goodness function may be rewritten in the form

$$G = (1 - b)g_1 + g_2 + g_3 + \dots + (1 + b)g_n.$$

This makes it look superficially like an instance of the weighted total theory. But actually it is not additively separable because the weight given to any particular time

depends on whether or not it is the last time. However, this theory is constant-length additively separable.

5.4 Nonadditively Separable Theories

Some theories are not additively separable at all. Some start out from the total theory, and modify it in one way or another, to take account of values it does not accommodate. One of these gives value directly to evenness:

$$G = g_1 + g_2 + g_3 + \dots + g_n - cI(g_1, g_2, g_3, \dots, g_n).$$

$I()$ is some measure of unevenness in the life's temporal goodnesses; it is a measure of inequality among the temporal goodnesses in the life. Various measures could be used: the variance, the Gini coefficient, and so on, and c is a parameter that assigns a weight to evenness. This formula values evenness more directly than the prioritarian formula does.

Other theories of lifetime goodness are much more remote from the total theory, but nevertheless satisfy the assumption of distribution. One is the theory that the goodness of a life is given only by how good it is at its end:

$$G = g_n$$

Another is the theory that the goodness of a life is given by the best time in it:

$$G = \max\{g_1, g_2, g_3, \dots, g_n\}$$

Like the average theory, these theories are alternative expressions of the idea that only quality of life matters, and not quantity. They may be combined into the "peak and end rule."⁶ Either of them is able to order lives by their goodness so long as temporal goodnesses are ordered and comparable between different times. To order lives, neither requires a cardinal scale of temporal goodness. However, if the overall goodness of lives is to be on a cardinal scale, temporal goodnesses must be on a cardinal scale too.

I hope I have given enough examples now to illustrate the range of choice available among theories of lifetime goodness.

6. CONCLUSION

When you die, what you lose is neither nothing nor everything. It is the rest of your life. The badness of this loss is, seen differently, the goodness of rest of your life. More accurately, it is the difference between the goodness of the longer life you would have led, had you survived, and the shorter life you do lead. So the question of how bad is death transmutes into the question of how good is life.

I have not tried to answer this latter question, but I have outlined and classified some of the answers that are available.

NOTES

1. This is the way adopted by Thomas Nagel, 1970.
2. For instance, Ben Bradley, 2009, pp. 47–60.
3. A valuable recent discussion of how this might be done is Wlodek Rabinowicz, 2009.
4. This view is formalized in something I call “the principle of personal good.” See my 2004, p. 120.
5. The view that improvement is good is championed by David Velleman, 1991.
6. Kahneman, 1999. But note that Kahneman does not favor the peak and end rule as a formula for the value of a life.

REFERENCES

- Bradley, Ben. 2009. *Well-Being and Death*. Oxford: Oxford University Press.
- Broome, John. 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Epicurus. 1926. “Letter to Menoeceus.” In *Epicurus: The Extant Remains*, Cyril Bailey, trans., pp. 83–93. Oxford: Oxford University Press.
- Kahneman, Daniel. 1999. “Objective Happiness.” In *The Foundations of Hedonic Psychology*, Daniel Kahneman, Ed Diener, and Norbert Schwarz, eds., pp. 3–25. New York: Russell Sage Foundation.
- Murray, Christopher. 1994. “Quantifying the Burden of Disease: The Technical Basis for Disability-Adjusted Life Years.” *Bulletin of the World Health Organization* 72: 429–45.
- Nagel, Thomas. 1970. “Death.” *Noûs* 4: 73–80.
- Rabinowicz, Wlodek. 2009. “Incommensurability and Vagueness.” *Proceedings of the Aristotelian Society, Supplementary Volume* 83: 71–94.
- Velleman, David. 1991. “Well-being and Time.” *Pacific Philosophical Quarterly* 72: 48–77.
- Wittgenstein, Ludwig. 1961. *Notebooks 1914–16*. Oxford: Blackwell.