

What should I believe?

Today we begin a new topic. For the next few weeks, we will be investigating the question, **What should I believe?**

This is different than all of the questions we've asked so far. It is not a question about how the world is — like the questions about whether God exists, about whether we have free will, and about what we are. Instead, it is a **practical** question — a question about what we ought to do.

Moreover, I think that this is a kind of question about which most of us have lots of opinions. Let's look at some examples.

One kind of interesting case to think about is belief in conspiracy theories.

Some of these are silly and, perhaps, harmless. Examples might include the belief that moon landing was faked, or that the earth is flat. Believers in these kinds of conspiracy theories typically discount evidence which seems to count against their theories. (For example, the fact that ~400,000 people worked toward the Apollo moon landing, and none have admitted that it was faked, or photographs of the earth from space which seem to show it to be spherical.)



Other examples of conspiracy theories are definitely not harmless.

One example is the contemporary American political conspiracy theory QAnon, which holds (among other things) that Obama and Hillary Clinton were organizing a coup while running an international child sex ring, and that Trump feigned collusion with Russia in order to get Robert Mueller's help exposing this. It is popular enough that an app devoted to distributing more information about the conspiracy theory's claims was in 2018 for a time the 10th most popular paid app on the Apple App Store.



It seems pretty clear that people who believe conspiracy theories of this kind are making a mistake of a certain kind; they are believing something that they should not believe. Let's call these cases of **bad belief**.

And this is not just because the conspiracy theories are false. Intuitively, sometimes you can have very good reason to form a belief which turns out to be false (say, if you have misleading evidence). The mistake that conspiracy theorists are making is a different kind of mistake.

What are some other examples of people making mistakes of this kind — people believing things that they should not believe?

What are some other examples of people making mistakes of this kind — people believing things that they should not believe?

One very common kind of example comes from cases of **wishful thinking**. Example: me, every August, thinking about the upcoming Notre Dame football season.

Another kind of example: people who form beliefs about their future on the basis of the horoscopes published in the *Observer*.

These are all examples of people believing things they should not believe. What are some examples of the opposite phenomenon — people forming beliefs as they **should** form them?

These are all examples of people believing things they should not believe. What are some examples of the opposite phenomenon — people forming beliefs as they **should** form them?

The easiest examples are people who seem to weigh, and respond appropriately to, their evidence. Examples: Sherlock Holmes; responsible scientists; careful jurors.

Notice that none of these people are infallible; even responsible scientists make mistakes. But intuitively they are going about belief formation in the right way.

There are also plenty of simple and everyday examples of this kind of thing. Suppose that you see people walking around outside with umbrellas open over their heads, and form the belief that it is raining. Could your belief be false? Of course. But intuitively, given your experience of the world, you are forming the belief that you should form.

Let's call these cases of **good belief**.

These are all easy cases. But there are plenty of hard cases too — and, in fact, you might think, philosophy is a kind of machine for generating hard cases! Haven't we already seen lots of cases in which there are arguments on both sides of an issue, and where it is hard to tell which argument is better?

Here is a hard case of interest:

The Believer

I've always believed that there is a God. I never really thought about what my evidence is for this claim. But now I wonder whether I have good reason for my beliefs. Some of the arguments for God's existence sound good, but all face objections that I am not sure how to answer. Still, I continue to believe that God exists.

Is The Believer forming the beliefs he or she should form, or not? Is it a case of good belief, or a case of bad belief? What do you think?

The Believer

I've always believed that there is a God. I never really thought about what my evidence is for this claim. But now I wonder whether I have good reason for my beliefs. Some of the arguments for God's existence sound good, but all face objections that I am not sure how to answer. Still, I continue to believe that God exists.



Start the presentation to activate live content

If you see this message in presentation mode, install the add-in or get help at PollEv.com/app



Is The Believer forming the beliefs he or she should form, or not? Is it a case of good belief, or a case of bad belief? What do you think?

It is worth emphasizing that all of us are like The Believer on some issues. Most of us have moral beliefs, or political beliefs, which we hold strongly but which we might find it difficult to argue for in a persuasive way.

Here's a way in which we might try to answer our question about whether The Believer should believe as s/he does. When we think about examples of good belief, and bad belief, the following thought seems very plausible:

It isn't just an inexplicable fact that horoscope beliefs are bad beliefs, and that responsible scientist beliefs are good beliefs. Instead, there are general principles which determine whether someone should, in a certain circumstance, form a certain belief, or not.

Let's call these general principles the [rules of belief](#).

It isn't just an inexplicable fact that horoscope beliefs are bad beliefs, and that responsible scientist beliefs are good beliefs. Instead, there are general principles which determine whether someone should, in a certain circumstance, form a certain belief, or not.

Let's call these general principles the **rules of belief**.

These rules might come in two flavors. One kind of rule might be a rule which tells you that in certain circumstances you **should** form a belief. A second kind of rule might tell you that in certain circumstances you **should not** form a belief.

It seems very plausible that there must be rules of this kind which explain the difference between cases of good belief and cases of bad belief.

And it also seems plausible that, if we can figure out what these rules are, we'll be able to figure out whether The Believer should believe what s/he does.

The attempt to figure out the rules of belief is part of the field of philosophy known as **epistemology** — so called because 'epistêmê' is the ancient Greek word for knowledge, and questions about what we should believe are connected with questions about what we can know.

Our first attempt to formulate a rule of beliefs comes from a text familiar from our discussion of the nature of the self: Descartes' *Meditations*.

Some years ago I was struck by the large number of falsehoods that I had accepted as true in my childhood, and by the highly doubtful nature of the whole edifice that I had subsequently based on them. I realized that it was necessary, once in the course of my life, to demolish everything completely and start again right from the foundations if I wanted to establish anything at all in the sciences that was stable and likely to last.

Descartes was aware that, by accepting common sense and the opinions of others, he could come to believe falsehoods. He therefore undertook to try to find a method — a way of forming beliefs — which would provide a secure foundation for belief.

Descartes was aware that, by accepting common sense and the opinions of others, he could come to believe falsehoods. He therefore undertook to try to find a method — a way of forming beliefs — which would provide a secure foundation for belief.

Reason now leads me to think that I should hold back my assent from opinions which are not completely certain and indubitable just as carefully as I do from those which are patently false. So, for the purpose of rejecting all my opinions, it will be enough if I find in each of them at least some reason for doubt.

He hit upon the **method of doubt**: for any belief which he could coherently doubt to be true, he would give up that belief. Only then could he be certain not to believe any falsehoods.

He then applies this method to one of the most fundamental kinds of belief we form: beliefs about our environment formed on the basis of sensory experience.

He then applies this method to one of the most fundamental kinds of belief we form: beliefs about our environment formed on the basis of sensory experience.

He raises the question: how can I tell whether a given sense experience of mine is accurate?

After noting that only 'madmen' doubt the reliability of their sense experiences, Descartes notices something about his own experiences:

As I think about this more carefully, I see plainly that there are never any sure signs by means of which being awake can be distinguished from being asleep.

As you know, Descartes goes on to argue that, although he must doubt the reliability of his sense experiences, he cannot doubt that he exists. But rather than going on to think about that aspect of his views (as we did in our discussion of dualism), today I want to focus on Descartes' central point about his sensory experiences of the world: namely, that "there are never any sure signs by means of which being awake can be distinguished from being asleep."

The key point is the following claim:

The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

We are all familiar with experiences which seem not to represent our environment accurately.

Some are every day experiences.




The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

We are all familiar with experiences which seem not to represent our environment accurately.


Others are intentionally constructed illusions which are used in vision science to study our mechanisms for representing the world around us.



The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

We are all familiar with experiences which seem not to represent our environment accurately.



The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

But these are just specific examples. Is it really true that for **any** experience, we can imagine a matching — i.e., indiscriminate — illusion?

A number of different thought experiments suggest that we can.

One, which Descartes mentions, is the possibility that we are simply dreaming.



The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

But these are just specific examples. Is it really true that for **any** experience, we can imagine a matching — i.e., indiscriminate — illusion?

A number of different thought experiments suggest that we can.

Another, which Descartes also discusses, is the possibility that we are being deceived by an evil demon.



The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

But these are just specific examples. Is it really true that for **any** experience, we can imagine a matching — i.e., indiscriminate — illusion?

A number of different thought experiments suggest that we can.

We might also imagine that we are simply brains in vats which are being stimulated to cause illusory sense experiences as part of some nefarious scientific experiment.



The matching hallucination assumption

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

Once we notice this, though, this point can be used to generate a powerful argument for the conclusion that we cannot know anything around us on the basis of sense experience.

This is because the following principle seems very plausible:

If I cannot distinguish between two situations, then I cannot know which of them is real.

If I cannot distinguish between two situations, then I cannot know which of them is real.

Suppose that I tell you that, behind the lectern, I have an object. It is either a triangle or a circle.

Given that it is behind the lectern, the two different possibilities are indistinguishable to you.

Does it follow that you cannot know whether there is a triangle or a circle behind the lectern?

If, as is plausible, you think 'No', that might seem to provide reason to believe that if you cannot distinguish between two situations, then you cannot know which is real.

The matching hallucination assumption



For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

If I cannot distinguish between two situations, then I cannot know which of them is real.

Sense experience is my only way of knowing whether there is an external world.

I can never know whether any sense experience of mine is accurate.

I do not know whether there is an external world.



1. For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.
2. If I cannot distinguish between two situations, then I cannot know which of them is real.
3. I can never know whether any sense experience of mine is accurate. (1,2)
4. If I cannot know whether any of my sense experiences are accurate, I cannot know whether there is an external world.

C. I do not know whether there is an external world. (3,4)

Skepticism about some domain is the claim that one cannot have knowledge about that domain. This is an argument for **skepticism about our knowledge of the external world.**

We've already seen that there are strong reasons for accepting premises (1) and (2), and premise (4) seems quite plausible.

1. For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.
2. If I cannot distinguish between two situations, then I cannot know which of them is real.
3. I can never know whether any sense experience of mine is accurate. (1,2)
4. If I cannot know whether any of my sense experiences are accurate, I cannot know whether there is an external world.

C. I do not know whether there is an external world. (3,4)

You might be tempted to reply like this:
'OK, this shows that I can't **know** that there is an external world. But I should still **believe** that there is one.'

The problem is that a parallel argument seems to rule even this out.

1. For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.
2. If I cannot distinguish between two situations, then I cannot **know** which of them is real.
3. I can never **know** whether any sense experience of mine is accurate. (1,2)
4. If I cannot know whether any of my sense experiences are accurate, I cannot **know** whether there is an external world.

C. I do not **know** whether there is an external world. (3,4)

1. For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.
- 2*. If I cannot distinguish between two situations, then **I should not believe** that one but not the other is real.
- 3*. I **should never believe** that any sense experience of mine is accurate. (1,2*)
- 4*. If I should never believe that any sense experience of mine is accurate, I **should never form beliefs** about the external world.

C*. I **should never form beliefs** about the external world. (3*,4*)

Is the second argument as strong as the first?

It is tempting to think that we should be able to respond to Descartes by finding some way to show that certain experiences are not illusions.

For example, one might argue that, since our sense experiences are usually accurate, it is reasonable to form beliefs about the external world on their basis.

But how do we know that our sense experiences are usually accurate? Presumably on the basis of past sense experiences. And those experiences can be doubted just as much as our present experiences. Any attempt to respond to Descartes seems to assume the very thing we are trying to show.

If you think about it, it seems like any attempt to reply to Descartes' argument is going to face this kind of problem.

The seeming impossibility of replying to Descartes' argument — and hence the seeming impossibility of providing a satisfactory proof of an external world — was seen by the great German philosopher Immanuel Kant as a “scandal to philosophy.”

Descartes' reasoning relies on the following rule of belief:

Doubt \rightarrow No Belief

If you cannot distinguish between a situation in which P and a situation in which not-P, do not believe P.

By the line of reasoning just laid out, this rule seems to lead directly to the (very) surprising conclusion that you should never form beliefs about the objects you (seem to) perceive.

Indeed, to show that this rule leads to these kinds of surprising consequences, we don't even need to consider scenarios as extreme as Descartes' 'evil demon' scenario. You are not now in a position to distinguish between a situation in which your dorm room was robbed five minutes ago and one in which it wasn't. It then seems to follow from Doubt \rightarrow No Belief that you shouldn't believe that the possessions you left in your dorm room are safe.

Doubt \rightarrow No Belief

If you cannot distinguish between a situation in which P and a situation in which not-P, do not believe P.

Our second reading for today is from someone who has a very different perspective on our beliefs about the external world than Descartes did.

According to G.E. Moore, it is no “scandal to philosophy” that we cannot prove the existence of the external world — for in fact, he thought, proofs of this kind are extremely easy to give.

We can think of Moore as endorsing the following rule of belief:

Proof \rightarrow Belief

If you can prove P, believe P.

At first glance, our two rules look perfectly consistent. Moore tries to show that they are not.

He presents his proof of an external world in the following passage:

I can now give a large number of different proofs, each of which is a perfectly rigorous proof; and that at many other times I have been in a position to give many others. I can prove now, for instance, that two human hands exist. How? By holding up my two hands, and saying, as I make a certain gesture with the right hand, 'Here is one hand', and adding, as I make a certain gesture with the left, 'and here is another'. And if, by doing this, I have proved *ipso facto* the existence of external things, you will all see that I can also do it now in numbers of other ways: there is no need to multiply examples.



I can now give a large number of different proofs, each of which is a perfectly rigorous proof; and that at many other times I have been in a position to give many others. I can prove now, for instance, that two human hands exist. How? By holding up my two hands, and saying, as I make a certain gesture with the right hand, 'Here is one hand', and adding, as I make a certain gesture with the left, 'and here is another'. And if, by doing this, I have proved *ipso facto* the existence of external things, you will all see that I can also do it now in numbers of other ways: there is no need to multiply examples.



Moore's proof can be laid out as follows:

1. Here is one hand.
 2. Here is another hand.
-
- C. There are two hands. (1,2)

1. Here is one hand.
2. Here is another hand.
-
- C. There are two hands. (1,2)

It is, I think, safe to say that this is not the sort of proof that Moore's audience was expecting. We might ask: what does Moore mean when he says that this is a **proof**?

Moore tells us. He says that an argument is a proof if it satisfies three conditions:

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3) Its conclusion follows from its premises.

It is natural to find Moore's proof a little bit puzzling — and unsatisfying. But let's separate out two different questions which we can ask about his proof.

- 1. Here is one hand.
- 2. Here is another hand.
-
- C. There are two hands. (1,2)

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3) Its conclusion follows from its premises.

It is natural to find Moore's proof a little bit puzzling — and unsatisfying. But let's separate out two different questions which we can ask about his proof.

If an argument meets Moore's definition of a proof, does it provide knowledge of its conclusion?

Does Moore's argument meet his definition of a proof?

If an argument meets Moore's definition of a proof, does it provide knowledge of its conclusion?

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3) Its conclusion follows from its premises.

You might think, at first, yes: if we know the premises, and the conclusion follows from the premises, doesn't this give us knowledge of the conclusion?

Well, not quite. It might be the case that the conclusion actually follows from the premises, but that we don't know that it does. Here's an example:

- 1. There are infinitely many numbers.
- 2. A prime number is one whose only divisors are 1 and itself.
-
- C. There are infinitely many prime numbers. (1,2)

If an argument meets Moore's definition of a proof, does it provide knowledge of its conclusion?

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3) Its conclusion follows from its premises.

- 1. There are infinitely many numbers.
- 2. A prime number is one whose only divisors are 1 and itself.
-
- C. There are infinitely many prime numbers. (1,2)

This meets Moore's definition of a proof, and yet might not provide knowledge of its conclusion to someone who does not know that it is valid.

This suggests a slight modification of Moore's definition.

If an argument meets Moore's definition of a proof, does it provide knowledge of its conclusion?

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3*) Its conclusion is **known to follow** from its premises.

With this modification in hand, can we conclude that if an argument meets Moore's definition of a proof, then it provides knowledge of its conclusion?

Does Moore's
argument meet his
definition of a proof?

Moore's definition of a proof

- (1) Its premises are distinct from its conclusion.
- (2) Its premises are known to be true.
- (3*) Its conclusion is **known to follow** from its premises.

- 1. Here is one hand.
- 2. Here is another hand.
-
- C. There are two hands. (1,2)

Let's turn now to the question of whether Moore's argument does in fact meet his definition.

It obviously meets condition (1); and also pretty obviously meets condition (3*). So our question boils down to this one: does Moore really know the premises of his argument?

Does Moore's
argument meet his
definition of a proof?

1. Here is one hand.
2. Here is another hand.
-
- C. There are two hands. (1,2)

It is pretty easy to adapt our earlier argument for skepticism about the external world to make an argument that Moore does not know the premises of his argument.

1. I can imagine a situation which is indistinguishable from a visual experience of my hands but in which I have no hands.
2. If I cannot distinguish between two situations, then I cannot know which of them is real.
3. I can never know whether a visual experience of my hands is accurate. (1,2)
4. If I cannot know whether any of my sense experiences are accurate, I cannot know whether there is an external world.
-
- C. I do not know whether I have hands. (3,4)

Moore anticipates the objection that he does not know the premises of his argument, and responds as follows:

(2) I certainly did at the moment *know* that which I expressed by the combination of certain gestures with saying the words 'There is one hand and here is another'. I *knew* that there was one hand in the place indicated by combining a certain gesture with my first utterance of 'here' and that there was another in the different place indicated by combining a certain gesture with my second utterance of 'here'. How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case! You might as well suggest that I do not know that I am now standing up and talking—that perhaps after all I'm not, and that it's not quite certain that I am!

Moore is emphasizing the fact that, in ordinary life, we do take ourselves to know claims like the premises of his argument. So why should we now, once we start doing philosophy, discard these beliefs?

Here is a different way to put the same point. We have, it seems, a conflict between the following two claims:

2. If I cannot distinguish
between two situations,
then I cannot know which
of them is real.

I know that I have hands.

One can think of Moore as asking the proponent of our skeptical argument: which of these do you feel more sure of? Which, if you had to, would you bet your life on?

Which claim are you more confident is true?

That if you cannot distinguish
between two situations, then
you cannot know which of
them is real.

That you know that you have
hands.

Moore considers a second reason for thinking that he does not know the premises of his argument:

But another reason why some people would feel dissatisfied with my proofs is, I think, not merely that they want a proof of something which I haven't proved, but that they think that, if I cannot give such extra proofs, then the proofs that I have given are not conclusive proofs at all. And this, I think, is a definite mistake. They would say: 'If you cannot prove your premiss that here is one hand and here is another, then you do not know it. But you yourself have admitted that, if you did not know it, then your proof was not conclusive. Therefore your proof was not, as you say it was, a conclusive proof.'

This objection is based on the view that if you cannot prove something, then you do not know it.

This objection is based on the view that if you cannot prove something, then you do not know it.

This suggests the following rule of belief:

No Proof \rightarrow No Belief

If you can't prove P, don't believe P.

Since (as Moore concedes) he has no proof of his premises, this rule implies that you should not believe them.

The question is: is No Proof \rightarrow No Belief a plausible rule of belief?

While intuitively appealing, this principle faces two serious objections.

No Proof \rightarrow No Belief

If you can't prove P, don't believe P.

The first is that the principle seems to imply that we shouldn't believe anything.

Suppose (for *reductio*) that I should believe some claim P1. It follows from our rule that I must be able to prove P1; so it follows from our rule that there must be some other claims — call them P2 and P3 — which I should believe and from which P1 follows.

Let's focus on P2. If I know it, then from our rule it follows that I must be able to prove it. But then there must be some other claims — call them P4 and P5 — which I should believe and from which P2 follows.

Let's now focus on P4.

Actually, let's not. Can you see a pattern here?

No Proof → No Belief

If you can't prove P, don't believe P.

No Proof → No Belief says that, for every claim I should believe, there must be some other claims which I should believe which can be used to prove the first one.

But then one of two things must be true.

First option: this process never comes to an end. For any claim I should believe, there are infinitely many others that I believe and should believe. But I don't believe infinitely many things. So, if we take this first option, I shouldn't believe anything.

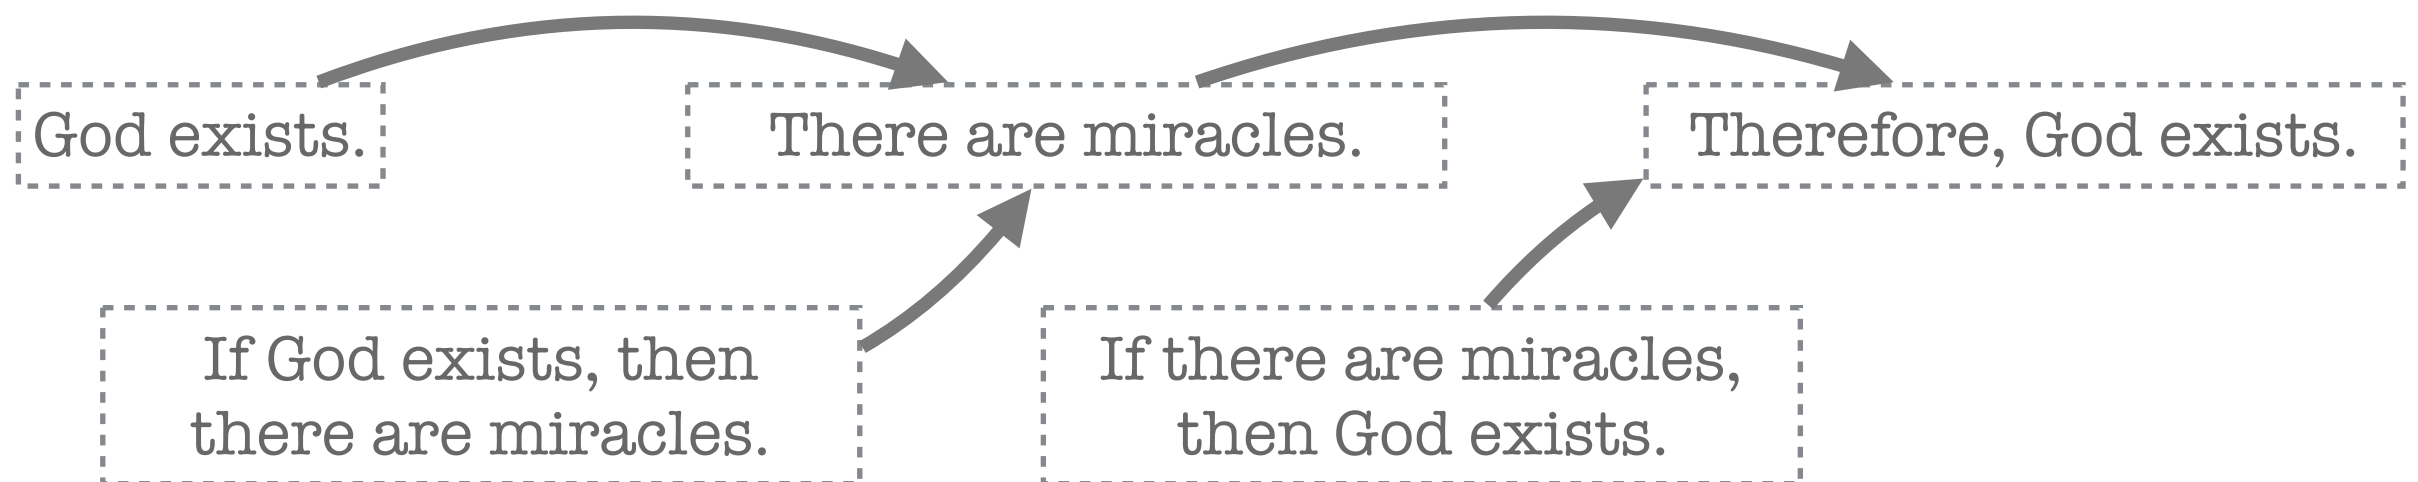
No Proof → No Belief

If you can't prove P, don't believe P.

But then one of two things must be true.

Second option: the process goes in a circle, so that (for example) P is used to prove Q, and Q is used to prove R, and R is used to prove P. But it does not seem as though this sort of circular reasoning can be a good reason to form a belief.

Imagine, for example, that one argued like this:



No Proof \rightarrow No Belief

If you can't prove P, don't believe P.

This might remind you a bit of Aquinas' first cause argument. Just like a chain of causes, every chain of reasoning must either be infinite, circular, or have some unproven premise. But the first two can't explain why I should believe anything, and, if No Proof \rightarrow No Belief is true, the last one can't either. So, if No Proof \rightarrow No Belief is true, I shouldn't believe anything.

No Proof \rightarrow No Belief

If you can't prove P, don't believe P.

Here is the second problem with using No Proof \rightarrow No Belief as an argument against Moore (or anyone else). It seems that one can legitimately use this principle in an argument only if one should believe it.

But if we should believe No Proof \rightarrow No Belief, then (by No Proof \rightarrow No Belief itself) one must have a proof of it.

But we have no proof of it.

So, in a way, No Proof \rightarrow No Belief is a principle which implies that we should not believe it. That is not a good quality for a principle to have!

Doubt \rightarrow No Belief

If you cannot distinguish between a situation in which P and a situation in which not-P, do not believe P.

Proof \rightarrow Belief

If you can prove P, believe P.

No Proof \rightarrow No Belief

If you can't prove P, don't believe P.

Let's take stock. We've now considered three candidate rules of belief.

We've seen that both of our "no belief" principles are open to substantial challenge. But surely, one might think, there must be **some** principle which explains why certain beliefs are bad beliefs.

This point can be brought out by examining one of the world's fastest growing religions: Pastafarianism.

This point can be brought out by examining one of the world's fastest growing religions: Pastafarianism.

Q: How do Pastafarians believe our world was created?

A: We believe the Flying Spaghetti Monster created the world much as it exists today, but for reasons unknown made it appear that the universe is billions of years old (instead of thousands) and that life evolved into its current state (rather than created in its current form). Every time a researcher carries out an experiment that appears to confirm one of these "scientific theories" supporting an old earth and evolution we can be sure that the FSM is there, modifying the data with his Noodly Appendage. We don't know why He does this but we believe He does, that is our Faith.



← A ChrisFSMas
tree



As you might guess, many Pastafarians take a somewhat less than serious attitude toward the tenets of Pastafarianism (though some apparently do not).

But suppose that someone were a serious Pastafarian. We would, I take it, be inclined to think that there is something irrational about his beliefs.

And this might be so even if we could not come up with any decisive argument against Pastafarianism.

In that case, it seems, we would want to provide some sort of standard for rational belief, and claim that Pastafarianism does not meet that standard.

In that case, it seems, we would want to provide some sort of standard for rational belief, and claim that Pastafarianism does not meet that standard.

No Proof → No Belief was a failed attempt to provide such a standard; can we do better?

To do so, it seems, we have to allow that it is sometimes rational to believe claims which one cannot prove. But which ones? A historically influential answer singles out two classes: claims which are self-evident, or obvious; and claims which your sense experiences tell you to be true.

This is the view which, in the third reading for today, Alvin Plantinga calls
foundationalism.



This is the view which, in the third reading for today, Alvin Plantinga calls **foundationalism**.



One way into this view begins with two candidate positive rules of belief:

Experience → Belief

If your sense experience tells you that P, and you have no reason to think that your sense experience is misleading, believe P.

Self—Evident → Belief

If P is self-evident, believe P.

Now recall the other positive rule of belief we discussed:

Proof → Belief

If you can prove P, believe P.

Experience \rightarrow Belief

If your sense experience tells you that P, and you have no reason to think that your sense experience is misleading, believe P.

Self-Evident \rightarrow Belief

If P is self-evident, believe P.

Proof \rightarrow Belief

If you can prove P, believe P.

The foundationalist says: these are the **only** cases in which you should form a belief. We can state this thought as follows:

No Foundations \rightarrow No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

No Foundations → No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

Foundationalism also seems to explain what is wrong with (serious) Pastafarianism. Given that there seem to be no good arguments in favor of the existence of the Flying Spaghetti Monster, we have no sensory evidence of its existence, and its existence is not self-evident, we should not be Pastafarians.

Could Foundationalism also be used as an argument against more standard forms of religious belief?

It can. (This is what Plantinga calls the 'evidentialist objection' to religious belief.)

No Foundations → No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

It can. (This is what Plantinga calls the 'evidentialist objection' to religious belief.)

1. No Foundations → No Belief.
2. We have no good argument for God's existence.
3. We have no sense experience of God.
4. God's existence is not self-evident.

C. You shouldn't believe that God exists. (1,2,3,4)

One might of course reject premise (2) of the evidentialist objection, if you found one of the arguments for the existence of God we discussed in class convincing. And you might reject (3) if you have had certain kinds of mystical experiences.

1. No Foundations → No Belief.
2. We have no good argument for God's existence.
3. We have no sense experience of God.
4. God's existence is not self-evident.

C. You shouldn't believe that God exists. (1,2,3,4)

One might of course reject premise (2) of the evidentialist objection, if you found one of the arguments for the existence of God we discussed in class convincing. And you might reject (3) if you have had certain kinds of mystical experiences.

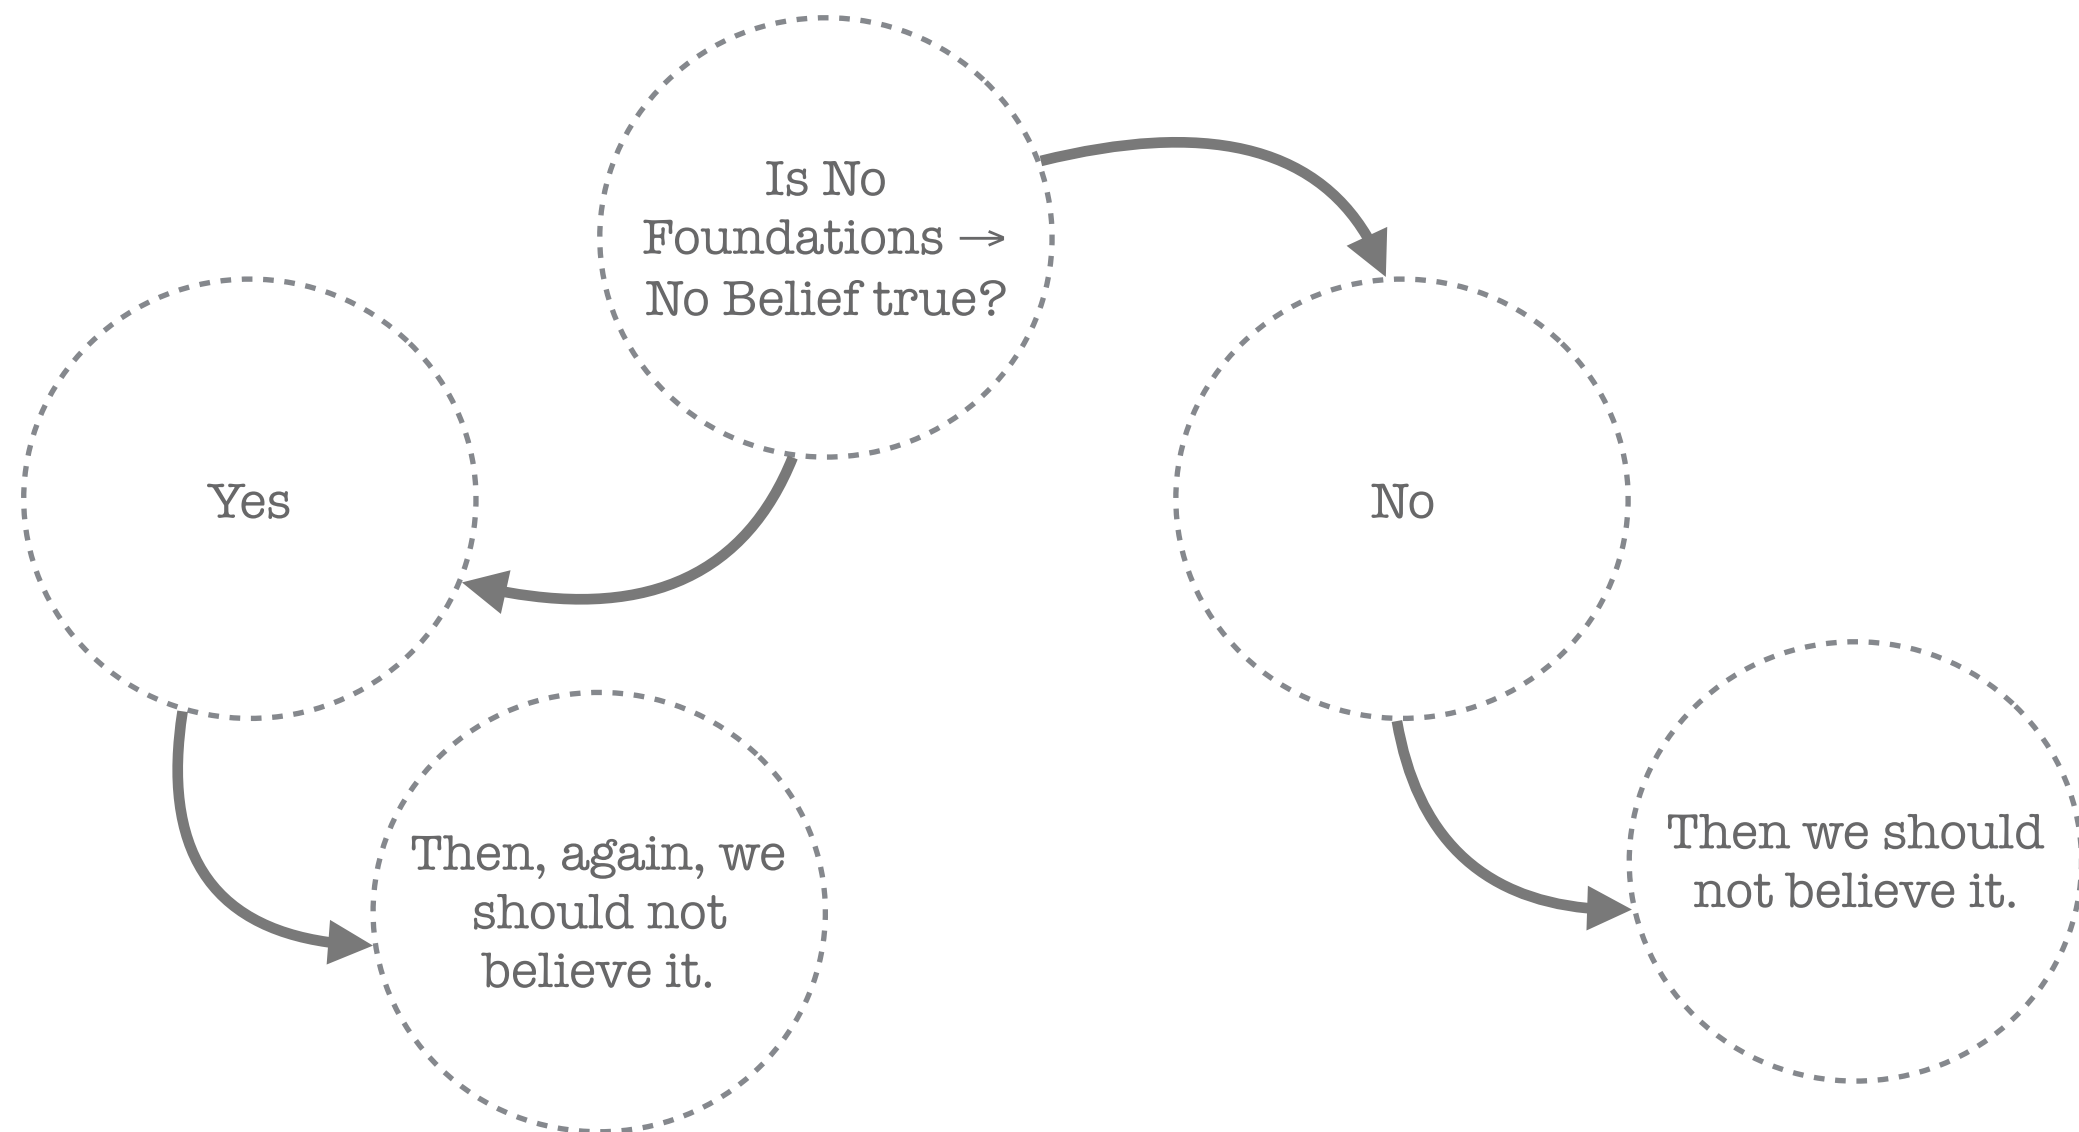
But set these aside for now. Our question is what you should do if you are in the position of The Believer — i.e., in a position where you find that you don't have a convincing positive case for some belief that you hold.

The key question is then: is our foundationalist rule of belief true?

No Foundations \rightarrow No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

Here is an argument by dilemma that we should not believe this principle.



No Foundations → No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

So we should not believe this principle.

But that principle was a premise of the evidentialist argument against belief in God:

1. No Foundations → No Belief.
2. We have no good argument for God's existence.
3. We have no sense experience of God.
4. God's existence is not self-evident.

C. You shouldn't believe that God exists. (1,2,3,4)


So, Plantinga concludes, the argument should be rejected.

No Foundations → No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

But even if this is an effective rebuttal to the evidentialist objection, it does not tell us whether Foundationalism is true or false. Plantinga's second argument is an attempt to show directly that Foundationalism is false.

This is based on the possibility that everyone besides you is a zombie.




This is based on the possibility that everyone besides you is a zombie.

A zombie (in the philosophical sense) is not a bloodthirsty undead monster.

A zombie is a creature who is externally indistinguishable from a human being, but lacks consciousness.

We can ask: how do you know that everyone besides you is not a zombie, in this sense? This question is sometimes called **the problem of other minds**.



Your senses don't tell you one way or another whether the person to whom you are talking is conscious. And it is not self-evident that the person is conscious.

So, if Foundationalism is true, it looks like we can know that other people are conscious only if we can give a good argument for the claim that they are conscious. Can we?

So, if Foundationalism is true, it looks like we can know that other people are conscious only if we can give a good argument for the claim that they are conscious. Can we?

Here is an argument you might give:

I know that I am conscious, and I observe that in my case there is a correlation between my conscious states and my outward bodily movements. But I also notice that the outward movements of the bodies of other people are similar to my own. So it is reasonable for me to believe that, just as there is a correlation between outward movements and conscious states in my case, so there is such a correlation in the case of other people. Hence it is reasonable for me to believe that they too are conscious.

This argument — which is sometimes called the [argument from analogy](#) — sounds plausible. But it faces a serious problem.

An **inductive argument** is an argument which generalizes from cases. Here is an example of an inductive argument:

1. The sun came up today.
 2. The sun came up yesterday.
 3. The sun came up the day before yesterday.
 -
 -
 -
 -

C. The sun will come up tomorrow.

Is this argument valid?

In general, inductive arguments are not valid — but it does seem as though they can give us good reason to believe certain claims which go beyond our sense experience.

The argument from analogy for the conclusion that other people are conscious seems to be an inductive argument: it generalizes from my own case to the case of other people.

But it is a very weird argument of this sort: it is induction **from a single case**. Is this sort of inductive reasoning a good way to reason? Compare the following:

Yesterday, I saw my first sushi roll. It had salmon in it. So, I think that all sushi rolls must have salmon in them.

This is pretty clearly a bad piece of reasoning. But then the question is: why isn't the inductive argument for the conclusion that other people are conscious just as bad?

But it is hard to see how we could argue that other people are conscious, other than on broadly inductive grounds.

No Foundations → No Belief

If P is not self-evident and your senses don't tell you that P and you can't prove P, don't believe P.

So it seems as though, if No Foundations → No Belief is true, we should not believe that other people are conscious. But that, Plantinga thinks, is very implausible. Hence, he thinks, this rule of belief should be rejected.

This is good news for someone who wants to oppose the evidentialist objection to religious belief. But it leaves us without the thing we wanted: some explanation of why Pastafarianism is irrational.

We have two different claims for which we lack good arguments: the claim that other people are conscious, and the claim that there is a Flying Spaghetti Monster. And yet it is reasonable to believe the first, but not the second. What explains the difference? (And which one, the traditional religious believer might ask, is the belief that God exists more like?)

Let's say, borrowing a term from Plantinga, that a belief which is not based on argument is a **basic belief**. We know that some basic beliefs are rational (like belief in other minds) and that other basic beliefs are irrational (Pastafarianism). Let's call a rational basic belief **properly basic**. Then our question is what makes some beliefs but not others properly basic.

This is a difficult question to answer. Here is what Plantinga says about it:

[one] can properly hold that belief in the Great Pumpkin is not properly basic, even though he holds that belief in God is properly basic and even if he has no full fledged criterion of proper basicity. Of course he is committed to supposing that there is a relevant difference between belief in God and belief in the Great Pumpkin, if he holds that the former but not the latter is properly basic. But this should prove no great embarrassment; there are plenty of candidates. ... [he] may concur with Calvin in holding that God has implanted in us a natural tendency to see his hand in the world around us; the same cannot be said for the Great Pumpkin, there being no Great Pumpkin and no natural tendency to accept beliefs about the Great Pumpkin.

Here Plantinga seems to be suggesting that a belief is properly basic if we have a natural tendency to believe it. This seems to put belief in God — though perhaps not specifically Christian belief — on the side of other minds rather than on the side of Pastafarianism, which is what Plantinga wants.

But of course this is — as Plantinga recognizes — too simple. After all, we seem to have a natural tendency to believe that the sun moves around the earth — but we can hardly rationally take that on board as a basic belief.

The reason why is obvious: we have a great deal of evidence that this belief is false. Let's call this evidence a **defeater** for the belief that the sun moves around the earth.

Then we might reformulate Plantinga's suggestion as follows: a belief is properly basic if we have a natural tendency to believe it, and it has no defeaters (or, if it does, that those defeaters are outweighed by reasons counting in favor of the belief).

This suggests the following rule of belief:

Inclination + No Defeaters → Belief

If you are inclined to believe P, and have (on balance) no defeaters for believing P, you should believe P.

Inclination + No Defeaters → Belief

If you are inclined to believe P, and have (on balance) no defeaters for believing P, you should believe P.

Our question, then, is whether belief in God has defeaters. And one might think that it does: one might think, for example, that the amount and kind of evil we find in the world is a defeater for the belief that there is an omnipotent and all-good being. In this case, we might think that belief in God is properly basic for children and adults who have never thought the problem of evil through, but not properly basic for intellectually sophisticated adults like the students in this class.

We've now found a candidate criterion to distinguish belief in God from belief in the flying spaghetti monster. The claim is that we have a natural tendency to believe in God and no defeaters for this belief (or defeaters that are outweighed). One might reject the idea that there are no defeaters for belief in God — but then one is arguing that religious belief is irrational, not because of a lack of evidence, but rather because there are arguments against it.

But we already knew that one could challenge the rationality of religious belief on **that** basis. That's consistent with claiming that there is no special problem for religious belief which follows from a lack of evidence for God's existence.

Inclination + No Defeaters → Belief

If you are inclined to believe P, and have (on balance) no defeaters for believing P, you should believe P.

Important questions remain. Does any inclination count? Plantinga talks about natural inclinations; but is it obvious that we have a natural (as opposed to socially influenced) inclination to believe in God? And why should the fact that someone is inclined to believe something be a reason to believe it?

Perhaps more importantly, we've failed to come up with a general negative rule of belief — a rule which tells you when not to believe something. You might think about how we could improve upon the foundationalist's efforts to provide such a rule.