

PUZZLES OF SURVIVAL

tele-
transportation

My Division

the survival
spectrum



Our starting point today is a thought-experiment, from Derek Parfit.

“I enter the Teletransporter. I have been to Mars before, but only by the old method, a space-ship journey taking several weeks. This machine will send me at the speed of light. I merely have to press the green button, Like others, I am nervous. Will it work? I remind myself what I have been told to expect. When I press the button, I shall lose consciousness, and then wake up at what seems a moment later. In fact I shall have been unconscious for about an hour. The Scanner here on earth will destroy my brain and body, while recording the exact states of all of my cells. It will then transmit this information by radio. Traveling at the speed of light, the message will take three minutes to reach the Replicator on Mars. This will then create, out of new matter, a brain and body exactly like mine.

It will be in this body that I shall wake up.

Though I believe that this will happen, I still hesitate. But then I remember seeing my wife grin when, at breakfast today, I revealed my nervousness. As she reminded me, she has been often teletransported, and there is nothing wrong with *her*. I press the button. As expected, I lose and seem at once to regain consciousness, but in a different cubicle. Examining my new body, I find no change at all. Even the cut on my upper lip, from this morning’s shave, is still there.”


A continuation of Parfit's case shows that it is not as unproblematic as it at first appears.

“Several years pass, during which I am often Teletransported. I am now back in the cubicle, ready for another trip to Mars. But this time, when I press the button, I do not lose consciousness. There is a whirring sound, then silence. I leave the cubicle, and say to the attendant, ‘It’s not working. What did I do wrong?’ ‘It’s working,’ he replies, handing me a printed card. This reads: ‘The New Scanner records your blueprint without destroying your brain and body. We hope that you will welcome the opportunities which this technical advance offers.’

The attendant tells me that I am one of the first people to use the New Scanner. He adds that, if I stay for an hour, I can use the Intercom to see and talk to myself on Mars.


‘Wait a minute,’ I reply. ‘If I’m here I can’t *also* be on Mars.’

The teletransporter was invented as a way of traveling quickly from Earth to the now-colonized planet of Mars. One simply steps into the teletransporter on Earth, at which time all of the data about my cells is recorded and transmitted near the speed of light to Mars, at which time that data is used by the teletransporter there to reconstitute me.




This looks pretty unproblematic, from the point of view of the psychological theory. After all, the being that emerges on Mars has exactly the memories and personality as the person who stepped into the machine on Earth. So it is the same person.

But problems are not far away. What happens if the teletransportation machine on earth, after copying all of the information about the cells of the person who steps into the teletransporter, simply leaves the body in the teletransportation machine untouched? This is the case of the 'New Scanner.'



Or we can imagine that there is another teletransportation machine located on the surface of Venus, to which the machine on earth simultaneously transmits the relevant cellular information.



We appear to face a problem which is in some ways similar to the problem posed by the Ship of Theseus. Let's focus on the version of the story in which there are teletransportation machines on Mars and Venus.

Let's call the person who steps into the teletransporter on Earth 'Earthy,' the one who steps out on Mars 'Marsy,' and the one who steps out on Venus 'Venusy.'

We have already seen that, if the psychological theory is true, then the idea that a single person can travel (and continue to exist!) via teletransportation is unproblematic. So we know that, if the psychological theory is true, then:

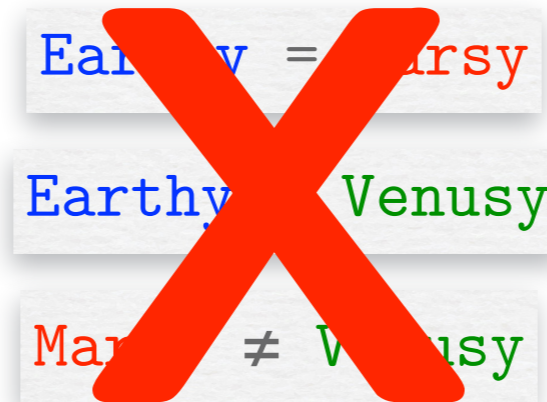
Earthy = Marsy

Earthy = Venusy

But the following seems clearly true:

Marsy ≠ Venusy

But, for reasons we have already discussed — namely, the fact that identity is transitive — these three claims do not sit well together. So it appears that the psychological theory implies a contradiction.



Earthy = Marsy
Earthy = Venusy
Marsy ≠ Venusy

Basically the same point could be made about the version of the story on which, after the transmission to Mars, the individual who steps into the teletransporter on Earth steps back out. To tell that version of the story, we'd just need to introduce two names — Earthy-1 and Earthy-2 — for the individual on earth pre-teletransportation, and the individual who exists after the teletransportation.

It is easy enough to turn this into an argument against the psychological theory of survival.

It is easy enough to turn this into an argument against the psychological theory of survival.

THE TELETRANSPORTATION ARGUMENT

1. Teletransportation is possible.
 2. If teletransportation is possible and the psychological theory of survival is true, then it is possible that Earthy = Marsy and Earthy = Venusy and Marsy \neq Venusy.
 3. It is not possible that Earthy = Marsy and Earthy = Venusy and Marsy \neq Venusy.
 4. Either teletransportation is impossible or the psychological theory of survival is false.
(2,3)
-
- C. The psychological theory of survival is false.
(1,4)

These cases look bad for the psychological theory of survival. One might think that they provide a reason to favor a materialist view. What would a materialist say about teletransportation?

These cases look bad for the psychological theory of survival. One might think that they provide a reason to favor a materialist view. What would a materialist say about teletransportation?

This kind of example is called a case of **fission**, because (on the psychological view of survival) it looks like a case in which one person becomes two people.

At a first glance, this looks problematic only for the psychological view of survival. But, as it turns out, cases of fission can be used to make trouble for other theories of survival as well.

Let's return to a case we discussed earlier: the example of the Brain Transplant.

Let's return to a case we discussed earlier: the example of the Brain Transplant.

“Suppose first that I am one of a pair of identical twins, and that both my body and my twin's brain have been fatally injured. Because of advances in neuro-surgery, it is not inevitable that these injuries will cause us both to die. We have between us one healthy brain and one healthy body. Surgeons can put these together.

If all of my brain continues both to exist and to be the brain of one living person, who is psychologically continuous with me, I continue to exist. This is true whatever happens to the rest of my body. ...”

As Parfit points out, we can also imagine a case in which only one hemisphere of your brain survives.

As Parfit points out, we can also imagine a case in which only one hemisphere of your brain survives.

“It is in fact true that one hemisphere is enough. There are many people who have survived, when a stroke or injury puts out of action one of their hemispheres. With his remaining hemisphere, such a person may need to re-learn certain things, such as adult speech, or how to control both hands. But this is possible. ... [So] I would survive if my brain was successfully transplanted into my twin’s body. And I could survive with only half my brain, the other half having been destroyed. Given these two facts, it seems clear that I would survive if half my brain was successfully transplanted into my twin’s body, and the other half was destroyed.”

Let’s call this the case of **Hemisphere Transplant**. It seems, for reasons Parfit gives, that if you could survive Brain Transplant, you could also survive Hemisphere Transplant.

Let's call this the case of **Hemisphere Transplant**. It seems, for reasons Parfit gives, that if you could survive Brain Transplant, you could also survive Hemisphere Transplant.

But if you are on board with this, that leads to a case about which it is really hard to know what we should say.

“My Division. My body is fatally injured, as are the brains of my two brothers. My brain is divided, and each half is successfully transplanted into the body of one of my brothers. Each of the resulting people believes that he is me, seems to remember living my life, has my character, and is in every other way psychologically continuous with me. And he has a body that is very like mine.”

Let's call the resulting people Lefty and Righty. Then we seem to have four options.

Let's call the resulting people Lefty and Righty. Then we seem to have four options.

- (1) You survive as Lefty.
- (2) You survive as Righty.
- (3) You survive as both Lefty and Righty.
- (4) You do not survive.

It is hard to see how (1) or (2) could be true, since nothing seems to favor one over the other.

But (3) cannot be true, since Lefty \neq Righty.

So it looks like (4) must be true: you do not survive.

Here's the puzzle: in both Hemisphere Transplant and My Survival, one of your hemispheres is successfully transplanted into a body which goes on living. Given that, how could you survive in one case but not the other?

Here's the puzzle: in both Hemisphere Transplant and My Survival, one of your hemispheres is successfully transplanted into a body which goes on living. Given that, how could you survive in one case but not the other?

This looks puzzling from the point of view of both materialist views and the psychological view. (We'll ask what the dualist should say about this case in a second.)

It would appear that these theorists have two options.

Say that you survive neither Hemisphere Transplant nor My Division.

Try to find some relevant difference between Hemisphere Transplant and My Division which could explain why you survive one but not the other.

Say that you survive
neither Hemisphere
Transplant nor My
Division.

Let's try out the first option first.

One person who has a principled reason for making this move is the proponent of the idea that I survive only if the organism which I am survives; as we've already seen, this view seems committed to denying that you survive Brain Transplant (let alone Hemisphere Transplant).

The proponent of brain survival might also say that one hemisphere is not enough for the brain to survive; removing half of the brain, they might say, is not a "small change." So they might say that you survive Brain Transplant but **not** Hemisphere Transplant or My Division.

On the other hand: it really does seem like you could survive Brain Transplant. And if you could survive Brain Transplant, it looks like you could survive Hemisphere Transplant!

Say that you survive
neither Hemisphere
Transplant nor My
Division.

Neither of these options is open to the proponent of psychological survival. It is very hard to deny that the result of Hemisphere Transplant is a person who stands in the right kind of psychological relations to you.

So let's consider the second kind of response to My Division.

Try to find some relevant difference between Hemisphere Transplant and My Division which could explain why you survive one but not the other.


So let's consider the second kind of response to My Division.

What could the relevant difference be? One difference is that in the case of Hemisphere Transplant (but not My Division) the psychological connection between the individuals is a **non-branching connection**. So some psychological theorists say that for you to survive at some later time, you must stand in a certain psychological connection to that thing and **to nothing else**.

This would explain why you can survive Hemisphere Transplant but not My Division.


What could the relevant difference be? One difference is that in the case of Hemisphere Transplant (but not My Division) the psychological connection between the individuals is a **non-branching connection**. So some psychological theorists say that for you to survive at some later time, you must stand in a certain psychological connection to that thing and **to nothing else**.

This would explain why you can survive Hemisphere Transplant but not My Division.



What could the relevant difference be? One difference is that in the case of Hemisphere Transplant (but not My Division) the psychological connection between the individuals is a **non-branching connection**. So some psychological theorists say that for you to survive at some later time, you must stand in a certain psychological connection to that thing and **to nothing else**.

This would explain why you can survive Hemisphere Transplant but not My Division.



This would explain why you can survive Hemisphere Transplant but not My Division.

This also gives the psychological theorist the resources to resist the Teletransportation Argument. On this kind of view, premises (1) and (2) are false, because the psychological connections to Marsy and Venusy are branching connections.

Notice that a materialist could also adopt this kind of view. They could say that what is required for your survival is a non-branching material relation to some future thing.

But here is an odd consequence of this view. Suppose that you wake up, and are told that you are the result of a hemisphere transplant operation. Next to you on another hospital bed is a body into whom the other hemisphere was transplanted.

You do not yet know whether that body will awake.

You think that you know that you are the person you remember being. But, if that other body survives to become a living person, you will not be the person you remember being.

But how could your identity — who you are — depend on what happens with that other body?

But how could your identity — who you are — depend on what happens with that other body?

This is an argument against the idea that we can save the psychological theory of survival by adding on a non-branching clause. Intuitively, the idea is that whether you are the same person as X depends **only** on relations between you and X — it can't possibly depend on relations to some other thing!

We can turn this into an argument against the psychological theory:

THE NO-NON-BRANCHING ARGUMENT

- (1) If some version of the psychological theory of survival is true, it must be a “non-branching” theory.
- (2) Whether I am the same person as X depends only on relations between me and X.
- (3) No “non-branching” theory can be true. (2)
-
- (C) The psychological theory of survival is false.

THE NO-NON-BRANCHING ARGUMENT

- (1) If some version of the psychological theory of survival is true, it must be a “non-branching” theory.
 - (2) Whether I am the same person as X depends only on relations between me and X.
 - (3) No “non-branching” theory can be true. (2)
-
- (C) The psychological theory of survival is false.

We've seen that My Division can be used to make trouble for materialist and psychological theories of survival.

What should the dualist say about this kind of case?

What should the dualist say about this kind of case?

At first glance, the dualist might seem to have a much easier time than other views. After all, according to the dualist, survival is a matter of the survival of an immaterial soul. And immaterial souls (unlike brains) cannot be split into two.

But My Division can still be used to pose a challenge for the dualist. It seems that the dualist has two options:

In My Division, the psychophysical laws imply that your soul gets connected to neither of the resulting bodies, so that neither Lefty nor Righty is you.

In My Division, the psychophysical laws imply that your soul gets connected to one of Lefty or Righty (but not both).

In My Division, the psychophysical laws imply that your soul gets connected to neither of the resulting bodies, so that neither Lefty nor Righty is you.

But neither looks especially appealing.

On the view that neither Lefty nor Righty is you, we face the same challenge that the materialist faced: the challenge of explaining why you survive Hemisphere Transplant but not My Division.

And the options for responding to this challenge are the same. We can deny that you survive Brain Transplant or Hemisphere Transplant — but that seems implausible.

Or we can say that neither Lefty nor Righty is you because the physical connection is branching. But that is open to the challenges faced by all “non-branching” theories.

In My Division, the psychophysical laws imply that your soul gets connected to one of Lefty or Righty (but not both).

Could we say instead that in this case one of Lefty or Righty would be you? Perhaps in this kind of case the psychophysical laws randomly connect the soul to one of the two hemispheres.

But suppose that you are Lefty, and you are wondering whether you are the same person as the person whose body was fatally injured. You know that Righty is wondering the same thing. At most one of you is the person you seem to remember being. Will you ever be able to discover who really is that person?

But suppose that you are Lefty, and you are wondering whether you are the same person as the person whose body was fatally injured. You know that Righty is wondering the same thing. At most one of you is the person you seem to remember being. Will you ever be able to discover who really is that person?

It seems that you will never be able to figure out which one of you really is that person. And yet this fact would seem to be a fact of great importance. Suppose that the person whose body was destroyed was married; wouldn't his spouse want to know which person they are married to?

In this scenario, there is some temptation to believe that there cannot be a very important unknown fact about who Lefty and Righty are. It seems that once we know all of the facts about the physical and psychological relations between Lefty and Right and the person whose body was destroyed, we know all of the important facts. But that would seem to leave no role for immaterial souls.

We have seen that My Division can be used to present challenges for all of our theories of survival. It can also be used to challenge a fundamental assumption of our discussion so far: the assumption that whether we survive a given event is something which should be of great importance to us.

Consider again My Division. We have seen that a plausible case can be made that you are neither Lefty nor Righty, and so that the following is true:

You do not survive
My Division.

Yes, the dualist can say that you survive as an immaterial soul — just one that is no longer connected to any body. But even on this view My Division would mean your death (in the ordinary sense of that term).

By contrast, we have seen that the following is quite plausible:

You do survive Hemisphere
Transplant.

You do not survive
My Division.

You do survive Hemisphere
Transplant.

Ordinarily, if given a choice between two surgeries, one of which you might survive, and one of which you definitely won't survive, the choice would be clear.

But are matters so clear in this case? Suppose that your body has been destroyed and you have a choice to have one of your hemispheres transplanted or both. (For some reason a full brain transplant is impossible in this case.)

You know that there is a 50% chance of success for any given hemisphere transplant operation.

I think that many people would choose to have them try to transplant both hemispheres. But it is not clear why, if survival is what matters — the chance of survival in both cases is 50% (since if both are successful in the case where you choose two transplant attempts, you don't survive).

But are matters so clear in this case? Suppose that your body has been destroyed and you have a choice to have one of your hemispheres transplanted or both. (For some reason a full brain transplant is impossible in this case.) You know that there is a 50% chance of success for any given hemisphere transplant operation.

I think that many people would choose to have them try to transplant both hemispheres. But it is not clear why, if survival is what matters — the chance of survival in both cases is 50% (since if both are successful in the case where you choose two transplant attempts, you don't survive).

If that is true, that's some evidence that — at least in this case — we care more about having someone psychologically connected to us survive than we do about someone numerically identical to us surviving.

My Division is a puzzling case along a number of different dimensions. A different kind of puzzling case can be used to cast doubt on another fundamental assumption about survival: the assumption that there must always be a fact of the matter about whether some future person is me.

We can introduce this kind of case by example. Suppose that I am an impoverished philosophy professor, and definitely not rich.

Now suppose that a wealthy benefactor who loves philosophy decides to give me some money. But he does this in an eccentric way: by adding 1 cent to my bank account every second.

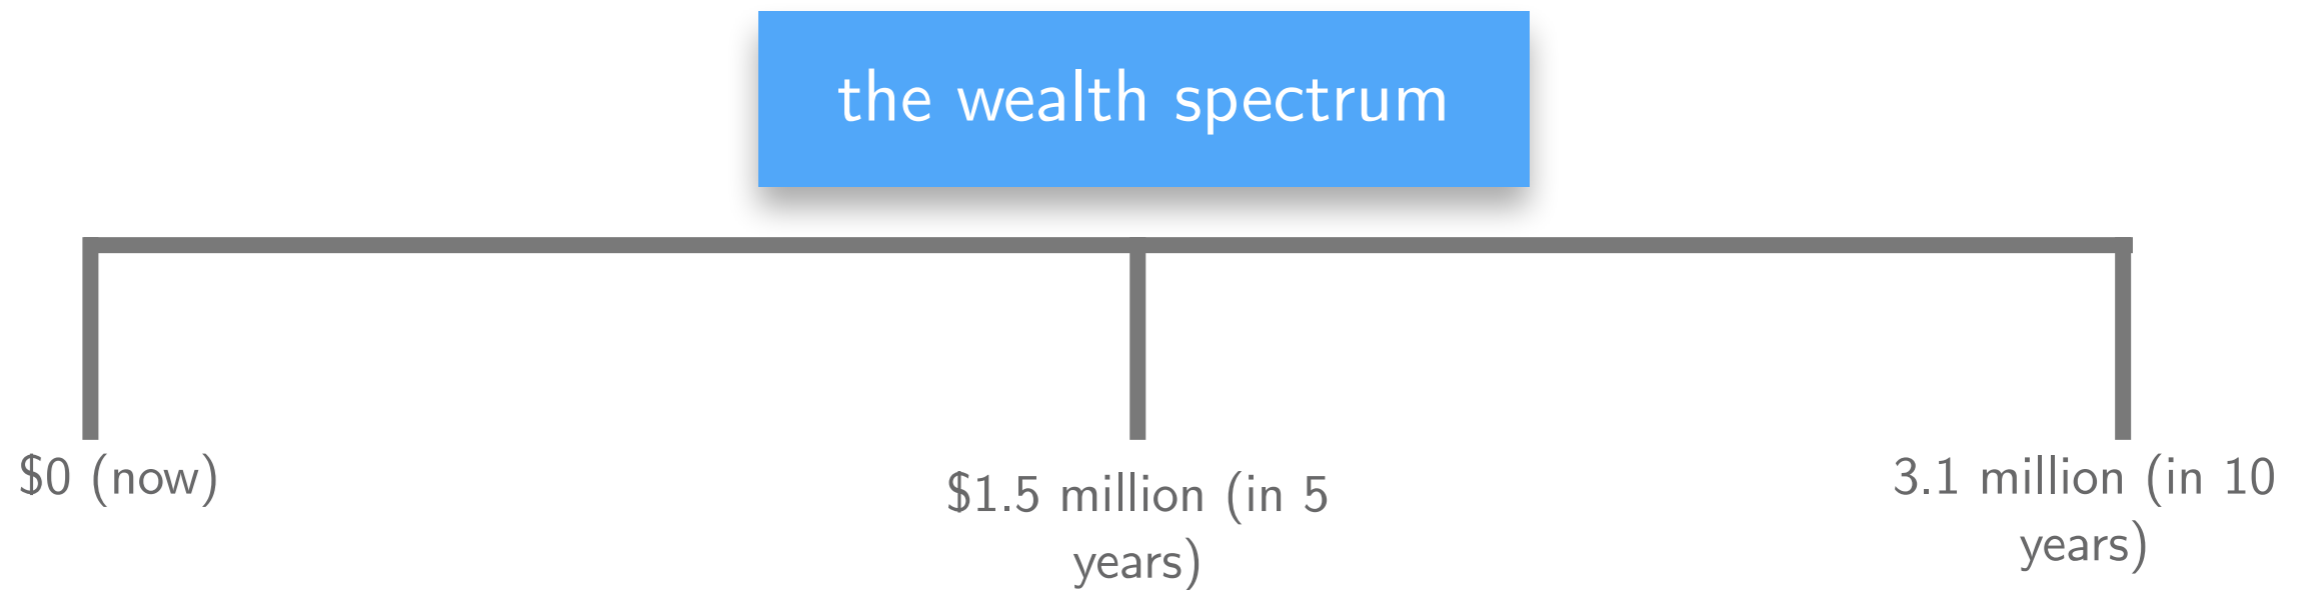
At the end of 10 years, I will have \$3.1 million in my bank account, and will be rich.

We can chart my progress using the wealth spectrum.

Now suppose that a wealthy benefactor who loves philosophy decides to give me some money. But he does this in an eccentric way: by adding 1 cent to my bank account every second.

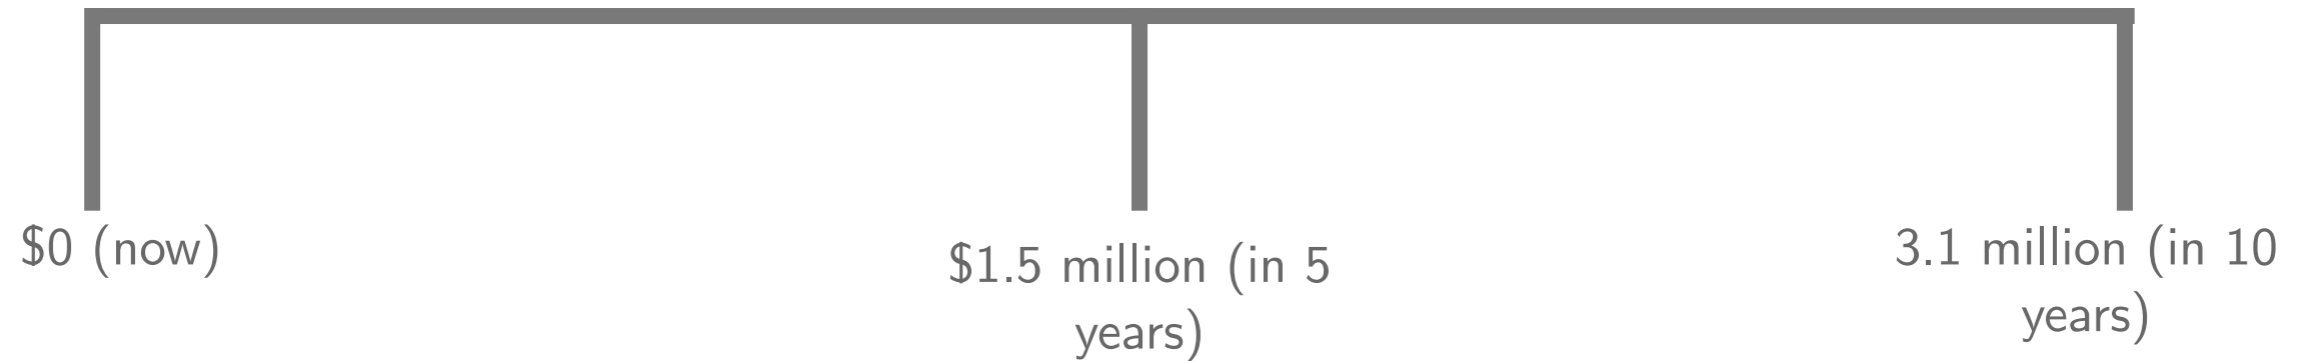
At the end of 10 years, I will have \$3.1 million in my bank account, and will be rich.

We can chart my progress using the wealth spectrum.



When did I become rich?

the wealth spectrum



When did I become rich?

It appears that there are exactly three ways to answer this question.

Sharp Cut Off

There is a precise point in the spectrum at which I switched from being non-rich to being rich.

Indeterminacy

At the beginning I was non-rich; at the end I am rich; but there is no sharp cut off. Instead, there is a range of cases in which it is not determinately true either that I am rich or that I am non-rich.

Never rich

Even at the end of the spectrum, I am still not rich.

the wealth spectrum

\$0 (now)

\$1.5 million (in 5
years)

3.1 million (in 10
years)

Sharp Cut Off

There is a precise point in the spectrum at which I switched from being non-rich to being rich.

Indeterminacy

At the beginning I was non-rich; at the end I am rich; but there is no sharp cut off. Instead, there is a range of cases in which it is not determinately true either that I am rich or that I am non-rich.

Never rich

Even at the end of the spectrum, I am still not rich.

Which answer is most plausible in the case of the wealth spectrum?


Let's now look at a different "spectrum" example, more relevant to our present purposes. (This is also due to Derek Parfit.)

“At the near end of this spectrum is the normal case in which a future person would be fully continuous with me as I am now, both physically and psychologically. This person would be me in just the way that, in my actual life, it will be me who wakes up tomorrow. At the far end of this spectrum the resulting person would have no continuity with me as I am now, either physically or psychologically. In this case the scientists would destroy my brain and body, and then create, out of new organic matter, a perfect Replica of someone else. Let us suppose this person to be Greta Garbo. We can suppose that, when Garbo was 30, a group of scientists recorded the states of all the cells in her brain and body.”

In the intermediate stages, the person is to some degree physically like you and to some degree physically like Garbo, and to some degree psychologically like you and to some degree psychologically like Garbo. Let's call this **the survival spectrum.**

In the intermediate stages, the person is to some degree physically like you and to some degree physically like Garbo, and to some degree psychologically like you and to some degree psychologically like Garbo.

the survival spectrum



We again have just three choices.

the survival spectrum

0% of your memories, personality traits, etc. are changed, and 0% of the material composing your body is replaced

50% of your memories, personality traits, etc. are changed, and 50% of the material composing your body is replaced

100% of your memories, personality traits, etc. are changed, and 100% of the material composing your body is replaced

Sharp Cut Off

There is a precise point in the spectrum at which, for the first time, I would not survive the surgery. Perhaps it is when 43.13% of your psychological traits have changed and the same percentage of the matter composing the organism is replaced.

Indeterminacy

In the first cases I survive; in the last cases I do not survive; but there is no sharp cut off. Instead, there is a range of cases in which it is not determinately true either that the person is me or that the person is not me.

Survive All

Even in the cases at the right edge of the spectrum, I survive.

We can all agree that Survive All looks pretty implausible. If your body were destroyed and replaced with a completely different body, you would not survive.

Sharp Cut Off

There is a precise point in the spectrum at which, for the first time, I would not survive the surgery. Perhaps it is when 43.13% of your psychological traits have changed and the same percentage of the matter composing the organism is replaced.

Indeterminacy

In the first cases I survive; in the last cases I do not survive; but there is no sharp cut off. Instead, there is a range of cases in which it is not determinately true either that the person is me or that the person is not me.

We can all agree that Survive All looks pretty implausible. If your body were destroyed and replaced with a completely different body, you would not survive.

Here is an argument against Sharp Cut Off. If Sharp Cut Off were true, then there are two adjacent procedures on the combined spectrum which are such that I should care an enormous amount which procedure happens to me. (After all, I would survive one but not the other.) But in reality it would never be rational to care which of two such similar procedures I should undergo.

Sharp Cut Off

There is a precise point in the spectrum at which, for the first time, I would not survive the surgery. Perhaps it is when 43.13% of your psychological traits have changed and the same percentage of the matter composing the organism is replaced.

Indeterminacy

In the first cases I survive; in the last cases I do not survive; but there is no sharp cut off. Instead, there is a range of cases in which it is not determinately true either that the person is me or that the person is not me.

Does Sharp Cut Off look more plausible if one is a dualist? Couldn't one then say that there is a point in the combined spectrum at which the soul would lose its connection to the body, and that this would explain the existence of a cut off point?

But even here there are puzzles. Suppose that you underwent one of the procedures in the middle of the combined spectrum. Could you tell afterwards whether you had survived?

And what should the dualist say about cases to the right of the cut off point (wherever that is) -- is a new soul created, or joined to the body for the first time, by the procedure?

Parfit thinks that the moral of the survival spectrum is not that the psychological theory is false, but that we should change a fundamental part of our view about what our own continued existence amounts to.

“[One] assumes that, in each of these cases, the resulting person either would or would not be me. This is not so. The resulting person would be me in the first few cases. In the last case he would not be me. In many of the intervening cases, neither answer would be true. I can always ask, ‘Am I about to die? Will there be some person living who will be me?’ But, in the cases in the middle of this Spectrum, there is no answer to this question.”

If this is right, then sometimes the answer to the question “Is that future person me?” is neither “Yes” nor “No” but “sort of.”

Could this be true?

If it could not be true, that can be used to construct an argument for soul survival. The common thread between materialist and psychological theories of survival is that they explain survival in terms of something **complex** — a material thing with many parts in one case, and a large collection of memories and psychological traits in the other.

It seems plausible that something like the survival spectrum can be constructed for any theory of survival which explains survival in terms of a complex thing.

What theory of survival doesn't do this? It seems that soul survival is the only one.

It seems plausible that something like the survival spectrum can be constructed for any theory of survival which explains survival in terms of a complex thing.

What theory of survival doesn't do this? It seems that soul survival is the only one.

That suggests the following argument:

THE NO-INDETERMINATE SURVIVAL ARGUMENT

- (1) It can never be indeterminate whether someone is me.
 - (2) If survival depended on the existence of something complex, it would sometimes be indeterminate whether someone is me.
 - (3) Survival cannot depend on the existence of something complex. (1,2)
 - (4) If survival cannot depend on the existence of something complex, soul survival is true.
-
- (C) Soul survival is true. (3,4)