What should I believe?

I should believe only what I can be certain of

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

This is a question in the field of philosophy known as epistemology — so called because ‘epistêmê’ is the ancient Greek word for knowledge.

Questions about what we should believe are bound up with questions about the scope of our knowledge, since it seems that one thing that we should try to do when forming beliefs is to obtain knowledge.

Today we return to a text familiar from our discussion of the nature of the self: Descartes’ *Meditations*.

Today we return to 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 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 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.

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

> As if I were not a man who sleeps at night, and regularly has all the same experiences[1] while asleep as madmen do when awake – indeed sometimes even more improbable ones. How often, asleep at night, am I convinced of just such familiar events – that I am here in my dressing-gown, sitting by the fire – when in fact I am lying undressed in bed! Yet at the moment my eyes are certainly wide awake when I look at this piece of paper; I shake my head and it is not asleep; as I stretch out and feel my hand I do so deliberately, and I know what I am doing. All this would not happen with such distinctness to someone asleep. Indeed! As if I did not remember other occasions when I have been tricked by exactly similar thoughts while asleep! 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. The result is that I begin to feel dazed, and this very feeling only reinforces the notion that I may be 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:

> 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.

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.

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.

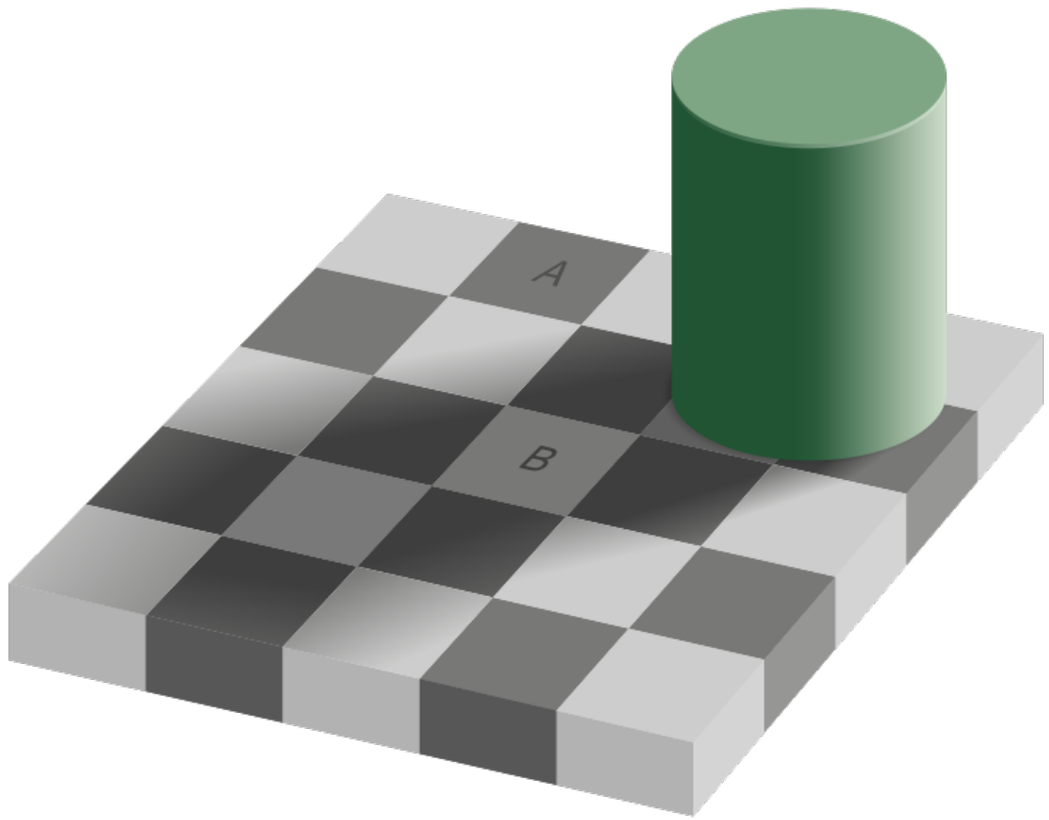

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

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

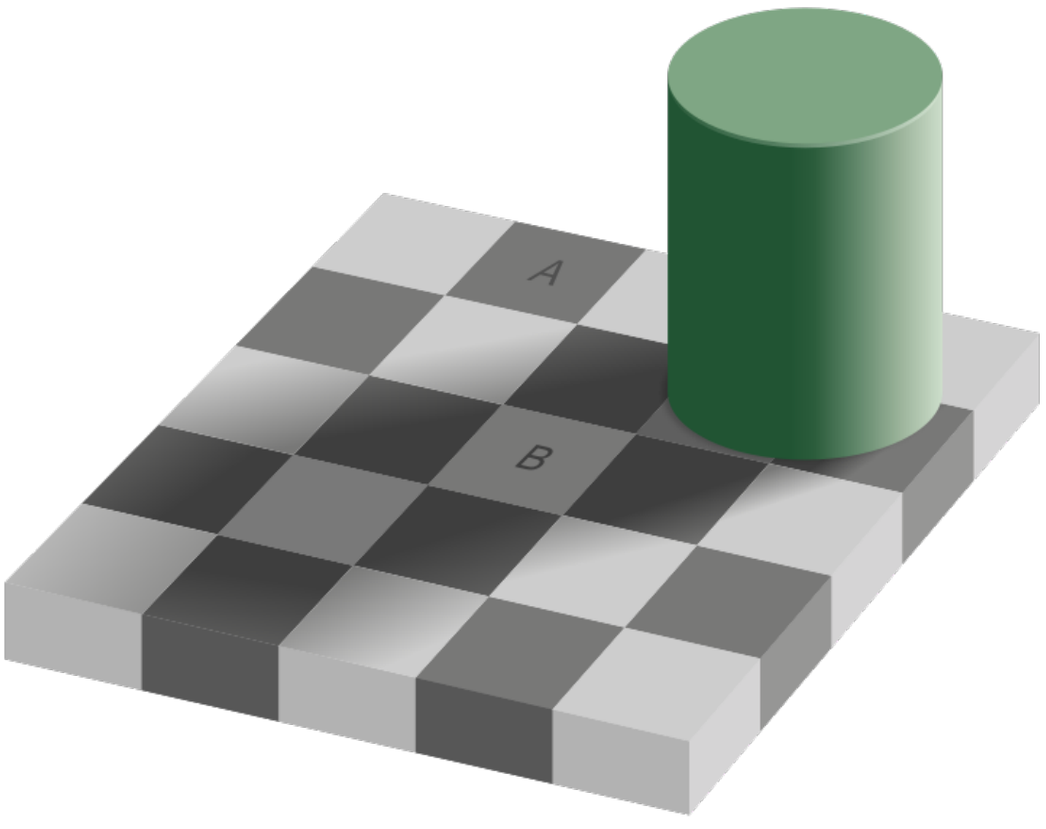

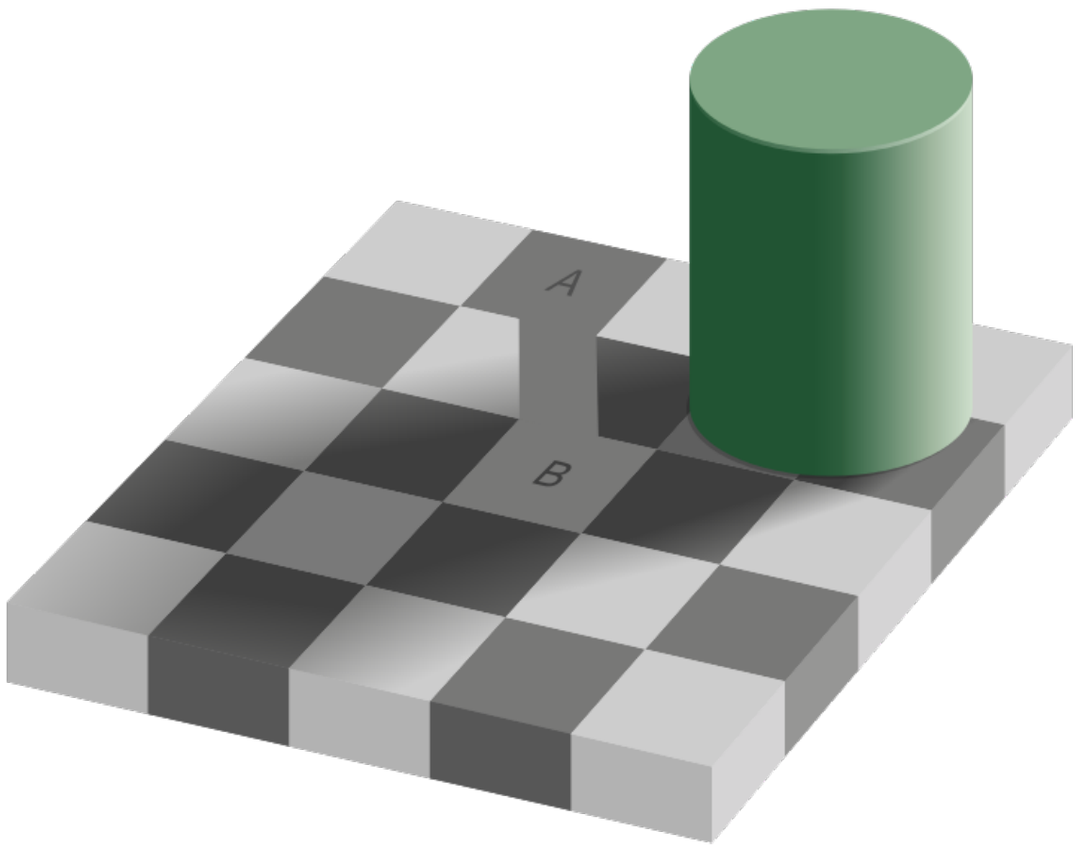

For any sense experience, I can imagine a situation which is indistinguishable from that sense experience but in which my environment is not as the experience says it is.

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.

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.

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.

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.

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.

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

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

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. Sense experience is my only way of knowing 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 premise (1), and premise (4) seems quite plausible; how about premise (2)?

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

———————————————-

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

It is important to see that this is not just an argument about what we can know. It can also, it seems, be turned into an argument about what we can reasonably believe.

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. Sense experience is my only way of knowing 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 it is not reasonable for me to believe that one but not the other is real.
3*. It is never reasonable for me to believe that any sense experience of mine is accurate. (1,2*)
4*. Sense experience is my only way of forming reasonable beliefs about whether there is an external world.

———————————————-

C*. It is not reasonable for me to believe that there is an 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 argue 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.

Indeed, as we saw in our second reading today, 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."

In response to this situation, the 20th century English philosopher G.E. Moore tried to show that we can, in fact, prove that there is an external world. (And that it turns out to be much easier than you might have thought.)

He presents his proof 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 three 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. Sense experience is my only way of knowing whether I have hands.

————————————————————————————-

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?

If your answer is ‘the second,’ then Moore would suggest that we respond to the skeptical argument by rejecting premise (2).

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 view — and the related view that if you do not have evidence for something, then you cannot know it — will be the subject of our discussion next time.