

How much Better than Death is Ordinary Human Survival?

Ivar R. Labukt

1. Introduction

It turns out that Peter, who is 20, has a neurological disease and will die a year from now. If we look at it impartially, this is bad news. Peter would otherwise have gone on to live a rich and rewarding life in which he also contributed to the happiness of several other human beings. Still, from the point of view of the universe, the death of an individual like Peter is relatively insignificant. From Peter's perspective, however, things look different. He takes the news to be terrible. Though he worries about the suffering he might have to go through, his primary concern is that he will die so young. He realizes that his death only makes the world a very slightly worse place, but thinks that, for *him*, it is a tragedy. He will, after all, miss out on almost a lifetime of happiness.

According to *Epicureanism*, it actually makes no sense to say that death is bad for the person who dies. There is thus no rational basis for Peter's despair. *Deprivationists*, on the other hand, think that death can be bad from a personal point of view if it deprives the person dying of life worth living. I will assume that the deprivationists are right about this.¹

However, I will argue that Peter and many other deprivationists, including some of the contributors to this volume, significantly overestimate the magnitude of the personal badness of death. It is actually not that bad. This may sound like excellent news. But death only fails to be particularly bad because what it deprives us of – ordinary human survival – fails to be

¹ For arguments against Epicureanism, see e.g. Solberg (this volume).

particularly good. Personally or egoistically speaking, survival is much closer to death than we think.

This claim is controversial and concerns highly complicated issues. I cannot address all details relevant for defending it in this chapter, but I hope to show that it is quite plausible. I will proceed by examining the most prominent philosophical theories of why survival matters. I argue that all these theories fail to offer something that is both deeply egoistically important and found in ordinary human survival. In the final section, I discuss how the claim might affect the justification of healthcare policies.

2. Non-reductionism

Before I began studying philosophy, I thought that my existence involved something permanent. I realized that my mental life changed a lot, but it still seemed to me that it was always the same entity – the real “me,” so to speak – that “had” this mental life. In technical terms, I was a non-reductionist about personal identity. I did not think that being “me” could be reduced to ordinary physical or psychological properties. It was, to use Derek Parfit’s expression, an important and unanalyzable “further fact” (Parfit 1984, 201).

If human beings have this kind of permanent irreducible self, that would arguably provide us with strong personal reasons to avoid death. There would be a robust and straightforward sense in which Peter could stay around and have his future happiness.² His death deprives him of this possibility and is, for that reason, a terrible thing from Peter’s perspective.

Now, like most philosophers, I am skeptical of non-reductive conceptions of personal identity. It is unclear precisely what the deep further fact would be, and how it would relate to

² Or perhaps not; see Johansson 2007.

ordinary physical and psychological facts. One way of trying to make sense of it would be to say that “I” am an immaterial soul, which resides in my body, or my brain, throughout my life. But this raises a host of new metaphysical difficulties. Dualism of this kind has largely been abandoned.

Though we may not be in a position to rule out non-reductionism entirely, I will, as is common in the literature, assume that it is false. In other words, if we are to come up with personal reasons for why survival is important and death is bad, these have to be based on ordinary physical or psychological properties. In the following, I examine three different attempts to do this.

3. Psychological Continuity

In ordinary human survival, there is a high degree of psychological continuity between different parts of our lives. For instance, we remember earlier experiences and make plans that we execute later. There are also psychological properties such as temper or cognitive aptitude that remain fairly stable throughout most of ordinary human lives. Some authors argue that such forms of psychological continuity are sufficient to justify egoistic concern about one’s survival and future. Let us call this the psychological continuity view.³

The psychological continuity view gives us a possible justification for Peter’s distress. If he had not been ill, there would, in the future, exist a person who would remember many of the events that took place in the first 20 years of Peter’s life. This person would do some of the things that Peter intends to do, and he would share many of Peter’s values and character traits. Peter has a strong personal reason to want this kind of person to exist (at least if he will

³ For a more carefully worked out version of the view, see e.g. Parfit 1984, part 3.

be happy, as we have assumed). His death precludes this from happening and is, therefore, bad.

When evaluating this proposal, we need to make sure that we do not find psychological continuity important because we confound it with other things that seem to matter. A good way of doing this is to imagine psychological continuity obtaining without other forms of continuity found in ordinary human survival. Let's say that some amazing new technology makes it possible to produce a complete physical copy of Peter, only without the neurological disease. The process would kill (the old) Peter, but the replica would have exactly the same memories and psychological traits, and it would go on to do the things that Peter would have done if his disease had been cured directly. We can now make the following argument:

- (1) If the psychological continuity view is correct, having a replica made is as good as ordinary human survival.
- (2) Having a replica made is not as good as ordinary human survival is typically taken to be.
- (3) If the psychological view is correct, ordinary human survival is not as good it is typically taken to be.

The first premise is clearly correct. Since there are no psychological differences between Peter and the replica, there would be the same degree of psychological continuity between Peter and the future replica as there would have been between Peter and the future Peter in the absence of the disease. (What would be lacking is physical continuity, since Peter and the replica are not made of the same physical matter.)

The second premise is intuitively plausible. We would not expect Peter to be as pleased with the replication technology as he would have been with a cure. Having the replica made could, of course, be a good idea for non-egoistic reasons. Its happiness would increase the

total sum of well-being in the universe. More specifically, it could take care of Peter's parents and bring joy to his friends. But we are focusing here on what is good *for Peter*. And then replication seems decidedly less attractive than we typically take ordinary survival to be.⁴

A stronger claim, which seems correct to me, is that replication does not have any egoistic merit whatsoever. I would feel great sympathy for a replica, but not egoistic concern. From a purely egoistic perspective, I think Peter would be making a mistake if he sacrificed some of his happiness to ensure that a replica was made. If this is correct, psychological continuity by itself is not a source of egoistic reasons at all. Several authors agree with this claim.⁵ Derek Parfit, the most prominent defender of the psychological continuity view, admits that he cannot refute it. However, he also thinks it is reasonable to hold that there are egoistic reasons to care about a replica – and, more generally, that psychological continuity does justify some special concern for one's own future (Parfit 1984, 307-12). I will not argue against this here. For my present purposes, all I need to claim is that the concern would be significantly weaker than the one displayed by most actual people. With this, Parfit agrees. He does not think that replication could be as good as we have taken survival to be. His view, rather, is that "*ordinary survival is about as bad as being destroyed and Replicated.*" (Parfit 1984, 280; italics in original)

I conclude that if survival and death matter as much from a personal perspective as we have assumed, something more than psychological continuity must be at stake.

4. Continuation of the Capacity for Consciousness

⁴ For a somewhat more elaborate defense of this claim, see McMahan 2002, 55-59.

⁵ For an overview, see Johansson 2007, 641-643.

Another possibility is that survival is important because it involves the continuation of one's consciousness. The problem with making the replica, we could then claim, is that it does not allow Peter's conscious life to continue. Instead, we just create another consciousness that is exactly similar to Peter's. What is important from Peter's perspective is that *his* consciousness should extend into the future. In order to evaluate this proposal, we need to be clear about exactly what it is that is supposed to continue.

Since ordinary human lives contain frequent episodes of unconsciousness, death cannot be (significantly) bad simply because it interrupts consciousness. Rather, the problem has been taken to be that it destroys our *capacity* for consciousness (see McMahan 2002, 67; Dainton and Baine 2005; Crisp 2006, 126-131). This capacity can be present even if there is no actual consciousness, and some form of it is found throughout all ordinary human lives. Thus, ordinary human survival would be important if a continued capacity for consciousness is important. But is it?

Let's return to Peter. After a visit to the doctor, he is informed that he possesses a unique immune system that will be able to deal with the disease. The process will take a few days, during which Peter will be in a coma. His brain will then be utterly incapable of producing consciousness. The authors who defend continuity of consciousness accounts do not specify exactly what it takes for the capacity of consciousness to disappear.⁶ But let's just say that the immune system causes the necessary changes, whatever they are. Then the brain gradually regains its former structure and functional capacities, including the ability to produce consciousness. When Peter wakes up, his brain contains the same physical matter organized in the same way as a few days before (except for the disease, of course).

⁶ McMahan's version of the account is the least vague on this point (see McMahan 2002, 66-69).

If continued capacity for consciousness is the crucial ingredient in survival, what Peter will be going through is as bad as death. But what Peter will be going through does not seem to be particularly bad. In a few days, there will be a healthy brain inside his skull, comprised of the same physical matter and giving rise to the same memories, psychological traits and conscious experiences that he would have had if the disease just magically disappeared. What's the problem?

Non-reductionists might legitimately worry about Peter's fate. For instance, the soul residing in the original Peter might leave during the coma, another one taking its place when the coma is over. We are, however, proceeding on reductionist assumptions. This makes it much harder to see how something very bad could be happening to Peter. It is not as if the molecules or neurons will somehow "remember" that they have been apart, thereby reducing or obliterating the value of their future joint existence from Peter's point of view.

Perhaps we should not put too much weight on a single imaginary example. Here is a more general argument. Supporters of the continued capacity for consciousness view want to be able to say that normal periods of unconsciousness, such as dreamless sleep, are perfectly fine from an egoistic point of view. But surely there is *some* sense in which the brain lacks the capacity to produce consciousness during dreamless sleep. So which notion of a capacity for consciousness is supposed to play the important ethical role? Consider the following cases:

(1) The consciousness-generating part of your brain remains entirely intact but is told to shut down by another part of the brain. It can't produce consciousness until something changes in the other part of the brain. (This appears to be roughly what happens during dreamless sleep.)

(2) A drug disrupts the activity in the consciousness-generating part of your brain. The brain is completely incapable of producing consciousness until the drug is metabolized. (This appears to be what happens during general anesthesia.)

(3) A neurosurgeon modifies small but functionally necessary parts of a number of neurons in the consciousness-generating part of the brain that is sufficient to preclude the brain from producing consciousness. He then reverses the modifications; the whole procedure takes less than an hour.

(4) A neurosurgeon kills a number of neurons in the consciousness-generating part of the brain that is sufficient to preclude the brain from producing consciousness. He then reassembles these neurons using the same atoms in the same places. The whole procedure takes less than an hour.

(5) A neurosurgeon removes the consciousness-producing part of your brain. He puts it in a blender and hits the puree button. He then puts all the atoms back exactly the way they were. The whole procedure takes less than an hour.

Which of these scenarios constitutes a break in the capacity for consciousness in the ethically relevant sense? No matter where we draw the line, the difference between the cases on each side of the line seems trivial. If the line separates death from survival, that makes death trivial, too.

To be sure, supporters of the continuity of consciousness view do not have to hold that there is a precise point where all basis for egoistic concern suddenly disappears. It could be gradually weakened. There might even be an area of indeterminacy where it is not clear

whether there is a basis for egoistic concern at all. We could then avoid implications such as “(3) is as good as ordinary survival while (4) is as bad as death”. But we would have to say things like: “in (3), it is determinate that you have egoistic reasons, but in (4) it is indeterminate whether you have egoistic reasons”, or: “in (3), you should retain 80% of normal egoistic concern for the future, but in (4), you should only retain 40%”. If survival matters as much as we think, these would be highly significant differences. It seems implausible that they could be based on distinctions between different notions of a capacity for consciousness as subtle as those above.

5. Physical Continuity

Instead of focusing on continuity in the *capacity* of consciousness, we could try to focus on continuity in the *physical basis* for consciousness – which for human beings is the brain. Thus, what matters in survival would not be that one’s brain at all times has some capacity to produce consciousness, but merely that at some point it will do so. Thomas Nagel seems to accept this view. He thinks that the basis for egoistic concern is the physical continuity of the brain and that this basis remains intact even in the case of a “radical break in the continuity of consciousness.” (Nagel 1986, 45)

This gives us another construal of Peter’s fear of death. The problem is that, after a year, his brain will never be conscious again. This explains why the replica is insufficient: the replica does not have Peter’s brain, but only one exactly like it. At the same time, we avoid all the problems discussed in section 4. Could this, then, be the important personal problem with death we have been looking for? To evaluate this proposal, we need to determine what we mean by a brain persisting.

One possibility is that the brain has to be composed of the same basic physical matter. If this requirement is satisfied, there is, arguably, a basis for egoistic concern. If I have a reason to care about the experiences produced by the physical matter in my brain right now, and if exactly the same physical matter will produce experiences at some point in the future, it seems plausible that I should care about those experiences, too.

The problem with this proposal is that brains are not composed of the same basic physical matter throughout ordinary human lives. Most neurons are formed in early childhood and survive into old age, but the molecules they consist of are continuously being replaced (see e.g. Spalding et al. 2005). I am not aware of any reliable estimates of the precise molecular turnover rate in brain cells, but it seems safe to assume that substantially more than half of the basic matter in an ordinary human brain is replaced each year.

If sameness of basic matter is required for egoistic concern, it follows that we should adopt a discount rate corresponding to the molecular turnover rate in our brains. Ordinary human survival would still matter from a personal point of view, but it would be far closer to death than we have assumed. Peter should not worry too much about his disease. In a year, most of him will be gone in any case.

In everyday language, we allow that physical objects can persist even if the basic matter they are composed of is gradually replaced. A classic example is a ship that goes through a large number of small repairs until one day there is nothing left of the original wooden material. In *this* sense, one and the same brain does exist throughout (most of) ordinary human lives. But does it matter whether our brains are physically continuous in this sense?

Let us look at one last version of Peter's story. Peter is contacted by a brilliant neurosurgeon who is capable of replacing his brain with a perfect healthy copy. Peter objects that this only gives him psychological continuity and no physical continuity. "What I want is for *my* brain to be healthy," he explains. "Relax," the neurosurgeon answers. "That won't be a

problem. Instead of transplanting the brain in one big operation, I will perform 365 small operations each day, removing and adding only a few grams of neuronal matter each time. The replacement will thus be so slow and incremental that your brain will still exist in a year. What do you say?"

It seems obvious to me that the surgeon's procedure fails to solve the problem. If replacing the brain in one operation does not provide Peter with a future, doing the same thing in several steps won't help either. The step-by-step procedure might be better for Peter during the next year since it allows for some of his brain to remain intact during this period. But it would not be better for Peter after a year has passed. No matter which procedure he chooses, he will then have the same psychology and the same conscious experiences, produced by the same physical matter organized in the same way.

Again, we must make sure that we do not slip into thinking of survival in a non-reductionist way. If there had been a soul residing in Peter's brain, we could perhaps avoid scaring it away by replacing the physical matter very carefully. But if there is no soul, or any other deep and significant further fact, it becomes something of a mystery how it could matter so much whether the replacement takes an hour or a year.⁷

Not everyone shares my intuitive response to this case. Jeff McMahan, for instance, thinks that the timing of the removals and insertions of brain matter would make a significant difference (McMahan 2002, 69-71). As in the discussion of the capacity of consciousness view, it might be helpful to consider precisely where the basis for egoistic concern would disappear.

Let's say that replacement of the brain involves m steps and n days. Since any difference between m and $m+1$ or n and $n+1$ is trivial, we cannot plausibly claim that there is a value of

⁷ Parfit makes similar observations about an analogous case (Parfit 1984, 474-5).

m and n such that m and n yield no egoistic reasons whatsoever while $m+1$ or $n+1$ yield strong egoistic reasons.⁸ In order to solve this problem, McMahan suggests that for some values of m and n , it may be indeterminate whether Peter has a basis for egoistic concern. Moreover, for some of the values of m and n where it is clear that there is such a basis, the concern it justifies is weaker than it would be if the numbers had been greater (McMahan 2002, 71-2). Even if McMahan thus avoids drawing a sharp and ethically arbitrary line between survival and death, it seems that his view will still end up with some rather implausible implications. For instance, the procedure might eradicate all basis for egoistic concern if m and n are between 1 and 25, but somehow neither clearly eradicate nor secure such a basis if m and n are 26. As m and n increase from 100 to 300, Peter's egoistic reasons might go from being very weak to having the strength found in ordinary human survival. A further increase in m and n from 300 to 500 will then be ethically irrelevant. Staying within a reductionist framework, it is not particularly credible that a gradual increase in the number of days or procedures could have such odd normative implications.

In other words, there does not seem to be any plausible way of drawing an ethically significant distinction between instant replication and the gradual replacement of matter found in ordinary human survival. If the former yields little or no basis for egoistic concern, so does the latter.

6. Taking Stock

⁸ Albeit in a somewhat different context, Theron Pummer (this volume) also rejects the idea that slight factual differences between two cases could lead to very large ethical differences.

None of the reductionist accounts I have examined appears capable of providing ordinary human beings with strong long-term egoistic reasons for avoiding death. Continuity in the capacity for consciousness and physical continuity of the kind that is compatible with gradual replacement do not seem to matter all. The same arguably holds for psychological continuity. If it doesn't, the egoistic concern this form of continuity justifies is still quite weak. Physical continuity in the sense of the brain being composed of the same basic physical matter might yield strong egoistic reasons for the near future, but the object of egoistic concern would be more or less completely gone in a few months (or at best a few years). If nothing besides these forms of continuity matters in human survival, there is little or no personal reason to worry about the fact that you will die in, say, 10 or 50 years. If you are about to die very soon, you have more reason to be distressed. But we're hardly talking about a tragedy.

There might, of course, be other reductionist forms of continuity that matter in human survival. Examining all the possibilities would require more space than I have in this chapter. I will therefore just state my opinion that there are no preferable reductionist alternatives. Assuming that non-reductionism is not an option either, I am led to the conclusion that from a personal point of view, ordinary human survival is not much better than death.

Some might find this conclusion too absurd to take seriously. However, I do not think that we should put too much weight on our pre-theoretical intuitions in this context. I suspect that our normative intuitions about death and egoistic concern to a great extent are based on a pre-theoretical commitment to some non-reductive conception of personal identity. If we are prepared to abandon the latter, we should be prepared to abandon, or at least revise, the former. We should not simply search for something other than an irreducible self that is found throughout ordinary human lives and assume that it must be just as important as we thought the irreducible self was. We should look carefully at what survival actually involves and ask

ourselves whether any of it is worth caring about deeply and personally, and we should do so with an open mind.

7. Implications for Health Priorities

So far, I have only discussed the badness of death from a personal perspective. However, it is natural to assume that the personal badness of death is also relevant from a policy perspective. In fact, the chapters by Ben Bradley, Espen Gamlund, Joe Millum and Andreas Mogensen in this volume all defend or proceed on the idea that the primary reason why we as a society should prevent a given death is that the death is bad for the person who dies. In this section, I examine how my claim about the personal badness of death affects this way of justifying priorities in mortality reduction.

The idea that we should prioritize in accordance with the personal badness of death can be cashed out in (at least) two different ways:

(1) From a policy perspective, the current importance of preventing a death that might occur at time t is a function of how bad the prospect of dying at t is *right now* for the person who might die.

(2) From a policy perspective, the current importance of preventing a death that might occur at time t is a function of how bad dying at t would be at t for the person who might die.

Let us begin with (1). If personal badness is crucial for health priorities, and if deaths are generally less bad for the dying person than we have assumed, it might seem obvious that we should spend fewer resources on preventing them. However, if we base policy choices on the present egoistic reasons of those affected when it comes to death, we should presumably do so

in other cases as well. The reasoning in the previous sections applies to all egoistic concern for the future, not just concern about survival and death. If we have weaker personal reasons than we think to make sure that we will be alive in 20 years, we also have less reason to, say, get an education or save for old age. Thus, on (1), we would not necessarily be forced to make any changes in how we prioritize between saving lives and other societal aims. However, when it comes to which deaths we should be most concerned with preventing, our thinking would have to change quite drastically. As noted in section 6, the things that might be worth caring about egoistically fade away quite rapidly with time in ordinary human lives. From a personal point of view, the prospect of dying soon is therefore much worse than the prospect of dying in a few years. This means that, if we accept (1), we should allocate most of our resources to the prevention of imminent deaths. It would be much more important to provide emergency health care than to, say vaccinate against deadly diseases or encourage healthy changes in life style. This way of making health priorities is clearly at odds with current practice. It seems objectionably short-sighted.

(2) avoids shortsightedness of this kind, since the personal badness of death is assessed at the time of death rather than in the present. We could take it as equally important to prevent a one-year-old from dying in 20 years as it is to prevent Peter from dying in one year. However, if my earlier arguments are sound, (2) does seem to entail that preventing deaths is generally much less important than we have assumed. At least this is so if policymakers are supposed to assess all goods and evils – and not just death – at the time they might occur. Nothing I have said in this chapter entails that goods and evils *within a life* are, at the time they occur, less important than what is typically assumed. I have only offered reasons for thinking that death is less important in this sense. Thus, on (2), my view implies that our commitment to preventing deaths should be revised downwards while our commitment to improving lives should remain the same. This would obviously affect our healthcare priorities. We should

transfer a significant portion of the resources spent on decreasing mortality to measures that reduce morbidity. Precisely where we would end up depends on what we take to be the correct starting point in the absence of worries about the object of egoistic concern, as well as empirical questions about the rate of decay of physical (and perhaps psychological) continuity. However, it seems clear that we would have to endorse some quite radical policies. Imagine, for instance, that we can choose between the following two treatments:

A: The patient's quality of life is improved from 0,5 to 1 for 20 years.

B: The patient's life is extended for 20 years with a quality of life of 1.

In the absence of any further morally relevant differences between the treatments, all existing proposals for priority setting in healthcare would favor B over A. Intuitively, this clearly seems the right answer. But if the death that is prevented in B is a lot less bad for the patient than we have assumed, and if this kind of badness is our only reason to prevent it, we would presumably have to reverse our judgment.

There might be other ways of basing healthcare priorities on the personal badness of death. However, I cannot think of any version of this idea where my view on the personal badness of death would not have rather drastically revisionist policy implications. One could respond to this in either of two ways. The first is to claim that if my view is correct – and death really is much less of a problem for the person who dies than we have assumed – radical policy changes do in fact seem warranted. The second, which I find more reasonable, is to conclude that the personal badness of death should not be the sole (or primary) basis for health priorities. There are other ways of justifying these priorities. We might, for instance, adopt the

point of view of the universe, as a total utilitarian would do.⁹ From this point of view, there is a reason to give Peter a long and happy life even if his present egoistic reasons only apply to a small part of this life.

References

- Bradley, Ben. A Gradualist View About the Badness of Death. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 9). New York: Oxford University Press.
- Crisp, Roger. 2006. *Reasons and the Good*. Oxford: Oxford University Press.
- Dainton, Barry & Bayne, Tim. 2005. "Consciousness as a Guide to Personal Persistence." *The Australasian Journal of Philosophy* 83, 4: 549-571.
- Gamlund, Espen. 2017. Age, Death and the Allocation of Scarce Health Care Resources. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 10). New York: Oxford University Press.
- Johansson, Jens. 2007. "Non-Reductionism and Special Concern." *Australasian Journal of Philosophy* 85, 4: 641-657.
- Kerstein, Samuel. The Badness of Death for Us, the Worth in Us, and Priorities in Saving Lives. 2017. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 13). New York: Oxford University Press.
- McMahan, Jeff. 2002. *The Ethics of Killing*. Oxford: Oxford University Press.

⁹ For some other suggestions, see Broome (this volume) and Kerstein (this volume).

- Millum, Joe. Putting a Number on the Harm of Death. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 11). New York: Oxford University Press.
- Mogensen, Andreas. Life Years at Stake: Justifying and Modelling Acquisition of Life-potential for DALYs. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 12). New York: Oxford University Press.
- Nagel, Thomas. 1986. *The View from Nowhere*. New York: Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Pummer, Theron. The Worseness of Nonexistence. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 2). New York: Oxford University Press.
- Solberg, Carl Tollef. Epicurean Challenges to the Disvalue of Death. In Espen Gamlund & Carl Tollef Solberg (Eds.), *Saving Lives from the Badness of Death* (Chapter 3). New York: Oxford University Press.
- Spalding, Kirsty L., Bhardwaj, Ratan D., Buchholz, Bruce A., Druid, Henrik & Frisé, Jonas. 2005. "Retrospective Birth Dating of Cells in Humans." *Cell* 122: 133-43.