CHAPTER 13

Other Aspects of Death

According to the deprivation account, the central way in which death is bad for us—when it is bad for us—is that it deprives us of something good. Initially, of course, I presented this idea by saying that death deprives us of the good things *in* life, but now we have seen that some may want to modify this slightly, so as to take note of the possibility that life itself may be good as well. But regardless of the details, we can capture the basic idea by saying that the central badness of death lies in the fact (when it is a fact) that it deprives me of life worth having.

But although I've been at pains to say this is the central or *fundamental* bad thing about death, I think that one could make the case that this isn't the *only* bad thing about death—not even if we continue to focus on the way in which death is bad for the person who *dies*. There are *other* features of death, as we experience it, that are separable from the deprivation that death involves, so we then have to ask the question, do these further features *add* to the badness of death? Alternatively, conceivably, some of these further features might mitigate it, reduce its badness in one way or another.

Here's an example. It is, of course, true of you that you are going to die. But more than that: it is *inevitable* that you are going to die. There's no avoiding it. Contrast the inevitability of death with the fact, say, that you are reading this book. You are, of course, reading this book, but it wasn't inevitable that you would read it. You had a choice. But death is different. It doesn't matter what you choose, you can't avoid dying. So it's not just

merely the case that in fact we are all going to die; it's a *necessary* truth that we're all going to die. Accordingly, we might ask, what about this inevitability of death? Does that make things worse? And here I want to distinguish between two issues: there's the individual issue (it's inevitable that *you* are going to die) and there's the universal issue (it's inevitable that we are *all* going to die).

Let's start by thinking about the fact that it's unavoidable that *you* are going to die. Does the unavoidability of death make it better or worse? The interesting thing is, I think it is easy to feel the pull of *both* possible answers here. On the one hand, you can imagine somebody who says, "Look, it's bad enough that I'm going to die, but the fact that there's nothing I can do about it just makes it worse. It's like adding insult to injury that I'm powerless in the face of death. I cannot escape the Grim Reaper. This sheer impotence with regard to a central fact about my existence makes things even worse."

Against that, however, there are those people who want to say that the inevitability of my death actually *reduces* the badness. To get a feel for this position, just think about the idea behind the expression "Don't cry over spilt milk." What's done is done; you can't change it. The idea, of course, is that when you focus on the fact that you can't change something, that thing loses some of its power to upset you. But if that's right, and if we then realize that there is nothing at all I can do about the fact that I am going to die, then perhaps some of the sting of that realization is eliminated as well.

Here's an especially clear example of this: try getting upset about the fact that two plus two equals four. Try feeling upset at your powerlessness to change that fact. Suppose you wanted two plus two to equal five. Can you work up anger and regret and dismay over that? I assume that you can't. You can't get worked up over something so clearly unalterable.

The philosopher Spinoza thought that if we could only recognize the fact—or, at least, what he took to be a fact—that *everything* that happens in life is necessary, then we would get a kind of emotional distance from it; it would no longer upset us. We would no longer be disappointed by things, because to be disappointed in something presupposes that it could have been some other way. And Spinoza thought that once you see that it *couldn't* go any other way, then you can't be sad about it. Well, then, if we see that our death is *inevitable* and we really internalize that fact, perhaps that would reduce the badness of it.

Maybe that's right. But I am not sure. Perhaps you have read Dostoyevsky's short novel *Notes from Underground*. The Underground Man is upset about the fact that two times two equals four. Or rather, he is upset

about the fact that there is nothing at all that he can do about the fact that two times two equals four. He *resents* being so impotent that he can't change the fact that two times two equals four. Similarly, when Descartes thought about God's omnipotence, he suggested that it wouldn't be good enough if God was unable to change the facts of mathematics. Descartes thinks it would be a sign of weakness for God to be stuck with necessities that he couldn't alter. So Descartes claims that God *could* have made it be true that two plus two equals five, but chose not to. In effect, then, Dostoyevsky takes that thought and runs with it. His Underground Man says that it doesn't actually help that something is inevitable. It makes it worse. As I say, I can feel the pull of both sides of this dispute. In different moods, I lean in different directions.

What about the fact that not only is it inevitable that I'm going to die, it's inevitable that we're *all* going to die. Does the *universality* of death make things better or worse? Here too, I find myself pulled both ways. On the one hand, I find myself wanting to say, it's bad that I'm going to die, but I'm not a *monster*. It makes me feel even *worse* that everybody else must die as well. Or perhaps, in light of our discussion of immortality, I should say that what is sad is the fact that we all (or, at least, most of us) die too soon. That makes it even worse.

On the other hand, let's be honest here, we also know the expression "Misery loves company." There's at least some comfort to be had—isn't there?—in the realization that this undesirable thing isn't just true for me. It's not like the universe has singled me out for the injury of dying too soon. That's something that it does to almost everybody. Perhaps there's some comfort in that fact.

Here's a different aspect of death worth thinking about: what about the *variability* of death? After all, it's not just the case that we all die. There's a great deal of *variation* in how much life we get. Some of us make it to the ripe old age of 80, 90, 100, or more. Others of us die at 20, or 15, or 10, or even younger.

Even if death were inevitable, life didn't have to come in different-sized packages. After all, it isn't as though death *has* to involve variability. We could imagine a world in which everybody dies at the same age, perhaps at a hundred. Does it make things worse or better that there's this kind of variability?

From the moral point of view, I suppose, it's fairly straightforward to suggest that it makes things worse. After all, most of us are inclined to think that inequality is morally objectionable. It's bad that, through no fault of their own, some people are poor while other people are rich. But if inequality is morally objectionable, then it's very likely that we're going to think that it is morally horrendous that there is this crucial inequality: some of us die at the age of five while others get to live to ninety. However, in keeping with the general focus of our discussion about the badness of death, I want to put aside the moral question and think about how good or bad it is for *me* that there's variability in death.

We can look at the situation from two basic perspectives: those who get less than the average lifespan, and those who get more than the average lifespan. From the point of view of somebody who gets *less*, this is obviously a bad thing. It's bad enough that I'm going to die too soon. But what's even worse is that I'm going to get even less than the average amount of life. That's clearly an extra bad. But we might then wonder, what about the people who are getting *more* than the average? Suppose we find the median lifespan, the exact length of life such that 50 percent of the people get less than this, and 50 percent get more. Then for every person who gets less than the median amount of life, there's another person who has *more* than the median amount of life. That person gets to say that although it is a pity that he is going to die too soon, at least he is getting more than the average. That seems to count as a plus.

Perhaps these two sides of the equation balance themselves out. There are people who are basically screwed by the fact that they get less than the average amount and there are people who are benefited by getting more than the average amount. Perhaps in terms of the individual badness of death that's a *wash*. Maybe. Except it seems to me there's a further relevant fact about human psychology: we care more about being shortchanged, as we might put it, than we do about being overcompensated. I rather suspect that when people have *less* than the average of something, the amount that this hurts them is *greater* than the amount of benefit that comes to people who have *more* than the average of something. And if that's right—and that seems likely to be the case, especially for something like death—the extra bad arising from the fact that there's variability, and so some people get less than average, outweighs the extra benefit of some people having more than average.

Here's another important feature of death. We considered inevitability; we looked at variability. What about *unpredictability*? Not only is it inevitable that you're going to die, and not only do some people live longer than others—it is also true that you don't know how much more time you've got.

You would be forgiven for thinking that we have already brought in unpredictability, once we start thinking about variability. But in fact, that's not the case. Logically speaking, although variability is a requirement for unpredictability, it doesn't suffice to guarantee unpredictability. You could, in fact, have variability with complete predictability. Imagine, for example, that whenever a baby is born there is a natural birthmark on their wrist that infallibly indicates the precise year, day, and time in which they are going to die. We can certainly imagine a world like this. Death remains inevitable; everybody's got some date marked on their wrist. And more to the point, there could still be variability. Some people live eighty years, some people live fifty-seven years, others only twenty. But there is no unpredictability. Because of the birthmark, everybody knows exactly how much longer they've got.

In our world, of course, we don't have that. In our world, not only do we have variability, we've got unpredictability. Does that make things better? Or does that make things worse? Would it be better to know *when* you were going to die?

One way in which unpredictability has at least the potential of making things worse is this: you don't know how much *more* time you have. So it is hard to make plans. Of course, you can make a guess based on statistics. You can calculate the average lifespan. Suppose that right now, in the United States, it's seventy-nine years. That means that if you are, say, in your late twenties, you have on average another fifty years. But as we have already pointed out, that average covers up a lot of underlying variability. So there you are, you're busy calculating all of this, you're walking across the street, and you get hit by a truck and you die. It could happen, right? Because of unpredictability, you can't really know. And because you can't really know, it's difficult to make the right kinds of plans.

In particular, it's hard to know how to pace yourself. You decide to go off to medical school, to become a doctor. And so not only do you put the time into college, you put the time into medical school, you put the time into your residency, and you put the time into your internship. And that's a very long commitment. It's a long-term plan, which can go wrong if you get sick and die in your early twenties. That's a rather dramatic example, of course, but the same sort of thing can happen in principle to all of us. You make a life plan, figure out what you want to accomplish in your life, and then—boom!—death comes, unexpectedly, screwing up your plans. If only you had known that you were only going to have twenty more years instead of fifty more years, you would have picked a different kind of life for yourself. The unpredictability makes it worse.

Less obviously, it can work the other way as well. You make a life plan, you carry it out, and then, expectedly, you don't die any time close to the time you guessed you would. You continue to stick around, and then your life has this feeling of being anticlimactic. You peaked too soon. You thought you were going to be like the actor James Dean, burn out fast and die young. But you were wrong. If only you had known that you had another seventy years, that you weren't going to die young—if only you had realized that you were going to live to the ripe old age of ninety-four, you would have picked a different life for yourself.

In thinking about these points, I am suggesting, in effect, that the overall value of your life can be affected by some features that we haven't yet talked about. We might put the idea this way: the overall shape of your life matters. Putting the same idea in slightly different language, we could say that "the narrative arc" of your life makes a difference to its overall value.

Let me illustrate the point with some very, very simple graphs. These are not meant to be realistic, but they'll give you the basic idea. The nineteenth-century American author Horatio Alger used to write stories about people who start out poor but who make their way (through hard work, dedication, and effort) to riches and success. Rags to riches—that's a wonderful, inspiring life.

Let's draw a graph of that life. (See Figure 13.1.) Let the Y axis represent well-being, how well off you are at a given moment, and let the X axis represent time. In the first graph in Figure 13.1, you start off with nothing and you end up incredibly well off. That's a great life. That's the Horatio Alger life.

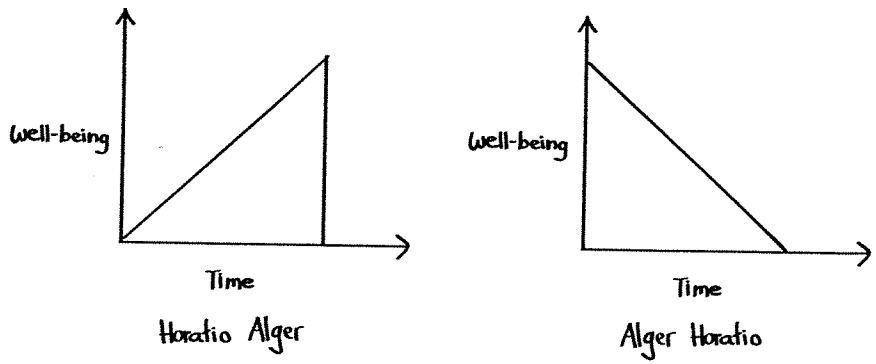

Fig. 13.1

Now consider a different story. This time, instead of going from rags to riches, the person goes from riches to rags. He starts off with everything but ends up with nothing. That's the Alger Horatio life. It is, of course, the reverse of the first story. In Figure 13.1, it's the graph on the right.

Now I doubt if there is anyone who is genuinely indifferent between these alternatives, indifferent with regard to the choice between these two lives. I imagine that pretty much everyone prefers the first life.[1] But notice that in terms of the contents of the life, or at least the *local* contents, it is hard to see why we should care which life we get. Both lives have the same amount of suffering; both lives have the same amount of success. The two graphs are, of course, mirror images of each other, and that means that for every good period in the one, there is an exactly similar good period in the other; for every bad period, there is a corresponding bad period. Roughly but intuitively speaking, the contents of the two lives are the same. (In mathematical terms, the areas under the two lines are the same.) And even if we accept the valuable container theory, and so say that being alive per se is worth something as well, that still doesn't give us reason to prefer one life over the other. Since the two lives last the same number of years, the same number of extra points gets added either way.

If we're not indifferent between these two lives, that seems to show that we think that not only do the various "local" goods make a difference to the value of a life—how happy or unhappy you are at various times—the overall *shape* of the life matters as well. The narrative arc matters. The story "bad to good" is the kind of story we want for ourselves, while the story "good to bad" is the kind of story we don't want for ourselves.

That raises an interesting question: why do we care? And this, of course, should remind us of the puzzle from Lucretius: why do we care more about future nonexistence than past nonexistence? The reason isn't obvious, but the fact remains that when the bad is behind us, that seems less bothersome than when the bad is in front of us. Similarly, it seems, we want the bad to come sooner rather than later. (Recall, as well, the story from Derek Parfit about the painful operation. We're not indifferent between having the bad in the future or in the past.) Whatever the exact explanation may be, the simple fact seems to be that we care about the overall shape and trajectory of our life.

Given that this is the case, however, we have to worry about the possibility that, thanks to the unpredictability of death, our lives may end up lacking the ideal shape. Consider the life shown in Figure 13.2. Here, the problem is that the person peaks too soon. We peak, but then we stick around too long after the climax. A lot of us, I imagine, might be unhappy

with a life like that. Think about your life as though it were a novel, where the graph of your life is like the plot of a great story. It isn't as though we think the denouement must occur on the very last page. It's okay to stick around for a while after. But if the high point of the story occurs in chapter 2, and then there are another sixty-seven chapters after that, you're likely to find yourself thinking that this is not a well-constructed novel.

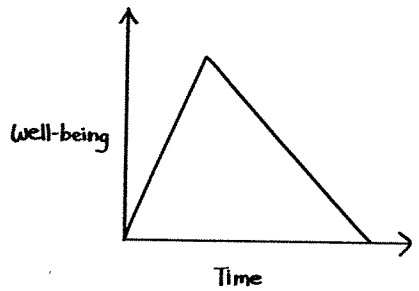

Fig. 13.2

Insofar as we care about the overall shape of our lives, we might worry about life having the *right* shape overall. Where and when do you want your life to peak in terms of your accomplishments? That clearly matters to us, but the trouble is, without predictability you don't know where to put the peak. If you aim for peaking later, you might not make it that long. If you put the peak too soon, you might be around for too long afterwards. All of this suggests, then, that the unpredictability of death adds an extra negative element. It makes it harder to plan what the best way to live my life would be. And from that perspective, it looks as though it would be better to know how much time you've got left.

But then we have to ask, would it really be better to know? Would you really want to know exactly how much time you still have? Suppose we were born with the kinds of birthmarks I was describing earlier, so that you always knew exactly how much longer you had before you were going to die. If you had that kind of birthmark, you would face your entire life with the burden of knowing: I've got 50 years left, I've got 49 years left, 48 years left, 47 . . . Many of us, I think, would indeed find that a burden—something hanging constantly over us, interfering with our ability to enjoy life.

Let's alter the story. Instead of imagining a birthmark, readily visible and interpreted, suppose that there were some sort of genetic marker that could be examined with the right kind of testing. If you wanted to, you

could have your DNA examined, and then you would know exactly how much time you had left. Would you want to get that testing done? Of course, that's science fiction, and I presume it is going to stay science fiction. But the fact of the matter is, as we learn more and more about the various genes that carry different diseases, more and more of us face the question of whether we want to get tested for those diseases.

Imagine that there is a horrible genetic disorder that kills its victims at age forty. You are twenty years old, and you already know that you have a 50 percent chance of having the gene (one of your parents had the gene and died young), but you don't yet know whether you yourself have it. If you do, you are going to die in another twenty years. Would you want to be tested? Would you want to know?

Here's a closely related question: if you did know how much time you had left, would you act differently from what you are doing now? Would that knowledge lead you to refocus your attention on doing the things that are most important to you? Thinking about this question can provide a useful way of recognizing what it is in life that you most value. Ask yourself, what would you choose to do if you knew you only had another year, another five years, another ten years?

There's an old *Saturday Night Live* routine where one of the actors is in the doctor's office, and the doctor gives him the very sad news that he's got two minutes left to live. And he says, "I'm going to pack a lifetime of enjoyment into those two minutes." And then, of course, the point of the skit is that he presses the down button on the elevator and a minute and a half goes by while he's waiting for the elevator to come.

If you knew you had a year left or two years left, what would you do with that time? Would you be in school? Would you travel? Would you spend more time hanging out with your friends? An extremely moving instance of someone having to face this question occurred in the class on death that I teach at Yale. There was a student in that class some years ago who was dying. And he knew that he was dying. He'd been diagnosed with cancer as a freshman. His doctor had told him that he pretty much had no chance of recovery, and indeed he had only a couple more years to live. Faced with that knowledge, he had to ask himself, "Well, what should I do with my remaining years?"

He decided that what he wanted to do was finish his Yale degree. He set himself the goal of graduating college before he died. And there he was, second semester of his senior year, taking my class on death. (It was humbling for me to learn that someone in his situation had decided to take a

class on death, and then have me get up, week after week, and talk about how there is no soul, there is no afterlife, it's a good thing that we are all going to die . . .) So there he was, in my class—until spring break. By spring vacation he had gotten sufficiently sick that his doctor had told him that he could not continue in school anymore. He had to go home. In effect, his doctor had told him that it was time to go home to die. He went home, and his condition rapidly deteriorated after that.

The faculty members who were teaching his various courses that semester then all faced a question posed to them by the administration. Based on the work that he had done so far that semester, what kind of grades were we prepared to assign for the semester as a whole? Because, of course, depending on which of his classes he passed, and which he failed, he either was, or was not, going to be able to graduate. As it turned out, he had done well enough. And Yale, to its very real credit, sent a member of the administration down to his deathbed to award him his degree before he passed away.

It is a very moving and striking story. I'm not sure how many of us would decide that the thing we most want to do with our few remaining years is to spend them in college. But then, what is it that you *would* want to do? What would you choose? And returning to our original question, would knowing how much time you have left be something that would allow you to embrace those choices, so that you could finish your life in the most meaningful way? Or would it be a burden? That's the kind of question we have to face when we think about the fact that typically we don't *know* how much time we've got left. Is that something that increases the badness of death, or does it reduce it somewhat?

Here's yet another feature of death. In addition to the inevitability, in addition to the variability, in addition to the unpredictability, there's the fact that death is, as I like to think of it, *ubiquitous*. I don't just mean the fact that people are dying all around us. I mean, rather, that you yourself could die at any time. There's never any getting away from the possibility that you'll die *now*. Even if we had unpredictability, it wouldn't necessarily follow that death was pervasive in this way. The point I've got in mind here is this: even when you think you're perfectly safe, you could of course die of a stroke. You could die of a heart attack. Even somebody who's young and strong could die from an aneurysm.

Or—one of my favorite examples—you could be sitting in your living room, when suddenly an airplane crashes into your house, killing you. One reads about this sort of thing in the newspaper occasionally. You thought

you were safe. You were watching reruns on television—and the next minute, you're dead. This takes us beyond unpredictability. The fact that you don't know when you're going to die doesn't yet entail that you could die at any moment. But in fact, that's true of all of us as well.

Here is another example, also close to my heart. Once I was driving down the expressway and a car pulled into my lane without looking, clipping my car and causing it to go spinning out of control, careening across three lanes of traffic. The whole thing lasted only a few moments, but I remember thinking quite clearly, "I'm going to die." Now as it happens, I didn't die. I walked away from the accident, and the damage to my car was surprisingly minimal. But it could have happened like *that.*

Death—the possibility of death—is ubiquitous. It's pervasive. So we have to ask ourselves, does this make things worse? It certainly feels, to my mind, as though this is an additional bad aspect of death. It would be nice to get a breather. Imagine, if you will, that there were certain locations, certain vacation spots, where as long as you were there you couldn't die. Wouldn't it be nice to be able to go someplace and just for a period think to yourself, "Well, you know, right now I don't have to worry about that. It doesn't even have to cross my mind."

Of course, if there were death-free zones like this, they'd probably get rather crowded. So perhaps we should change the example. Instead of having death-free zones, imagine that there were death-free *times.* Just suppose that, for whatever reason, nobody could die between noon and one. You could just put it out of your mind. Wouldn't that be nice? To be sure, at one o'clock you would have to take the burden back on. But wouldn't it be nice to have a certain period of time every day when death wasn't so much as a remote possibility? Or suppose there were certain death-free *activities.* Maybe reading philosophy would be something that as long as you were doing it you couldn't die; or perhaps, as long as you were engaged in prayer you couldn't die. Wouldn't that be nice?

Or turn the entire thing the other way around. Suppose that most times and *most* activities were death free, but certain activities introduced the *possibility* of dying. So you couldn't die unless you were engaged in certain activities. You would be *potentially* immortal, in the sense that you *could* live forever, but you wouldn't be *forced* to live forever. Some activities—perhaps putting a gun to your head, for example—would put an end to your life. So even if immortality would be bad, there would be certain things that you could do that would end it. But beyond these special death-*ensuring* activities, imagine that there were various other activities that merely carried the

risk of dying (that is, they carried the level of risk that they actually carry in the real world): while doing these things, you would lose the guarantee that you would not die. And then ask yourself, what sorts of activities would you engage in if you knew that those activities carried with them the risk of dying?

What things are so important to you that you would be willing to suddenly risk death for the sake of doing those things? Perhaps you like art. Is art important enough to you that you would be prepared to look at a masterpiece, if you knew that while you were enjoying it you could die, but that you couldn't die otherwise? Is sex great enough that you'd be prepared to run the risk of dying while you were engaged in sex? We can discover what we think is *most* valuable by asking which activities are *so* valuable that you would be prepared to do them even if that would introduce what wouldn't otherwise be there, namely, the risk of death.

In posing the question that way, I've been assuming that these are things you would do *despite* the fact that they run the risk of death. I suppose, however, that there is a further question we have to ask as well: are there things that would be worth doing precisely *because* of the fact that they introduced the risk of death? To be sure, that new idea sounds rather bizarre. At least, it sounds bizarre if we put aside the possibility that we've now lived for a hundred thousand years and have exhausted all that life's got to offer for us. What seems hard to believe is that we might engage in activities now, while life still has so much *more* to offer, precisely *for* the chance of dying. And yet, it seems to me that there are activities—if not many, at least some—that people do exactly for that reason.

For example, let me tell you something that I know is going to shock you. Did you know that there are people who jump out of airplanes? Admittedly, when they jump, they've got this little piece of cloth that gives them a decent chance of not killing themselves. But these things do fail. Every now and then you read in the newspaper about somebody whose parachute failed to open and so they died. And I ask myself, why? What could possibly drive somebody to jump out of an airplane with nothing but a little piece of cloth between them and death? And the answer that strikes me as most plausible is this: it's the very fact that there's a significant chance of death that helps explain why people do this.

Of course, if you talk to some of these people, they'll say, "Oh, no, no, no. The views are so glorious," or something like that. But I think this is rather an implausible suggestion because you could have these glorious views just by going up in the airplane and looking down from the safety of the plane. Part of the thrill—or so it seems to me—has got to be the very

fact that they now have an increased risk of death. The chance of dying is part of what drives somebody to jump out of an airplane.

But if that's right, then maybe I was wrong earlier when I suggested that it would be nice to have a death-free time or a death-free location or death-free activities. Maybe I was wrong when I said that the pervasiveness of death, the ubiquity of death, is oppressive. If the chance of death adds a kind of zest, then perhaps the ubiquity of death is actually a *good* thing rather than a bad thing.

However, I'm inclined to think that that's not right, not even for those who are indeed drawn to parachuting because of the risk. I imagine, rather, that for such people the ubiquity of death is a kind of constant, unnoticed hum in the background. It's not really good enough to just have *some* risk of death—it's got to be greater risk than usual. What's attractive about jumping out of an airplane is that it *spikes* the risk of death. If that's right, though, then even for those death thrill seekers, the ubiquity of death won't be a particularly good thing precisely because of its ubiquity. The pervasiveness of the risk of death makes it disappear into the background.

There is one further aspect of death that I want to examine as well—the fact that death *follows* life. Arguably, this is the fundamental fact about the human condition. It isn't merely that we live, or that at some point in time we do not exist. Rather, what's true about humans is that we live and *then* we die. And what I want to ask is, what should we think about that fact? There is, after all, a kind of metaphysical compound here, a particular combination of life and death. We need to ask about the overall value of not just life, and not just death, but the entire combination.

One natural thought here, I suppose, would be that when we want to get clear about the value of a compound, we simply need to determine the value of its various component parts and then sum these values. Accordingly, if we want to understand the overall value of the human condition—life followed by death—we should first get clear about the value of life itself, then get clear about the badness of death, and then add the two up. Find the value of the two parts, and see what the sum comes to.

Of course, even given this strategy, people might disagree about the total. Optimists presumably will think that the sum comes out positive. "Yes," they might say, "death is bad. But life is good—sufficiently good to outweigh the badness arising from the fact that we are going to die. On balance, then, it is a good thing to be born." And pessimists will presumably insist that the sum is negative. "On balance," they will argue, "the bad of

death *outweighs* the good (if any!) of life." And moderates might suggest that the answer depends on the facts in the individual case.

But I think there is more to do than merely finding the sum. To evaluate the human condition as a whole, we need to do more than merely add the goodness of life and the badness of death. The issue is actually more complicated than this, because sometimes the value of a compound or combination is different from the total you would get by just thinking about each one of its parts in isolation and then adding up the values of the distinct parts. This simple, "additive" approach to the values of wholes will not always be correct.

Here's an example, to illustrate the point. My two favorite foods in the world are pizza, on the one hand, and chocolate, on the other. I have of course already shared my love of chocolate with you, but I don't think I've mentioned pizza previously. But there they are, my two favorite foods. Pizza—delicious! Chocolate—delicious! Now take these two delicious things and combine them into a chocolate-covered pizza. Yuck! The whole thing just sounds disgusting to me, utterly unappealing.[2] I hope you are with me in finding the idea disgusting. Still, you might not *notice* the disgustingness if you just thought about the value of pizza in isolation and the value of chocolate in isolation. The value of chocolate-covered pizza is not just a matter of summing up the value of the parts taken by themselves. You've got to think about what we might call the "interaction effects."

So let's ask ourselves, are there any interaction effects that need to be taken into account when thinking about the human condition—the fact that life is followed by death? Presumably, there are two main possibilities here. If there are indeed interaction effects, they might be negative—reducing the value of the combination overall—or they might, instead, be positive.

Let me start briefly by mentioning one possible candidate for a *positive* interaction effect. Given that you are going to die, it follows, trivially, that there is a *finite* amount of life that you are going to get. Life is a scarce resource. It's precious. And we might be attracted to the thought that the value of life is increased by its very preciousness. After all, it is a common enough thought that the value of something may be greater if it is fragile or rare. Perhaps, then, the very fact that life is precious, that it won't endure, actually increases its value.

There's a short story by the science fiction writer Orson Scott Card, where the basic idea of the story is that of all the life forms in the universe, we here on Earth are the only ones that are mortal.[3] And because of this we are

the envy of the rest of the universe. It's not so much that immortality is unattractive or boring. It's perfectly fine. But still, other beings throughout the universe envy us for our finite lifespans, because what we've got and they don't have is something that for each individual is precious—something that has to be valued during the limited time that we possess it. I'm not sure whether I agree with this idea, but I certainly can see its appeal. If it's right, then the very fact that we're going to die interacts with our life so as to make it more fragile, more ephemeral, and, as a result of that, more valuable.

Whether or not one accepts this first, positive, interaction effect, there remains the possibility of negative interaction effects as well. Here are two further suggestions that I often find plausible. The first idea I think of under the heading "A Taste Is Just a Tease." It starts with the observation that we live life for a while, getting a feel for all the wonderful things that life could offer us, and then a mere moment later, as it were, it's all snatched away from us. In a way, it adds insult to injury that we are offered just a whiff. It is as though somebody put a delicious meal before a hungry man, allowed him to see what it looked like, allowed him to smell the delicious aromas, perhaps gave him just one little tiny forkful to see just how beautifully delicious the food was. And then they snatched the whole thing away.

You can understand someone who says that it would be better never to have had the taste at all than to have the taste and then not be allowed to have the entire meal. Yet this negative feature is something that you might not notice at all, if you just focus on the intrinsic nature of the taste. In itself, after all, the taste of the meal is positive. Similarly, you might not notice this negative feature if you simply focus on the intrinsic character of *not* having the meal. Not having a meal, after all, is just an absence of a certain experience. The deprivation is a comparative bad, but it is not bad in and of itself. If you want to capture what's excruciatingly undesirable about being teased with a taste and then not being allowed to have more, you need to think about the two in combination. It's an interaction effect. And so we might think, similarly, that one of the bad things about the human condition is that we get a taste of life—nothing more—before it's snatched away. That's one possibility.

The second potential negative interaction effect that I want to mention falls under the heading "How the Noble Have Fallen." Right now, there's something amazing about you and me. We are *people*. In the universe as we know it, that is a remarkably rare and extraordinary thing to be. It is of course impossible to say with confidence what other forms of life there may be out there, but here on Earth, at least, we may well be the only people there are. (Who knows? Maybe dolphins or some of the great apes are people as

well—in the philosophical sense of the term. But at any rate, it's a rather select club.) Of course, according to the physicalist, a person is just a certain kind of machine. But as I have already explained, we're not just *any* old machine. We are *amazing* machines. We're able to love. We're able to write poetry. We're able to think about the farthest reaches of the universe and ask about our place in that universe. People are amazing. And yet, for all that, we end up rotting. We end up as corpses. For many people, there's something horrifying about the thought that something as amazing as us, as exalted and valuable as we are, could end up something as lowly and unimportant as a piece of rotting flesh.

Whenever I think about this idea, the image that always comes to mind is that of a deposed king who ends up waiting on tables to make a living in New York. You might think, reasonably enough, that the life of a waiter is not the worst thing in the world. But at the same time, there's an extra twist—an insult added to the injury. The waiter has to remember that he used to be something extraordinary, a ruler. Notice, however, that if you just think about life as a ruler—if you just think about that part of the combination in isolation—it is likely to seem pretty good. And even life as a waiter, viewed in isolation, isn't all that bad. So if you want to notice the *problem* with this fate, the potential extra negative feature, you've got to think about the fact that there is an entire package here being evaluated. There is, after all, something especially insulting about having gone from king to waiter. And that fate—or worse—is waiting for all of us. It's a fact about the human condition that the amazing things we are don't stay amazing. We turn into pieces of rotting, decaying flesh.

So there are at least three potential interaction effects worth thinking about when evaluating the human condition. On the one hand, there may be negative effects, arising from the fact that a mere taste can be a special form of torture or that there is something horrifying about going from person to corpse. And yet, at the same time, there may be a positive effect as well, arising from the sheer preciousness of life. In different moods I find myself inclined to accept one or another of these different ideas, sometimes all three. Beyond this, it isn't at all clear to me which of these effects, if real, would be greater.

Presumably, people will have different views on this question. Optimists will say that even when we throw in the negative interaction effects, the overall nature of the human condition remains positive. So it is, indeed, a good thing to be born, even though your life is going to be followed by death. In contrast, of course, pessimists will say that the negative aspects of

life are so great—especially when we throw in the negative interaction effects—that it would be better never to have been born at all. In effect, the pessimist's view is that the fact that we are going to die seeps in and poisons the nature of life, or perhaps the nature of the compound as a whole—life followed by death. On balance, they insist, the whole thing is a negative. Better not to have had any of it at all, better not to have been born, than to have this combination package of life followed by death. (Perhaps, then, instead of feeling sorry for Larry, the merely potential and never-to-be born individual from Chapter 10, we should *envy* him.)

Speaking for myself, I'm sufficiently optimistic to think that life can be pretty wonderful. Strictly speaking, though, I am not an optimist, but rather a moderate. There is no single overall value that we should place on the human condition, such that it would be appropriate to say of everyone that it is better to have been born, or better never to have been born. Sadly, it depends on the facts of the given individual's life. But still, it does seem to me that many of us attain lives very much worth living. And even in those moods where I am inclined to agree that we must not forget to take into account one or the other of the negative interaction effects, it still seems to me that for many of us—perhaps most of us—our situation is, on balance, a good one. For those of us fortunate enough to get a reasonable taste of what life can offer, then despite the fact that life must be followed by death, it is, I think, better to have been born than to have never been born at all.

Nonetheless, I want to emphasize the point that even if we were to *accept* the pessimist's conclusion that it would be better never to have been born at all, it wouldn't yet follow that the appropriate response to this realization is to commit suicide. That would take some further arguments.

It is, of course, tempting to think otherwise. That is to say, it is tempting to think that if we decide that it is better never to have been born, it follows without further ado that suicide is the appropriate response to the human condition. But in fact, as a matter of logic at least, that doesn't follow at all. Because if you think about it for a moment, you can readily see that suicide doesn't change the fundamental nature of the human condition—life followed by death. It isn't as though if you kill yourself you somehow bring it about that you've never been born at all! If there is something horrible about having just a taste, for example, then the fact remains that even if you kill yourself, all you have been given is a taste. Indeed, if you commit suicide, you've simply made it an even smaller taste. Similarly, if there is something degrading about being a person who is going to become a corpse,

committing suicide doesn't alter that fundamental fact either. It just makes the insult come all the sooner.

So even if we were to agree with the pessimist that it would be better never to have been born at all, we still need to say—in the language of the old joke—show me one person in a thousand who is so lucky! *We* have all been born. And from the fact, even if we were to agree that it is a fact, that it would have been better if we hadn't been born at all—even if that were true, it simply wouldn't follow that suicide was an appropriate response.

Of course, none of this shows that suicide *isn't* ever an appropriate response to one's situation. That's a topic we will turn to in our penultimate chapter. But let's hold off on thinking about that question for a bit longer. First, I think, we need to ask a more general question: how should one *live*, in light of the facts about death that I have been laying out up to this point? Indeed, we also need to ask: *should* death affect the way we live at all?