

Moore's proof of an external world &
the problem of other minds

Last time, we discussed two arguments for skepticism about the external world: arguments for the conclusion that we do not know anything about external objects.

In the reading for today, G. E. Moore, an English philosopher who worked in the first half of the 20th century, gives us an argument against skepticism of this sort: as against the skeptic, he claims that we can give a conclusive proof of the existence of an external world.

Moore begins his proof by explaining what an external object -- or, as he puts it, "a thing external to our minds" -- is. He says:

"to say of anything, e.g. my body, that it is external to my mind, means merely that from a proposition to the effect that it existed at a specified time, there in no case follows the further proposition that I was having an experience at the time in question."

So an external object is one which could exist without being perceived; and to prove the existence of external objects is to prove the existence of something which could exist without being perceived.

This suggests a two-part strategy for proving the existence of the external world: (1) find some kind of object which is such that if there is an object of that kind, then it can exist unperceived, and then (2) give an argument that there are objects of that kind.

Moore thinks that the first task is pretty easy. There are many, many kinds of objects which are such that any object of that sort would have to be able to exist unperceived. He uses the example of a soap-bubble:

e.g. consider the kind 'soap-bubble'. If I say of anything which I am perceiving, 'That is a soap-bubble', I am, it seems to me, certainly implying that there would be no contradiction in asserting that it existed before I perceived it and that it will continue to exist, even if I cease to perceive it. This seems to me to be part of what is meant by saying that it is a real soap-bubble, as distinguished, for instance, from an hallucination of a soap-bubble.

This, of course, does not show that there is an external world; it only shows that *if* there are soap-bubbles, then there are some external objects. We still need to figure out how Moore could manage stage (2) of his strategy for proving the existence of an external world.

So what Moore needs to do is prove the existence of some things which are such that, if they do exist, then external objects exist. His attempt to give such a proof is surprisingly simple:

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 is claiming to give a proof of the external world here, and a proof is just a certain sort of argument. So we should be able to separate out the premises and conclusion of his proof. Here is one way to think about it:

1. Here is one hand.
2. Here is another hand.

- C1. There are two hands. (1,2)
C2. There are external objects. (C1)

How can we tell whether Moore has succeeded in proving the existence of external objects? Moore himself gives us three conditions an argument must satisfy to be a genuine proof:

- (1) The conclusion must be distinct from the premises.
- (2) The premises must be known to be true.
- (3) The argument must be valid.

This gives us two questions to ask about Moore's proof: First, does it satisfy criteria (1)-(3)? Second, is meeting these criteria enough for an argument to count as a real proof?

Moore's argument clearly meets his first criterion: the conclusion is distinct from the premises. But does it satisfy (2)? Does Moore know that the premises are true?

He considers one objection to the idea that he knows the premises of his argument: he imagines an objector noting that Moore has provided no *proof* of the premises of his argument, and saying that it is impossible to know the premises of an argument unless you can prove them.

Moore grants that he has provided no proof of the premises of the argument, but is not convinced that this shows that he does not know them. One way to see the plausibility of this reply is to note that the objector is relying on some principle like the following:

If you can't prove something, you don't know it.

If the objection is going to be a good one, it seems that we must know that some principle of this sort is correct. But do we?

We will consider in a bit some further objections to the claim that Moore knows the premises of his argument; but the objection that he doesn't know them because he has not proven them does not look convincing.

Let's now consider criterion (3): the requirement that the conclusion of Moore's argument follow from his premises. Remember that the argument goes like this:

1. Here is one hand.
 2. Here is another hand.
-
- C1. There are two hands. (1,2)
C2. There are external objects. (C1)

It is uncontroversial that C1 follows from 1 and 2. But does C2 really follow from C1? Is it true, in other words, that if there are hands, they are the sorts of things which could exist without being perceived? Are they like soap-bubbles in this respect, or were we perhaps wrong to think that soap-bubbles could exist without being perceived?

We might try to give a proof of the principle that hands can exist without being perceived. For example:

- I have a hand right now (while hand is put behind one's back).
No one is perceiving my hand right now.
-
- Hands can exist without being perceived.

Is this a successful proof? Does it meet Moore's three criteria for a proof?

You might think that this proof, like Moore's, just goes wrong because we cannot know the premises. Just as Moore's proof assumed that we can know

Here is a hand.

This one assumes we can know

I have a hand right now (while hand is put behind one's back).

But these are exactly the sorts of things that the skeptic about the external world argues that we *can't* know. So Moore's proof just fails, since the skeptic has shown that I can't know, for example, that I have a hand.

Let's remind ourselves how a skeptical argument for such a conclusion would go.

One of the skeptical arguments we discussed last time was this one:

1. REALITY and DREAMWORLD are indistinguishable.
2. If I cannot distinguish between two situations, then I cannot know which of those situations obtains.
3. I cannot know which of REALITY and DREAMWORLD obtains. (1,2)

C. I cannot know whether the external world exists. (3)

We might adapt this argument to our purposes as follows:

1. A situation in which I am seeing my hand, and a situation in which I am hallucinating a hand, are indistinguishable.
2. If I cannot distinguish between two situations, then I cannot know which of those situations obtains.
3. I cannot know whether I am seeing my hand, or hallucinating a hand (1,2)

C. I cannot know whether I have a hand. (3)

Doesn't this just show that Moore does *not* know the premises of his argument?

It is interesting to see how Moore replies to the objection that he does not know the premises of his argument to be true. Here is what he says in defense of the idea that he does know the premises of his proof, and hence that his proof does satisfy his second requirement:

(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!

This looks a little bit puzzling. When called on to defend his view that he knows that he has hands, he seems simply to assert that he really does know this.

But there is a better way to understand what Moore is up to here.

Remember our skeptical argument:

1. A situation in which I am seeing my hand, and a situation in which I am hallucinating a hand, are indistinguishable.
2. If I cannot distinguish between two situations, then I cannot know which of those situations obtains.
3. I cannot know whether I am seeing my hand, or hallucinating a hand (1,2)

C. I cannot know whether I have a hand. (3)

Moore is insisting that it is one of our core beliefs about the world that the conclusion of this argument is false. Why might this be relevant?

A valid argument is one which is such that if its premises are true, its conclusion must be true as well. But in response to a valid argument, one need not believe the conclusion; one might also reject a premise. In fact, we have seen many examples of valid arguments to which this is clearly the correct response -- these were arguments by *reductio ad absurdum*, which had a clearly false conclusion.

So, in response to any valid argument, we can ask: does this argument show that the conclusion is true, or that one of the premises must be false?

One way to understand Moore's insistence that we do know that we have hands is that the above skeptical argument should be viewed as a *reductio*: it is a valid argument for a clearly false conclusion, and hence must have a false premise.

One way to press this point is to ask: are you more sure that you know that you have hands, or that the philosophical claim about knowledge expressed in premise (2) is true?

At this point, many people find the discussion a bit frustrating, and have something like the following response:

Look, maybe Moore's proof does meet his three conditions -- or at least I grant that there is no easy way to show that it does not meet those three conditions. But his argument would never convince a skeptic about the external world; the skeptic would just reject Moore's premises. But surely the point of an argument is to convince your opponent -- and doesn't this show that Moore's argument is a failure?

Could Moore's proof of an external world ever convince a skeptic?

Could Moore show that we should not be convinced by the skeptic? Are these the same thing?

This completes our discussion of arguments for skepticism about the external world, and Moore's response. We'll now turn to the second topic Nagel discusses: skepticism about our knowledge of other minds.

Nagel puts the argument as follows:

If we continue on this path, it leads finally to the most radical skepticism of all about other minds. How do you even know that your friend is conscious? How do you know that there are *any minds at all* besides your own?

The only example you've ever directly observed of a correlation between mind, behavior, anatomy, and physical circumstances is yourself. Even if other people and animals had no experiences whatever, no mental inner life of any kind, but were just elaborate biological machines, they would look just the same to you. So how do you know that's not what they are? How do you know that the beings around you aren't all mindless robots? You've never seen into their minds—you couldn't—and their physical behavior could all be produced by purely physical causes. Maybe your relatives, your neighbors, your cat and your dog have *no inner experiences whatever*. If they don't, there is no way you could ever find it out.

In a way, this is similar to the dream argument: Nagel points out that all of the beings which surround you could lack consciousness, and that this scenario is one which is indistinguishable

But in another way, this is more challenging than the dream argument. The dream argument relies on the possibility of a certain kind of very comprehensive hallucination -- it relies on the possibility of our senses deceiving us in every apparent case of perception of an external object. But the problem about knowledge of other minds arises *even if we assume that all of our sense perceptions are perfectly accurate*. After all, we never perceive the pain, or taste sensations, of another person; we only perceive the physical reactions which we take to indicate the presence of those sensations.

We can put the corresponding argument for skepticism about other minds as follows:

1. Perception is our only evidence for the existence of other minds.
2. We never directly perceive other minds, but only other bodies.
3. Our perception of other bodies is our only evidence for the existence of other minds. (1,2)
4. If we have no good evidence for the existence of other minds, we cannot know that there are other minds.
5. If perception of other bodies does not give us good evidence for the existence of other minds, we cannot know that there are other minds. (3,4)
6. Perception of other bodies does not give us good evidence for the existence of other minds.

C. We cannot know that there are other minds. (5,6)

One key premise of this argument for skepticism about other minds is the last one:

6. Perception of other bodies does not give us good evidence for the existence of other minds.

But why should we believe this? One might reason as follows: I observe in my own case that there is a correlation between certain physical events and certain mental events. So, it is reasonable for me to believe that in every case (or in the case of every creature 'similar enough' to me) there is such a correlation between these physical events and the corresponding mental events. So, when I observe a physical event of the right sort happening to someone else, this gives me good evidence that a mental event of the relevant sort is also going on.

This is an example of reasoning which begins from knowledge of particular cases and moves from this knowledge to a generalization, a claim about what happens in all cases. This is one kind of *inductive reasoning*, and we will talk about such reasoning in more depth in a few classes.

But whatever we think about such reasoning in general, the sort of reasoning above seems to be an especially bad, and irrational, form of inductive reasoning. Can you see why?