UPLOADING & LIFE AFTER DEATH

is life after death possible?

arguments for immortality

technological immortality

should I fear death?

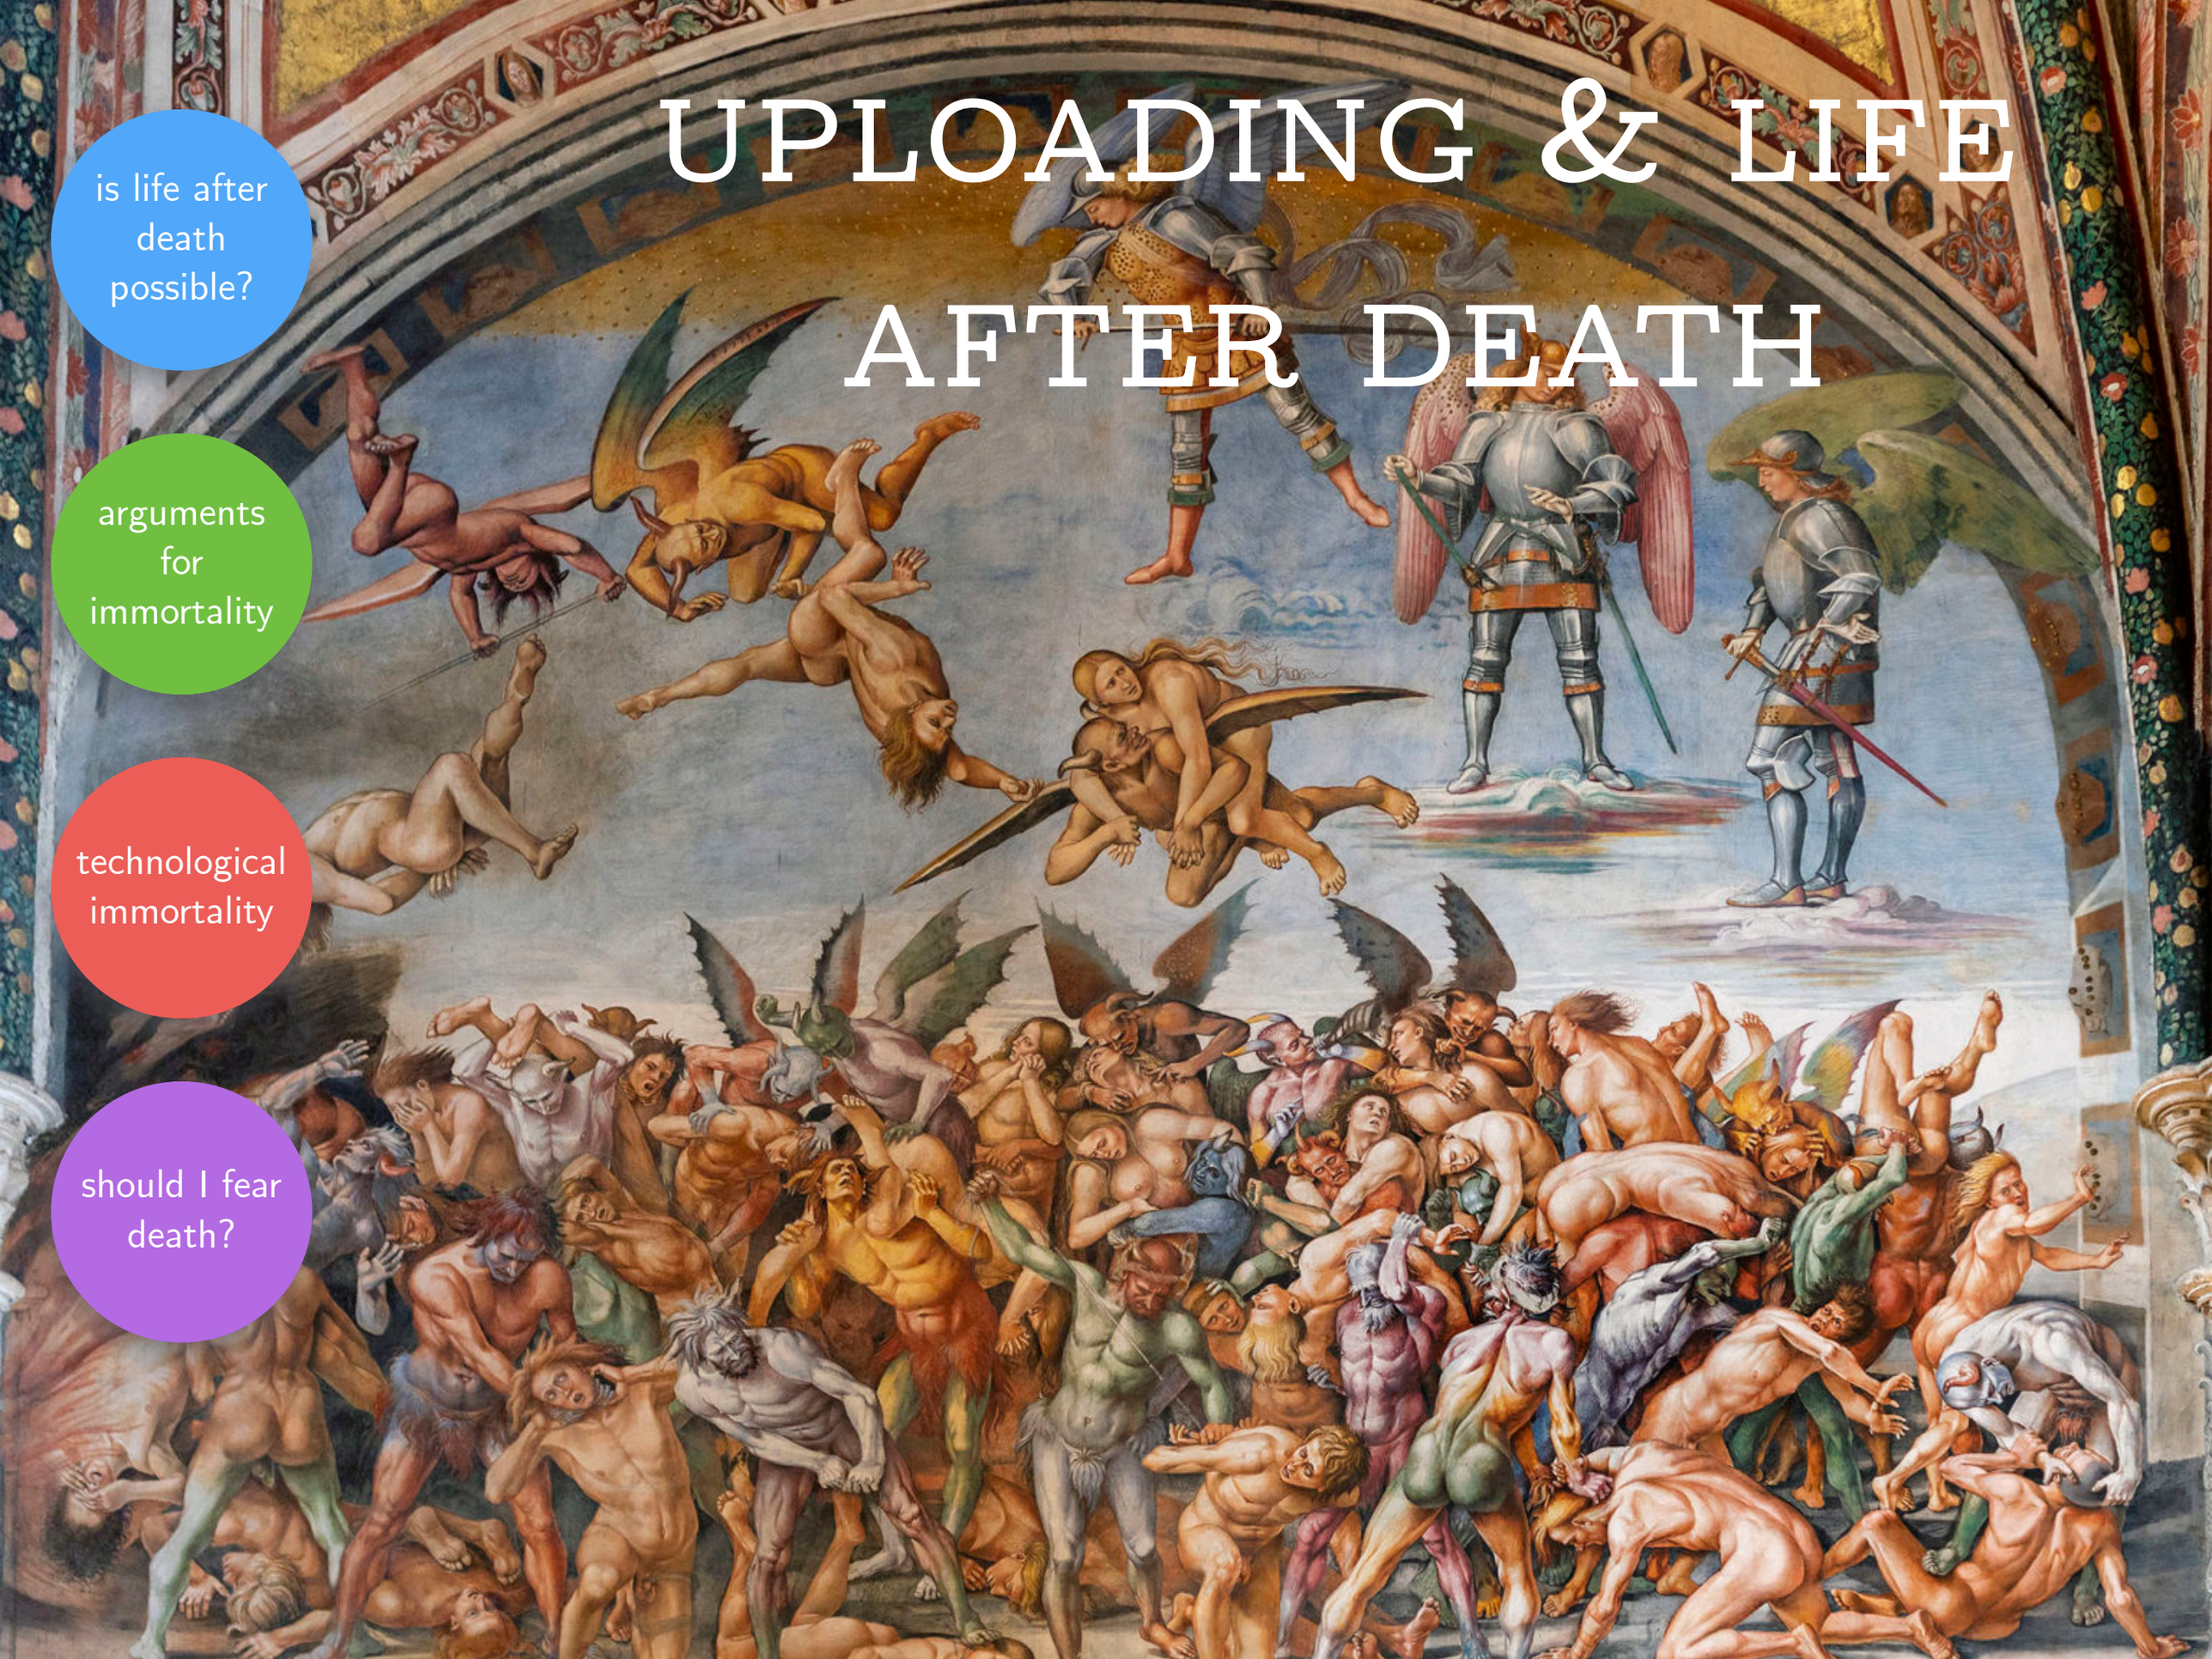

Our topic today is one on which we've already touched several times this semester.

Today we will discuss four different philosophical questions about death, and life after death:

Is life after death possible?

Can we prove that there is life after death?

Could technology allow me to indefinitely delay death?

If there is no life after death, should I fear death?

Is life after death possible?

For there to be life after death would be for someone to exist after your death who is **you**. A first question to ask about life after death is: is it possible?

The answer to this question clearly depends on our answer to the survival question: it depends on what it would take for some later individual to be you.

On two of our answers to the survival question, life after death appears to be clearly possible.

There seems no reason why your immaterial soul could not exist after the death of your body; so, if soul survival is true, life after death seems possible.

Similarly, there seems no reason why, after your death, it would be impossible for someone to exist who has your memories and stands in the right psychological relations to you. So, if psychological survival is true, it looks like life after death is possible.

Matters are less clear when it comes to our answers to the survival question which involve a materialist component.

These views differ in important ways, but they all agree that your survival depends on the survival of a certain physical thing.

Views of this kind face an obvious problem in making sense of the possibility of life after death. After death bodies decay (or are cremated). Doesn't that just show that the relevant physical objects won't survive your death — and so that, if one of the above views is right, you won't?

Views of this kind face an obvious problem in making sense of the possibility of life after death. After death bodies decay (or are cremated). Doesn't that just show that the relevant physical objects won't survive your death — and so that, if one of the above views is right, you won't?

This is a problem which has long troubled philosophers in the Christian tradition.

This might sound surprising. After all, don't Christians believe that one's immaterial soul is enough for life after death?

In fact, this is not the traditional Christian view. The following quote from Aquinas is representative:

In fact, this is not the traditional Christian view. The following quote from Aquinas is representative:

> "The necessity of holding the resurrection arises from this — that man may obtain the last end for which he was made. This cannot be accomplished in this life, nor in the life of the separated soul ... it is necessary for the selfsame man to rise again; and this is effected by the selfsame soul being united to the selfsame body. For otherwise there would be no resurrection properly speaking, if the same man were not reformed."

This is also implied by the use of the word "resurrection" — if life after death were simply the continued existence of an immaterial soul, then there would be nothing needing to be "raised up."

This is also implied by the use of the word "resurrection" — if life after death were simply the continued existence of an immaterial soul, then there would be nothing needing to be "raised up."

How can we make sense of this idea, given obvious facts about the decay of bodies and brains after death?

One idea (defended by Aquinas and others) is that at the time of resurrection our bodies could be **reassembled** by God. While it is true that our bodies decay after death, the material of which they were composed does not cease to exist. Surely an omnipotent being could then simply reassemble our bodies out of the matter from which they were composed when we were living.

Would this be enough for God to bring our bodies back into existence?

Would this be enough for God to bring our bodies back into existence?

This view of resurrection as reassembly seems to rely on the following thesis:

```
If x and y are material things,
   and x and y have the same
        parts, then x=y.
```

After all, if this were not true, there would be no reason to think that the reassembled body is in fact **my** body.

The problem is that this thesis called into question by an ancient paradox: the paradox of the Ship of Theseus.

The problem is that this thesis called into question by an ancient paradox: the paradox of the Ship of Theseus.

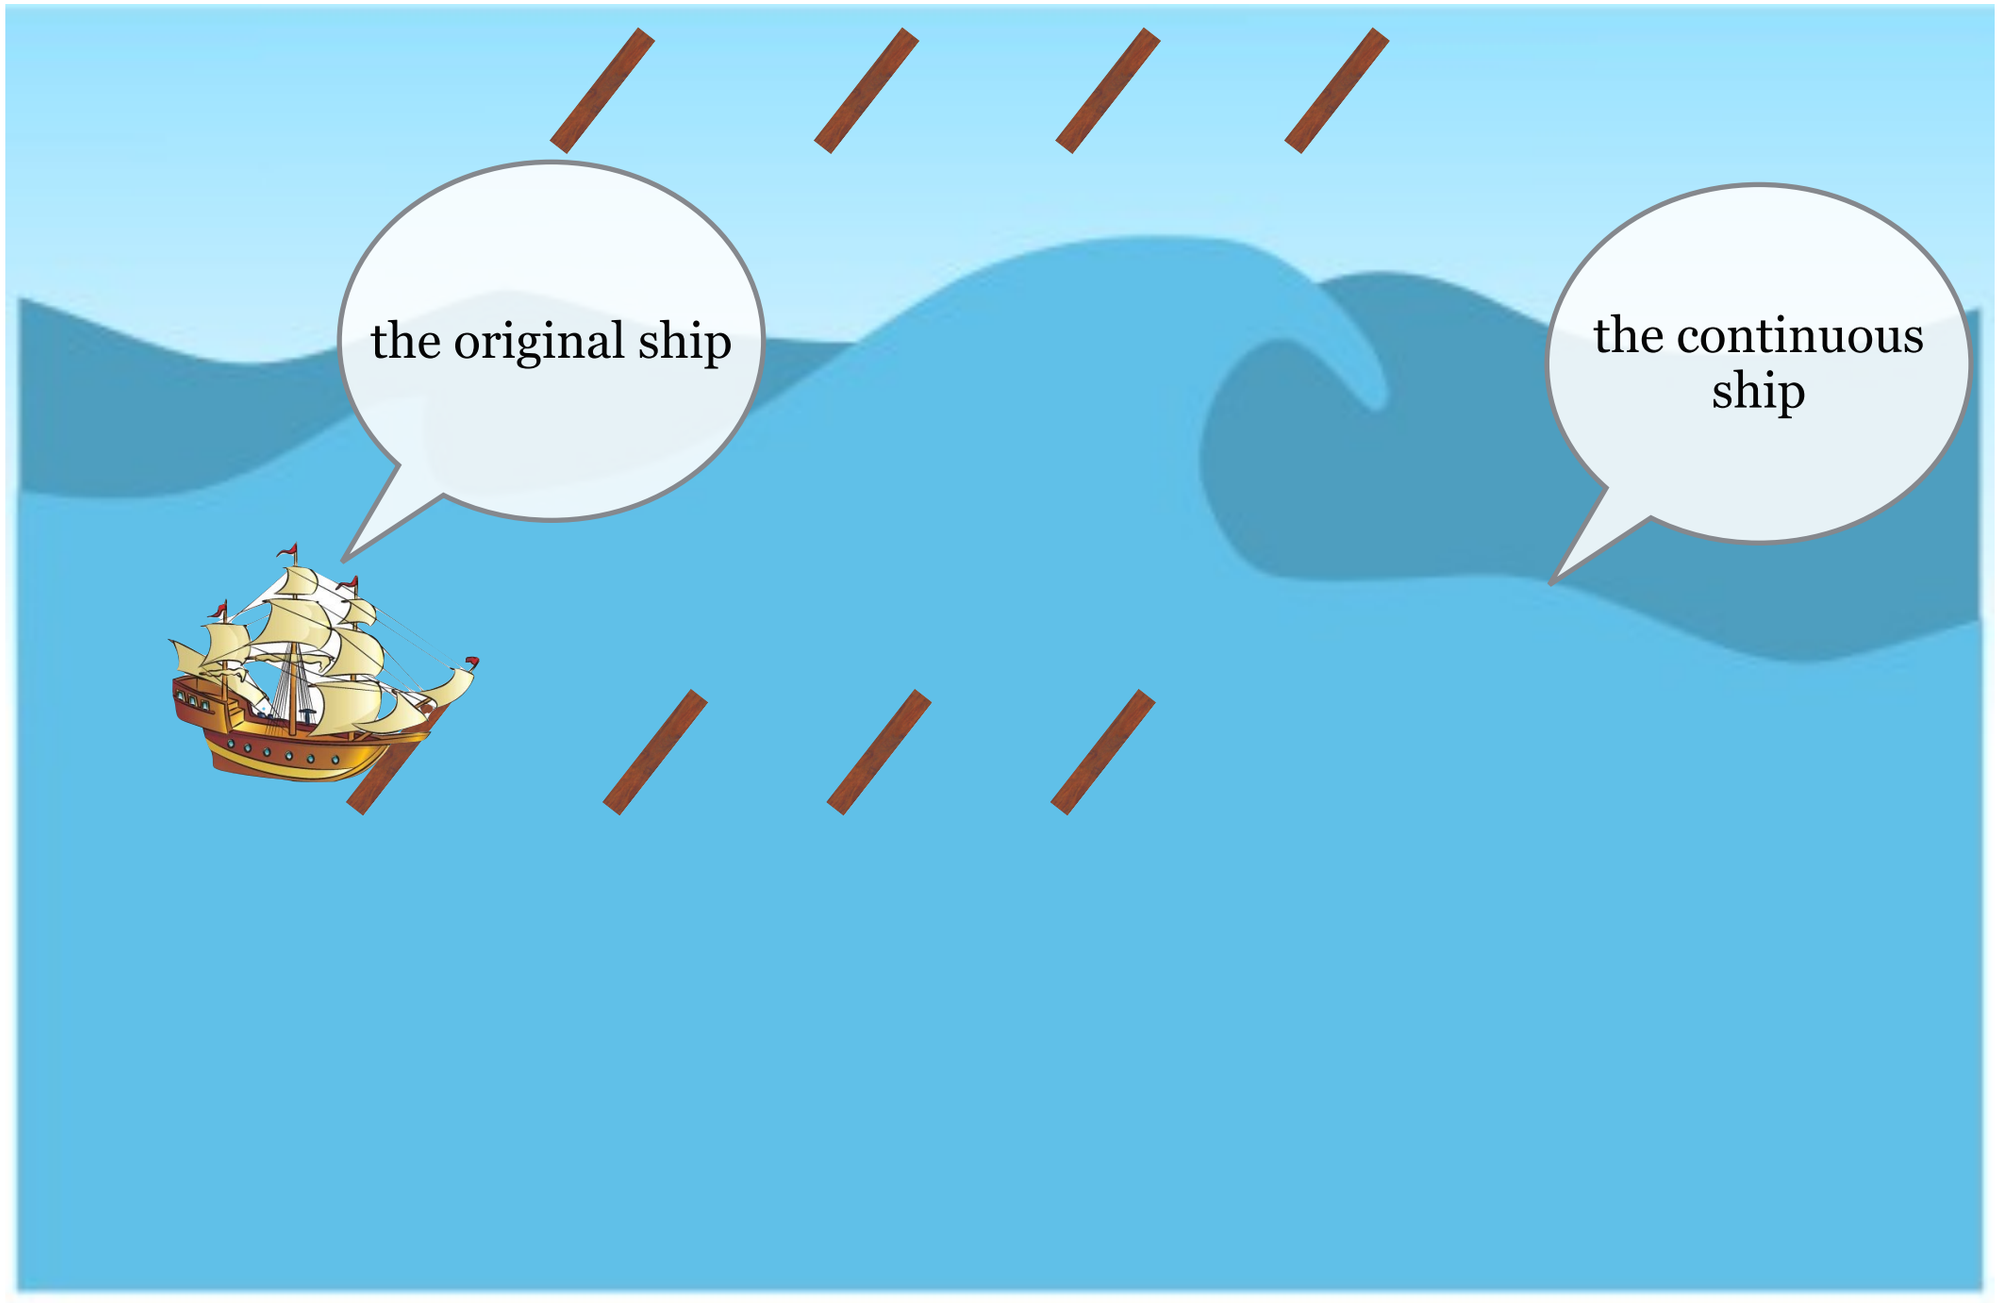

It seems plausible that the following claim is true:

```
Original Ship = Continuous Ship
```

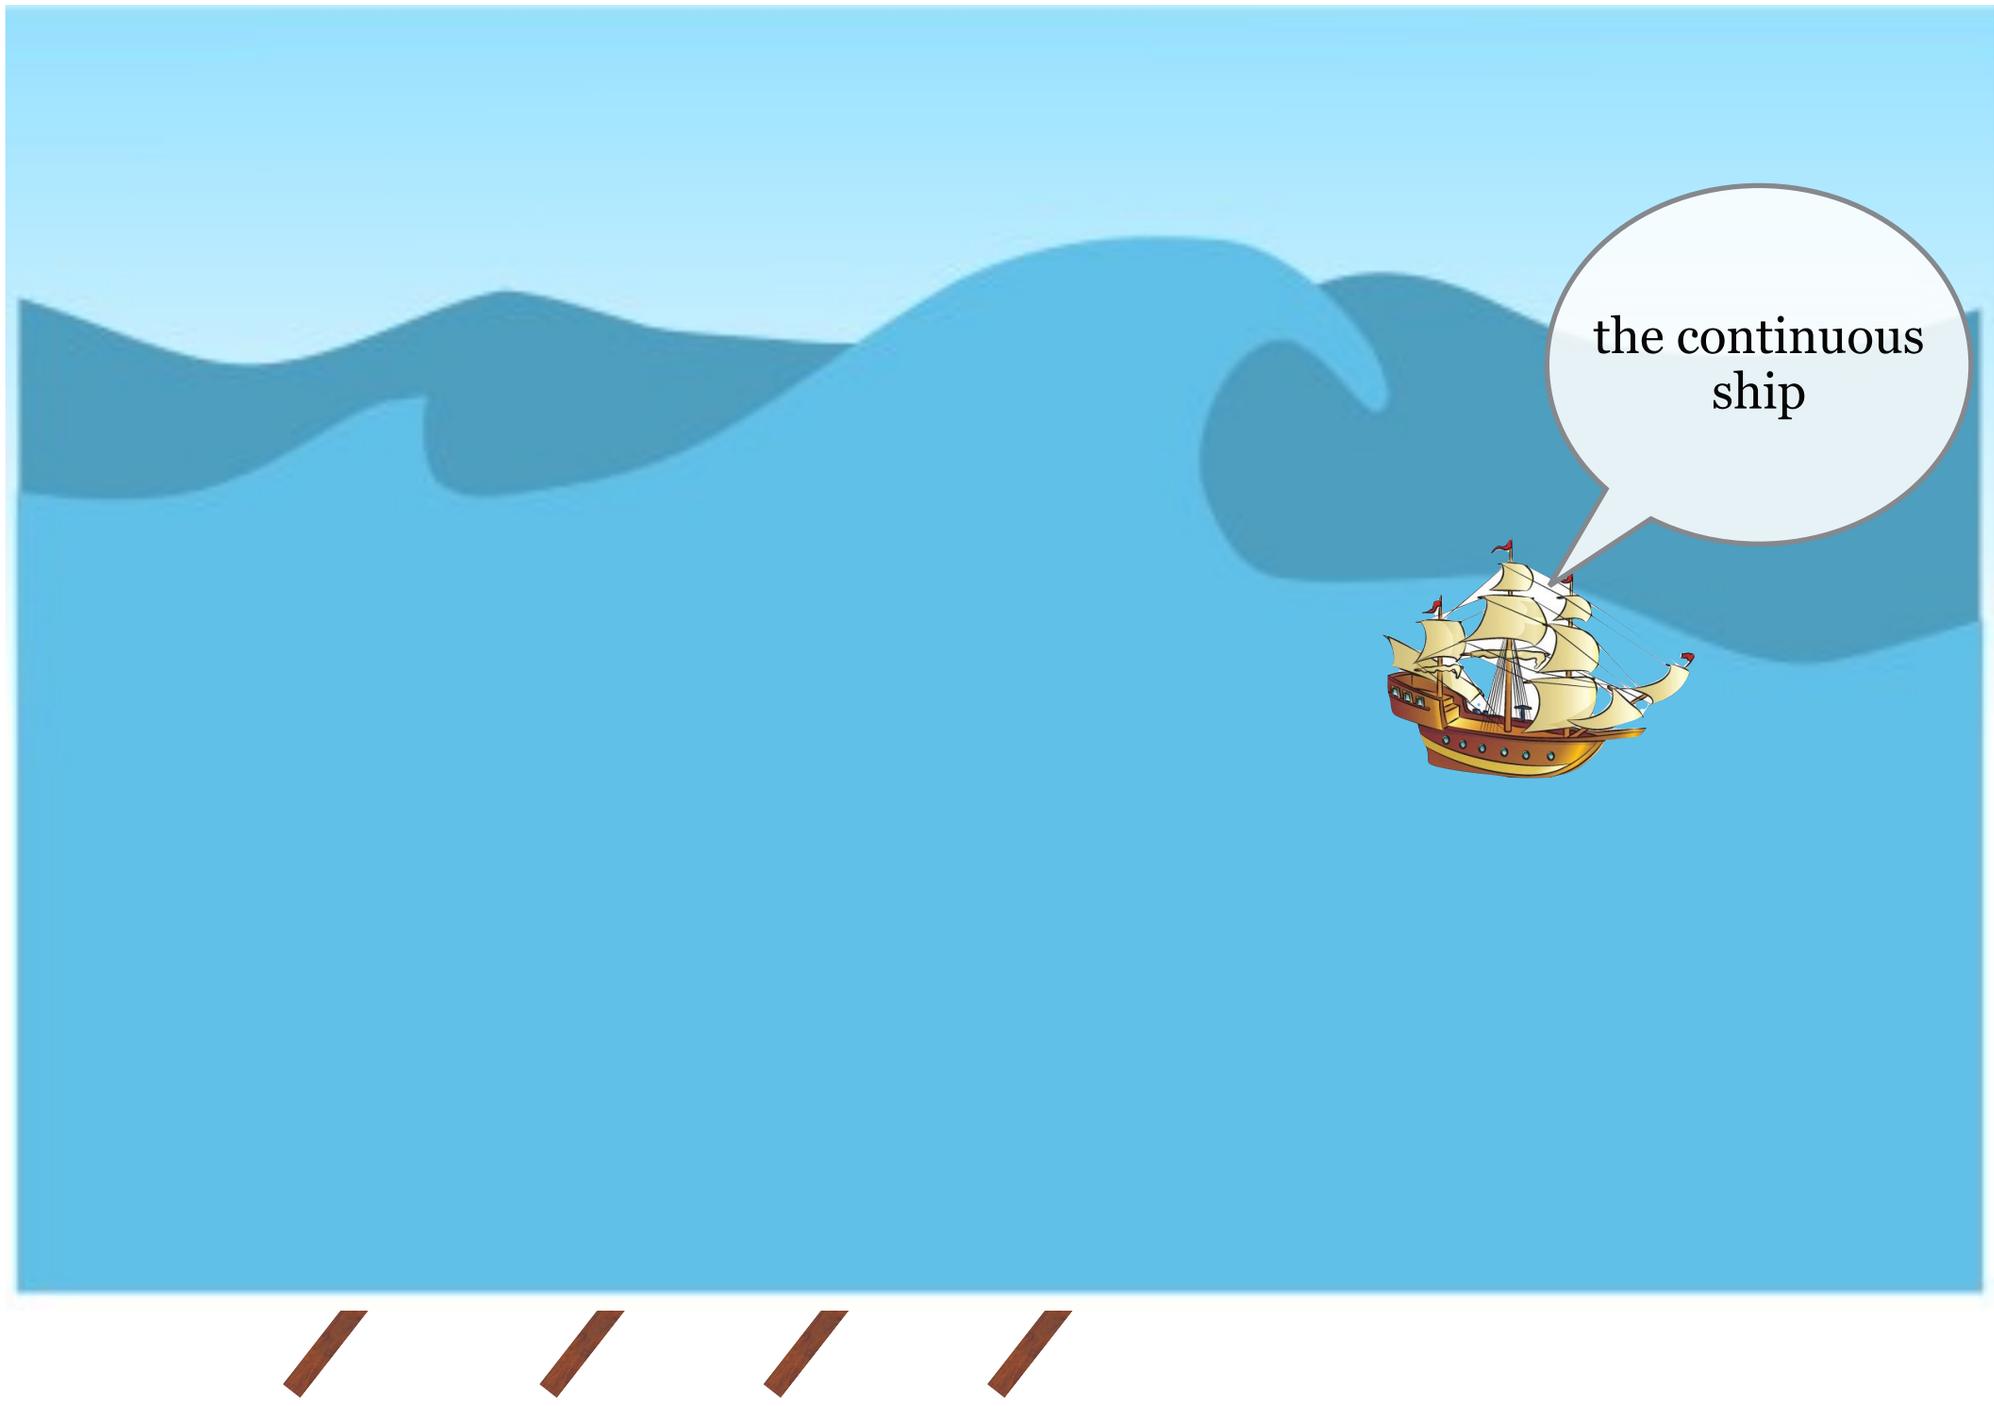

Our previous discussion strongly suggests that the following claim is true:

```
Original Ship = Continuous Ship
```

But now imagine that some enterprising person gets the idea to rebuild the original Ship of Theseus from the wooden planks which have, over time, been replaced.

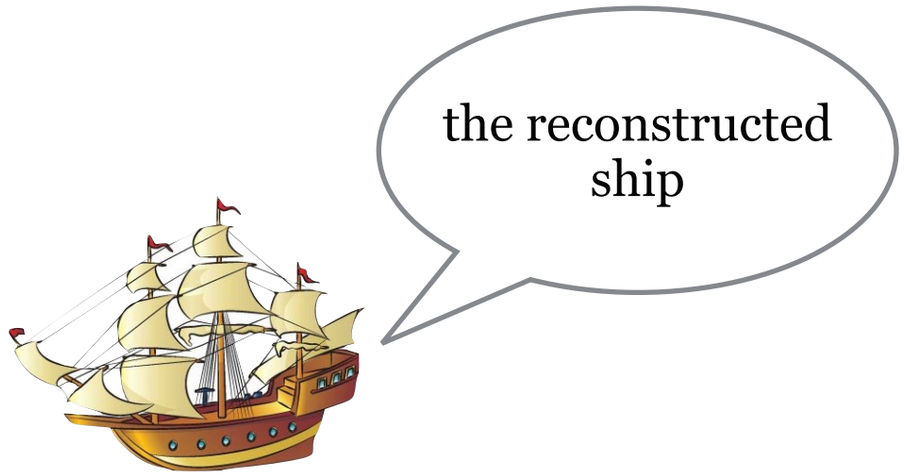

The following now seems plausible:

```
Original Ship = Reconstructed Ship
```

After all, Original Ship and Reconstructed Ship are made of exactly the same materials organized in exactly the same way!

Original Ship = Continuous Ship

Original Ship = Reconstructed Ship

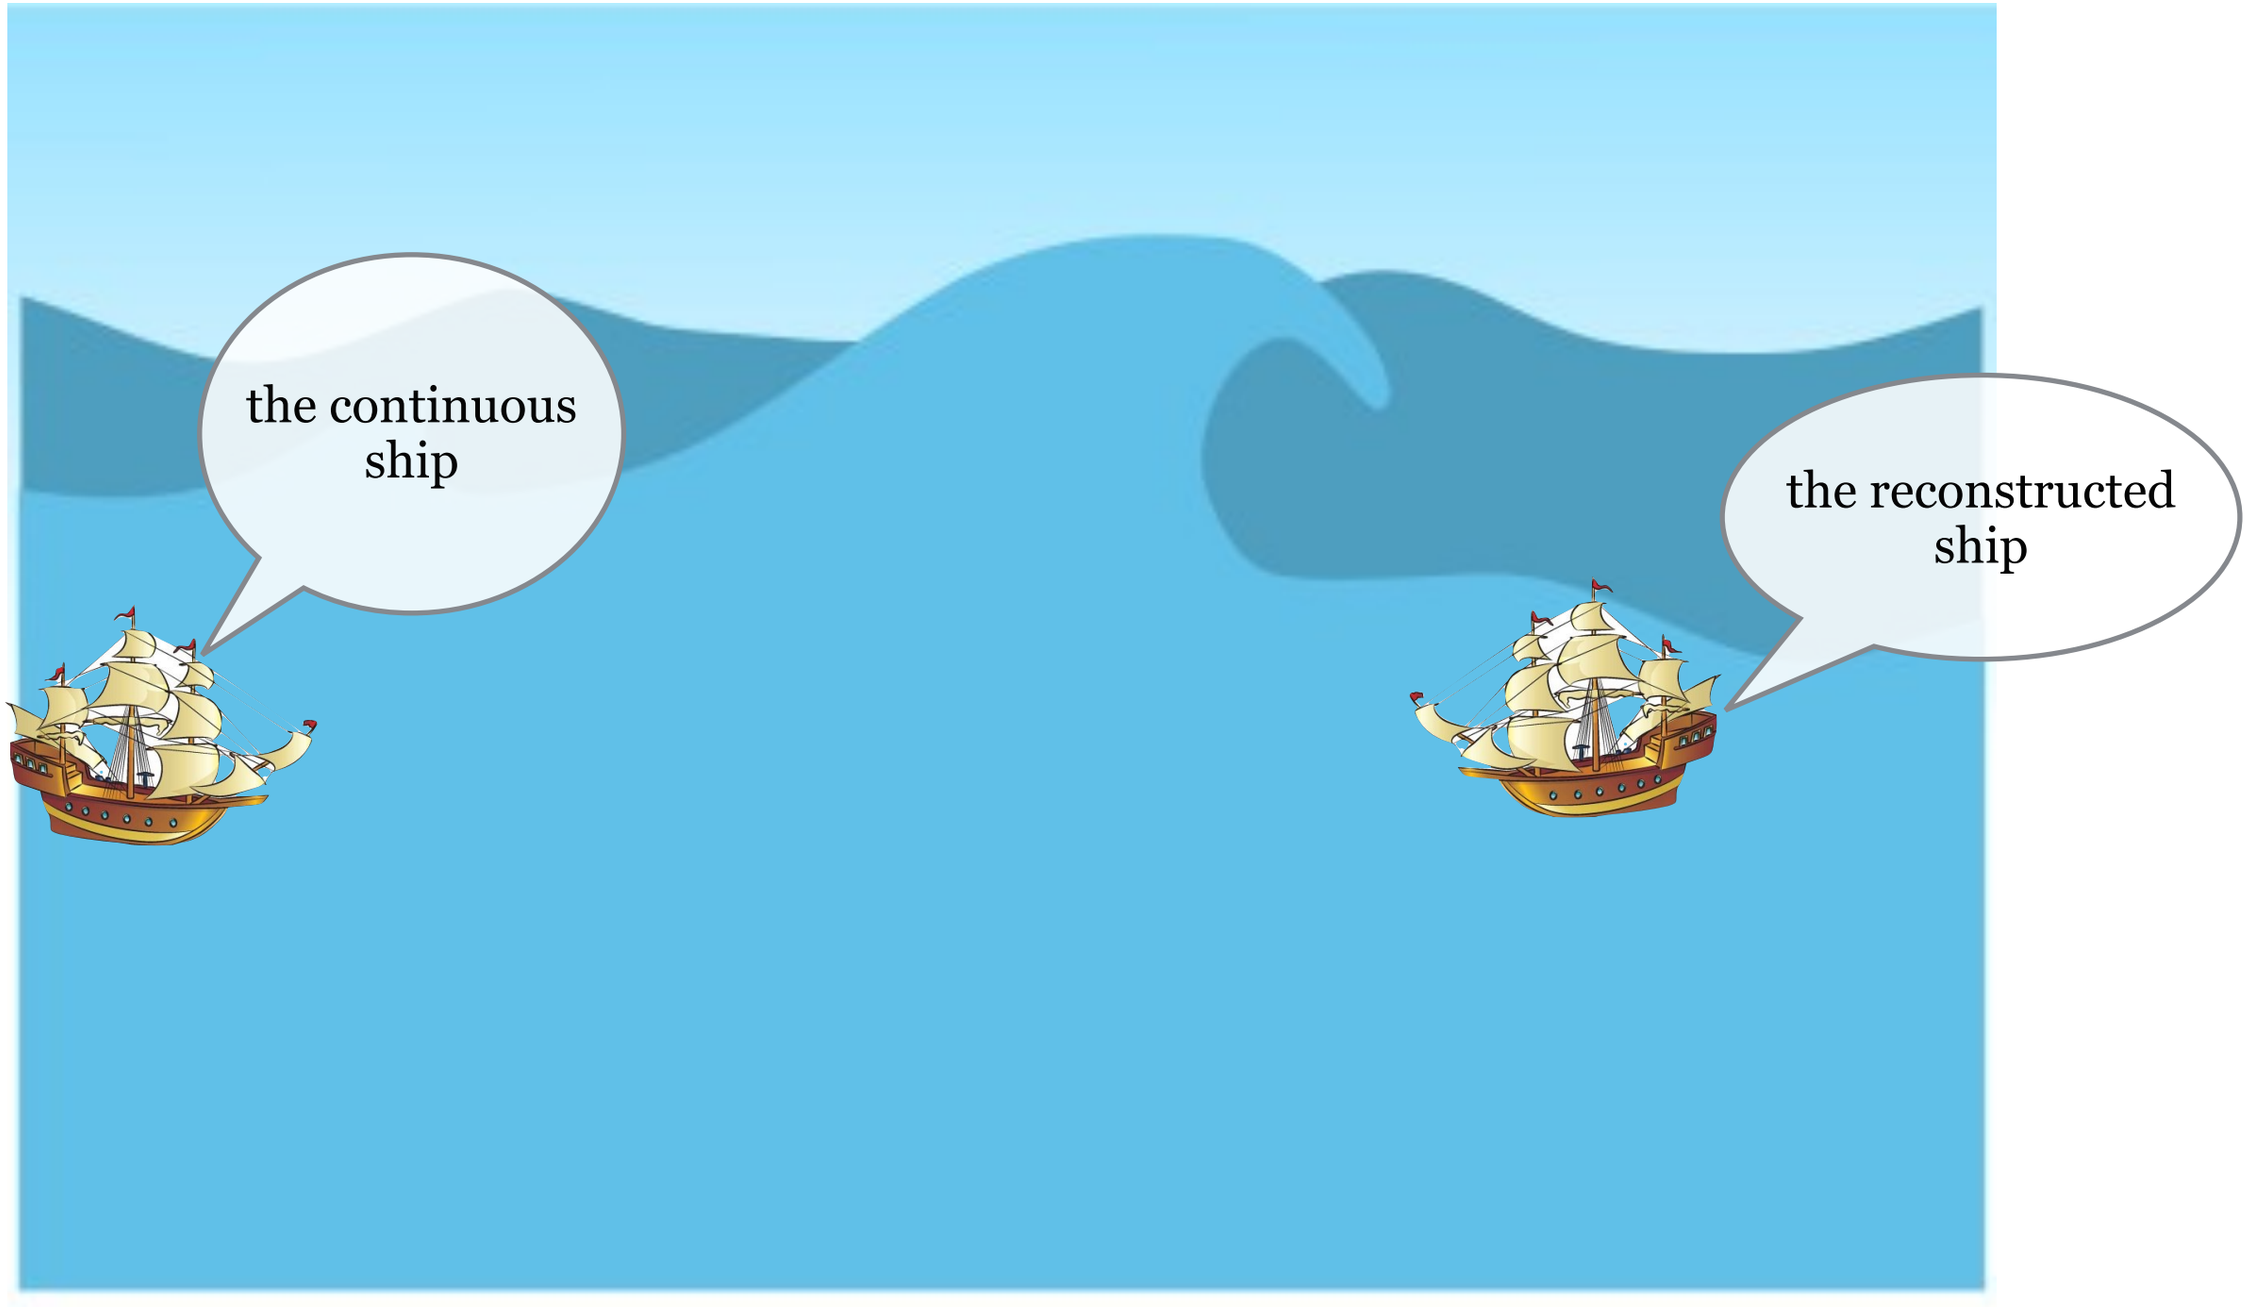

But suppose that we take our reconstructed ship for a cruise.

$$\text{Original Ship} = \text{Continuous Ship} \qquad \text{Original Ship} = \text{Reconstructed Ship}$$

This is not a story of a ship crashing into itself; so it seems fairly clear that:

$$\text{Continuous Ship} \neq \text{Reconstructed Ship}$$

The problem, though, is that these three claims are inconsistent. This is due to the transitivity of identity: if A=B, and B=C, then A=C.

What's the best way out of this paradox?

What's the best way out of this paradox?

One natural thought is that we should reject the claim that Original Ship is the same as Reconstructed Ship. On this view, if you find all of the parts that composed some thing, and put them back together, that is not enough to reconstitute the thing. Rather, on this view, material objects survive via a series of causal connections over time, perhaps with the requirement that only relatively small changes at one time are possible.

But this raises an obvious problem for the idea of resurrection as reassembly. It seems to imply that, even if God were to reassemble the particles which composed my body, the thing which resulted would not be my body. And it follows from this — plus the assumption that one of our materialist theories is true — that I would not exist either.

How might the proponent of resurrection as reassembly reply?

How might the proponent of resurrection as reassembly reply?

One idea would be to deny that Original Ship = Continuous Ship, on the grounds that they have different parts.

But this leads to a dilemma.

If material things can't gain and lose parts, we only exist for a very short time.

If they can gain and lose parts, then reassembly is not a model for how life after death might be possible.

Can the materialist provide some other model of how life after death might be possible?

It is fair to say that there is no generally accepted answer to this question.

But here is a strategy that some Christian materialists have tried. First, we point out that there is no reason why we should be able to figure out how God brings about the resurrection. (It is, after all, supposed to be a miracle.)

Second, we try to give some story — no matter how unlikely — which shows that there would be no impossibility in doing this. Here's an example, from Peter van Inwagen.

Second, we try to give some story — no matter how unlikely — which shows that there would be no impossibility in doing this. Here's an example, from Peter van Inwagen.

"... I proposed a solution to this problem that has, let us say, not won wide assent. ... I suggested that God could accomplish the resurrection of, say, Socrates, in the following way. He could have, in 399 BC, miraculously translated Socrates' fresh corpse to some distant place for safe-keeping (at the same time removing the hemlock and undoing the physiological damage it had done) and have replaced it with a simulacrum, a perfect physical duplicate of Socrates' corpse; later, on the day of resurrection, he could reanimate Socrates' corpse, and the reanimated corpse, no longer a corpse but once more a living organism, would be Socrates."

The idea is not that this is what in fact happens — the aim is the more limited one of showing that there is no impossibility in it happening in this way.

The idea is not that this is what in fact happens — the aim is the more limited one of showing that there is no impossibility in it happening in this way.

If this is conceded, then one might argue that there are probably plenty of other ways — including lots that we can't think of — in which this could take place.

The aim here is not the ambitious one of providing a satisfactory theory of how the resurrection of the dead works. The aim is just to show that, even if survival requires the survival of some physical part of us, life after death is not impossible.

Can we prove that there is life after death?

But of course even if life after death is possible, that doesn't tell us whether or not there is life after death. Let's turn to that question now.

There are a number of arguments for and against life after death that I'll mention only briefly and then set aside.

The first is what might be called the argument from religion. There are as many versions of this argument as there are religions; here is one:

```
1. Christianity is true.
2. If Christianity is true, there is
   life after death.
---------------
C. There is life after death.
```

I'm not setting aside this kind of argument because it is bad. Rather, I am setting it aside because a discussion of the first premise would take us too far afield. We have already discussed arguments relevant to it — the arguments for God's existence, and the argument from evil.

The second argument I am going to mention and then set aside is the argument from near death experiences.

Many people who come very close to death report similar kinds of experiences — a feeling of looking down at one's body, of feeling disembodied, of moving towards a light. One might argue from these experiences as follows:

```
1. People have near death experiences.
2. If there were no life after death, people would
   not have near death experiences.
---------------
C. There is life after death.
```

It is worth noting that this is an argument for life after death, but cannot in any obvious way (unlike the argument from religion) be turned into an argument for immortality.

The key premise here is obviously the second one. A serious assessment of it would have to look at the details of the kinds of near death experiences people report, and consideration of the possible explanations of these experiences.

Instead we will look at the main philosophical argument for life after death.

The first argument has its origins in Plato's *Phaedo*. This is a dialogue which takes place between Socrates and his friends, after Socrates has been sentenced to death for corrupting the youth of Athens.

Socrates is unworried, explaining to his friends that death is nothing to be afraid of; death is just the death of the body, and not the death of him.

He gives a few arguments in favor of this view; the most influential is contained in the following passage.

He gives a few arguments in favor of this view; the most influential is contained in the following passage.

"We ought," said Socrates, "to ask ourselves this: what sort of thing is it that would naturally suffer the fate of being dispersed? For what sort of thing should we fear this fate, and for what should we not? When we have answered this, we should next consider to which class the soul belongs; and then we shall know whether to feel confidence or fear about the fate of our souls."

"Quite true."

"Would you not expect a composite object or a natural compound to be liable to break up where it was put together? And ought not anything which is really incomposite be the one thing of all others which is not affected in this way?"

"We ought," said Socrates, "to ask ourselves this: what sort of thing is it that would naturally suffer the fate of being dispersed? For what sort of thing should we fear this fate, and for what should we not? When we have answered this, we should next consider to which class the soul belongs; and then we shall know whether to feel confidence or fear about the fate of our souls."

"Quite true."

"Would you not expect a composite object or a natural compound to be liable to break up where it was put together? And ought not anything which is really incomposite be the one thing of all others which is not affected in this way?"

Socrates begins by asking what sorts of things can be 'dispersed.' He considers two categories of things: composite things, which have parts, and incomposite things, which are simple and have no parts.

It seems clear that composite things can be dispersed, whereas simple things cannot. Being dispersed, after all, is just a matter of having your parts taken out of connection with each other, and simple things have no parts.

Socrates begins by asking what sorts of things can be 'dispersed.' He considers two categories of things: composite things, which have parts, and incomposite things, which are simple and have no parts.

It seems clear that composite things can be dispersed, whereas simple things cannot. Being dispersed, after all, is just a matter of having your parts taken out of connection with each other, and simple things have no parts.

But, one might think, this shows that only composite things can be destroyed; for how can you destroy something other than by breaking it up into its parts?

The key question, then, is: are we composite, or simple?

Plato was, like Descartes, a proponent of soul survival — he held that we are immaterial souls. If we assume this view, then the question is whether immaterial souls are composite or simple.

A reasonable argument can be made that immaterial souls are simple rather than composite. For, arguably, we have no grip on what it would take for am immaterial thing — which is not extended in space — to have parts.

A reasonable argument can be made that immaterial souls are simple rather than composite. For, arguably, we have no grip on what it would take for am immaterial thing — which is not extended in space — to have parts.

We can then give the following argument for immortality:

ARGUMENT FROM THE SIMPLICITY OF THE SOUL

```
1. Persons are immaterial souls.
2. All immaterial things are simple.
3. Only composite things can be destroyed.
4. Immaterial souls cannot be destroyed. (2,3)
----------------
C. Persons cannot be destroyed. (1,4)
```

We've already considered defenses of the first premise. Obviously, materialists and psychological theorists who reject those defenses are unlikely to be persuaded by this argument. But should proponents of soul survival be convinced by it?

One might question either premise (2) or premise (3).

ARGUMENT FROM THE SIMPLICITY OF THE SOUL

```
1. Persons are immaterial souls.
2. All immaterial things are simple.
3. Only composite things can be destroyed.
4. Immaterial souls cannot be destroyed. (2,3)
---------------
C. Persons cannot be destroyed. (1,4)
```

One might question either premise (2) or premise (3).

The Scottish philosopher David Hume gives an interesting reply to this argument:

> "What is incorruptible must also be ingenerable. The soul, therefore, if immortal, existed before our birth. And if the former existence nowise concerned us, neither will the latter."

"What is incorruptible must also be ingenerable. The soul, therefore, if immortal, existed before our birth. And if the former existence nowise concerned us, neither will the latter."

This begins with the plausible thought that if something cannot be destroyed then it also cannot be created. So, if we are things that cannot be destroyed, then we are also things that cannot be created. So, just as (according to this argument) we will exist after our death, so we must have existed before our birth.

This poses a dilemma for the defender of the simplicity argument.

On the one hand, she can deny that we preexisted our births. But then she needs to explain why the argument for life after death is stronger than the argument for preexistence.

On the other hand, she can accept preexistence. (This was Plato's view.) But how good was your life before you were born? If life after death is just like the 'life' you had before you were born, then it does not seem to be a kind of life after death worth wanting.

This poses a dilemma for the defender of the simplicity argument.

On the one hand, she can deny that we preexisted our births. But then she needs to explain why the argument for life after death is stronger than the argument for preexistence.

On the other hand, she can accept preexistence. (This was Plato's view.) But how good was your life before you were born? If life after death is just like the 'life' you had before you were born, then it does not seem to be a kind of life after death worth wanting.

This is a serious problem for the main philosophical argument for immortality. Of course, if your main reasons for believing in life after death are due to belief in a religion like Christianity, this should strike you as somewhat unsurprising. If life after death is due to a miraculous intervention of God, it is perhaps not too surprising if there is no good purely philosophical argument for life after death.

Could technology allow me to indefinitely delay death?

For biological systems, death seems inevitable. Biological systems age and decay over time, and there is (so far as we know) no way to prevent this from happening.

Recent technological developments, though, have led many to pose the question of whether we might be able to indefinitely delay death, not by stopping or slowing the biological aging process, but by escaping our biological nature entirely.

We might attempt to do this by becoming, not biological systems, but synthetic artificially intelligent systems.

'Artificial intelligence' is a term for the ability of machines to perform tasks intelligently: for example, to strategize and to solve problems.

So defined, artificial intelligence is now all around us. There are plenty of examples of AI systems which are vastly better than humans at performing various tasks.

Could technology allow me to indefinitely delay death?

So defined, artificial intelligence is now all around us. There are plenty of examples of AI systems which are vastly better than humans at performing various tasks.

What does not yet exist is a **general** artificial intelligence: an artificial intelligence capable of doing all or almost all of the things that an ordinary adult human being can do. No machine in existence (that we know of) has general artificial intelligence.

One very interesting question is whether, and when, we will develop human level artificial intelligence. A recent survey of researchers in the field gave an average guess of the year 2100 — but opinions vary widely.

Today we will focus on one way in which you might become a non-biological entity with general artificial intelligence — a pathway which may include choices which will be available to you in your lifetime.

Today we will focus on one way in which you might become a non-biological entity with general artificial intelligence — a pathway which may include choices which will be available to you in your lifetime.

Suppose that it is the year 2045. We have now developed silicon devices which replicate but improve upon the functioning of neurons or clusters of neurons. The silicon devices do just the same things as the neurons they replace, but more quickly and more efficiently.

You have the opportunity to have part of your brain replaced with silicon devices of this kind. Lots of your friends have done this, and they can process information much more quickly than they used to be able to. You find yourself consistently underperforming relative to your peers who have had the synthetic replacement done — and you suspect that your newly super-smart friends are beginning to find it kind of boring to talk to you.

If given the opportunity to go in for partial synthetic replacement, would you do it?

You have the opportunity to have part of your brain replaced with silicon devices of this kind. Lots of your friends have done this, and they can process information much more quickly than they used to be able to. You find yourself consistently underperforming relative to your peers who have had the synthetic replacement done — and you suspect that your newly super-smart friends are beginning to find it kind of boring to talk to you.

If given the opportunity to go in for partial synthetic replacement, would you do it?

Once you have part of your brain replaced in this way, it seems to be irresistible to gradually have all of your brain replaced in this way (assuming that the surgery is affordable). Why would you want to keep part of your underperforming biological brain around?

Suppose that you were now given the opportunity to have your synthetic brain supplemented with improved memory, so that more of your memories could be reliably stored and retrieved. Would you opt for that as well?

Suppose that you were now given the opportunity to have your synthetic brain supplemented with improved memory, so that more of your memories could be reliably stored and retrieved. Would you opt for that as well?

You would now have become, at least in part, an AI system with greater than human level intelligence. Your intelligence would be in many ways like human intelligence — but you would have much faster processing speed and much better memory.

Having traded in your brain for an artificial system, you might become annoyed with the limitations of your other biological parts.

For example, we could presumably replace all of your organs and body parts with synthetic systems which were not subject to decay, and which worked much better than your current biological parts. Perhaps you would no longer have to sleep or eat (though you might have the option to do so).

Having traded in your brain for an artificial system, you might become annoyed with the limitations of your other biological parts.

For example, we could presumably replace all of your organs and body parts with synthetic systems which were not subject to decay, and which worked much better than your current biological parts. Perhaps you would no longer have to sleep or eat (though you might have the option to do so).

This might make you effectively immortal (barring some disaster). After all, replacement of any of your failed parts would now be a straightforward matter.

Would you trade in the rest of your biological parts for synthetic replacements? (Again, it may help to imagine that your friends have all done this, and are now annoyed with your "biological" limitations.)

At this point it seems that you would have become an artificial intelligence. You would no longer be a biological organism.

At this point it seems that you would have become an artificial intelligence. You would no longer be a biological organism.

The scenarios just laid out show that it is not wildly implausible to think that you will be faced with choices like this in your lifetime, and that it is not wildly implausible to think that decisions which lead to this outcome would be very tempting.

But at this stage it is natural to pose the following question: would the synthetic being which results from these changes be **you**? Would you survive?

Let's look at some examples which will help make the issues clear.

Maria is considering whether to "go synthetic." Being a cautious person, she does this gradually. At t1, she has one neuron replaced by a silicon device which replicates the functioning of that neuron.

Would she notice a change? It seems that she would not.

So now suppose that she has a second neuron replaced. Would she notice a change? Again, it seems that she would not.

This process might continue until all of Maria's neurons have been replaced. Gradually, this synthetic system inside her head could then be supplemented in ways which gave it more memory and greater processing speed. Here Maria would notice a difference — she would be able gradually so solve problems faster, and remember much more. But it does not seem as though changes of this kind could make it "no longer Maria."

Once we have gone his far, it seems pretty clear that we could provide synthetic replacements of all of Maria's body parts without her ceasing to exist. Surely replacing Maria's index finger with a synthetic replacement need not involve a change in identity!

Once we have gone his far, it seems pretty clear that we could provide synthetic replacements of all of Maria's body parts without her ceasing to exist. Surely replacing Maria's index finger with a synthetic replacement need not involve a change in identity!

Let's call the outcome of this procedure "Digi-Maria." It seems plausible that Maria could survive this process; so it seems plausible that Maria=Digi-Maria.

Let's call this procedure **slow gradual destructive uploading.** It is destructive because Maria's neurons (and other parts of her body) are destroyed as they are replaced. It is gradual because this happens, not all at once, but one by one. It is slow because it takes place over a long period of time.

Our discussion so far suggests:

You can survive slow gradual destructive uploading.

Our discussion so far suggests:

> You can survive slow gradual destructive uploading.

Now imagine a procedure just like this one, but in which the replacements happen more rapidly. Maybe each replacement takes just a small fraction of a second. This would be a case of **fast** gradual destructive uploading.

Surely this makes no difference; the speed with which the replacement of the parts makes no difference. That suggests:

> If you can survive slow gradual destructive uploading, you can survive fast gradual destructive uploading.

You can survive slow gradual destructive uploading.

If you can survive slow gradual destructive uploading, you can survive fast gradual destructive uploading.

You can survive fast gradual destructive uploading.

Let's now look at a very different example.

Caleb is considering whether to go synthetic. But he does not have Maria's patience, and is nervous about having parts of his body destroyed.

So a synthetic version of Caleb — Digi-Caleb — is created while Caleb watches. Digi-Caleb is like Caleb in certain ways (just as DigiMaria is like Maria in certain ways) — but of course Digi-Caleb is much smarter than Caleb, and less prone to bodily damage of various kinds.

Suppose now that, satisfied with the production of Digi-Caleb, Caleb decides to take some cyanide, in the hopes that he will live on as Digi-Caleb. Does Caleb survive?

Surely not!

This is a case of **instant non-destructive uploading**. Our discussion suggests:

You cannot survive instant non-destructive uploading.

Let's look at one last example.

Mindful of Caleb's fate, Emily decides to take a different path. Like Caleb, she lacks the patience for gradual uploading. But she wants to become a synthetic thing, and knows that Caleb failed to achieve this.

So Emily decides to go for **instant destructive uploading**. In this process, Emily's body is destroyed, and right away a synthetic version — DigiEmily — is created.

Did Emily survive the procedure?

A strong case can be made that she did not, because Emily seems relevantly just like Caleb — the only difference is that Emily was destroyed earlier in the process than Caleb was. But why should that matter?

This suggests:

```
If you cannot survive
instant non-destructive
uploading, you cannot
   survive instant
destructive uploading.
```

If you cannot survive instant non-destructive uploading, you cannot survive instant destructive uploading.

You cannot survive instant non-destructive uploading.

You cannot survive instant destructive uploading.

But now we have a puzzle. Compare the cases of fast gradual destructive uploading and instant destructive uploading. Could these really be that different?

After all, we can imagine that the fast gradual destructive uploading process is very fast indeed. Perhaps the entire process of destruction and replacement takes only a small fraction of a second.

Why should it matter if each step in the process takes place 0.0001 seconds after the previous one, versus them happening all at once?

This suggests:

> If you cannot survive instant destructive uploading, you cannot survive fast gradual destructive uploading.

You can survive slow gradual destructive uploading.

If you can survive slow gradual destructive uploading, you can survive fast gradual destructive uploading.

If you cannot survive instant non-destructive uploading, you cannot survive instant destructive uploading.

You can survive fast gradual destructive uploading.

You cannot survive instant destructive uploading.

You cannot survive instant non-destructive uploading.

You cannot survive fast gradual destructive uploading.

If you cannot survive instant destructive uploading, you cannot survive fast gradual destructive uploading.

The uploading paradox

1. You can survive slow gradual destructive uploading.
2. If you can survive slow gradual destructive uploading, you can survive fast gradual destructive uploading.
3. You can survive fast gradual destructive uploading. (1,2)
4. If you cannot survive instant non-destructive uploading, you cannot survive instant destructive uploading.
5. You cannot survive instant non-destructive uploading.
6. You cannot survive instant destructive uploading. (4,5)
7. If you cannot survive instant destructive uploading, you cannot survive fast gradual destructive uploading.
8. You cannot survive fast gradual destructive uploading. (6,7)

C. You can and cannot survive fast gradual destructive uploading. (4,8)

The uploading paradox

1. You can survive slow gradual destructive uploading.
2. If you can survive slow gradual destructive uploading, you can survive fast gradual destructive uploading.
3. You can survive fast gradual destructive uploading. (1,2)
4. If you cannot survive instant non-destructive uploading, you cannot survive instant destructive uploading.
5. You cannot survive instant non-destructive uploading.
6. You cannot survive instant destructive uploading. (4,5)
7. If you cannot survive instant destructive uploading, you cannot survive fast gradual destructive uploading.
8. You cannot survive fast gradual destructive uploading. (6,7)

C. You can and cannot survive fast gradual destructive uploading. (4,8)

The conclusion of the uploading paradox is a contradiction; so at least one of its independent premises must be false. So a question which any view of survival must face is: which one is it?

If there is no life after death, should I fear death?

One reason for interest in the question of whether life after death is possible is the thought that, if there is no life after death, then death would be a terrible thing.

But there is an ancient tradition which says that this is a mistake: that death, even if there is no life after death, is nothing to be feared. (Note that we should distinguish the fear of death from the fear of dying — no one disputes that dying painfully can be a bad thing.)

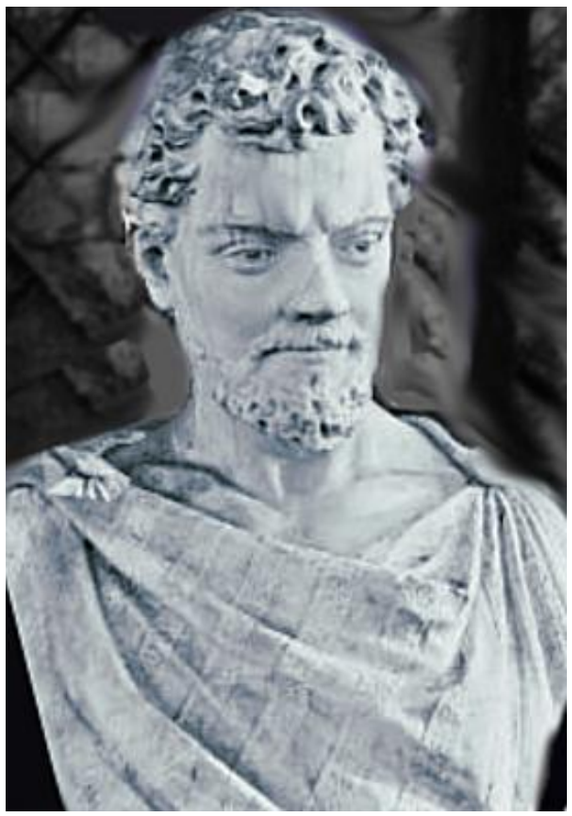

Lucretius, a Roman philosopher who lived in the first century B.C., was part of this tradition. In his poem *De Rerum Natura* (*On the Nature of Things*), he gave two short and intriguing arguments against the idea that death is at all a bad thing.

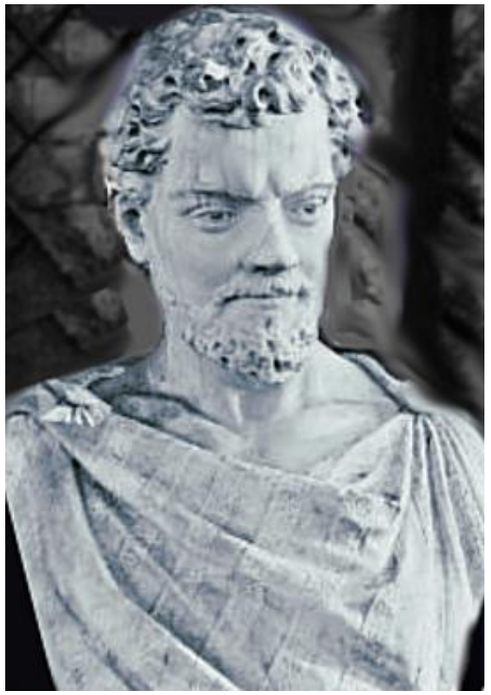

Lucretius, a Roman philosopher who lived in the first century B.C., was part of this tradition. In his poem *De Rerum Natura* (*On the Nature of Things*), he gave two short and intriguing arguments against the idea that death is at all a bad thing.

Here is the first:

"If it happens that people are to suffer unhappiness and pain in the future, they themselves must exist at that future time for harm to be able to befall them. Since death takes away this possibility by preventing the existence of those who might have been visited by troubles, you may be sure that there is nothing to fear in death. Those who no longer exist cannot become miserable, and it makes not one speck of difference whether or not they have ever been born once their mortal life has been snatched away by deathless death."

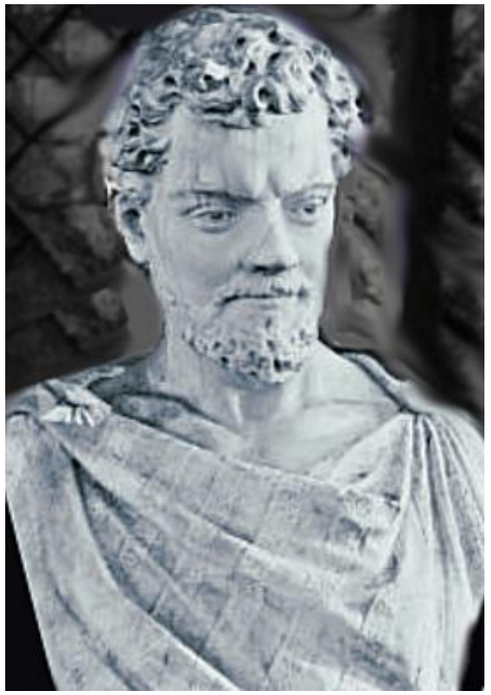

"If it happens that people are to suffer unhappiness and pain in the future, they themselves must exist at that future time for harm to be able to befall them. Since death takes away this possibility by preventing the existence of those who might have been visited by troubles, you may be sure that there is nothing to fear in death. Those who no longer exist cannot become miserable, and it makes not one speck of difference whether or not they have ever been born once their mortal life has been snatched away by deathless death."

Lucretius' idea is that after death we will not exist. But if we will not exist, it is impossible for us to be harmed in any way; and if this is right, there is nothing to fear from death.

In slogan form: 'If death is there, we are not, and if we are there, death is not.' So we have nothing to fear from death.

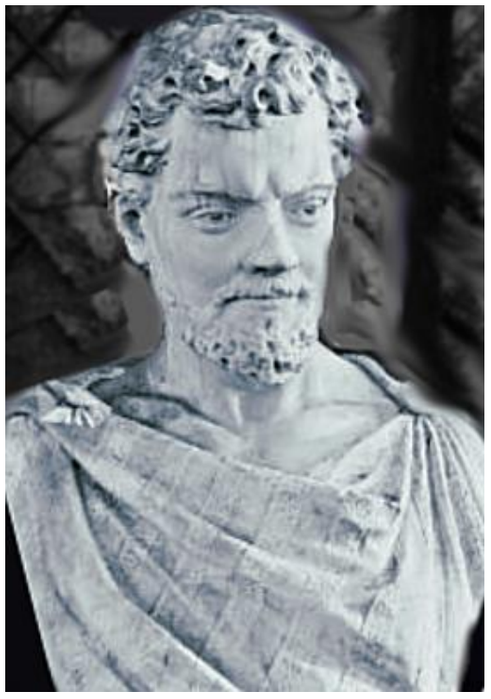

Lucretius' idea is that after death we will not exist. But if we will not exist, it is impossible for us to be harmed in any way; and if this is right, there is nothing to fear from death.

In slogan form: 'If death is there, we are not, and if we are there, death is not.' So we have nothing to fear from death.

Here is one way in which Lucretius' argument can be represented:

The Lucretian argument

```
1. The only things I should fear are
   experiences which I undergo.
2. When I am dead, I undergo no
   experiences.
-------------------------
C. I should not fear death. (1,2)
```

The Lucretian argument

```
1. The only things I should fear are
   experiences which I undergo.
2. When I am dead, I undergo no
   experiences.
-------------------------
C. I should not fear death. (1,2)
```

Of course, one might dispute the second premise — but here we are assuming for the sake of argument that there is no life after death. Is the first premise plausible?

Here is a natural response to the first premise: 'Yes, it is true that I will have no experiences after I die. But just that fact is part of what makes death so horrible. What is bad about death is that after death I will not exist — and my non-existence is the worst thing that can happen to me.'

Those who fear death because they fear the end of their existence are unlikely to be consoled by Lucretius' first argument.

One can think of Lucretius' second argument as a reply to this objection.

One can think of Lucretius' second argument as a reply to this objection.

> "Look back at time ... before our birth. In this way Nature holds before our eyes the mirror of our future after death. Is this so grim, so gloomy?"

Here Lucretius points out that we are already familiar with times at which we do not exist: namely, all of those times before our birth. When you think about times before your birth, are you filled with horror? Lucretius thinks not. But then you should not fear times after your death, because those will be just the same.

Because this draws a kind of parallel between pre-birth and post-death times, this is sometimes called the **symmetry argument**.

> "Look back at time ... before our birth. In this way Nature holds before our eyes the mirror of our future after death. Is this so grim, so gloomy?"

Because this draws a kind of parallel between pre-birth and post-death times, this is sometimes called the **symmetry argument**.

A natural reply is to say: 'OK, I agree that there is nothing especially fearsome about my past nonexistence. But future nonexistence is different; I should fear my future nonexistence even if I do not fear my past nonexistence.'

Most of us have a negative feeling about future nonexistence which we do not have about past nonexistence. Lucretius' challenge is to justify this difference in our attitudes.

A natural reply is to say: 'OK, I agree that there is nothing especially fearsome about my past nonexistence. But future nonexistence is different; I should fear my future nonexistence even if I do not fear my past nonexistence.'

Most of us have a negative feeling about future nonexistence which we do not have about past nonexistence. Lucretius' challenge is to justify this difference in our attitudes.

But why **shouldn't** we feel about post-death nonexistence the same way we feel about pre-birth nonexistence? Why would it be rational to have very different attitudes toward two equivalent states of affairs just because they happen to occupy different locations in time? We don't, after all, make parallel distinctions between events occurring in different locations in space.

The fact is that people do systematically exhibit time bias: they prefer good things to be in their future and bad things in their past. The interesting question raised by the symmetry argument is whether this feature of human thinking is a rational one, or one we should attempt to overcome. If the latter is correct, then the symmetry argument has considerable force.