

Skepticism about the external world & the
problem of other minds

So far in this course we have, broadly speaking, discussed two different sorts of issues: issues connected with the nature of persons (a topic in metaphysics), and the existence of God (a topic in the philosophy of religion). Today we begin a new section of the course. In this section, we'll be discussing the nature and scope of human knowledge. This area of philosophy is called *epistemology*.

For the next three classes we will be discussing the question: What sorts of things, if any, can we know?

Someone who claims that we know quite a bit less than we usually take ourselves to know is a *skeptic*. In today's reading, Nagel presents several skeptical arguments.

But it is important to be clear at the outset about a difference between two different ways in which the word "skeptic" or "skepticism" is used. In the first sense, a skeptic about X is someone who believes that there is no such thing as X -- so, for example, in this sense a skeptic about God is someone who believes that there is no such thing as God.

This is *not* what we mean by "skepticism." In the sense that we are interested in, a skeptic about God is someone who denies that we know that God exists. This is a less ambitious position -- you might be an agnostic, for example, and take no stand on the existence of God, but still hold that we have no knowledge of God's existence.

The first sort of skeptical argument he discusses is an argument which aims to show that we can have no knowledge of the external world: no knowledge about the existence and natures of things outside our own minds. So, as Nagel says, a skeptic of this sort is someone who thinks that we can't know about "the existence of the floor under your feet, or the tree outside the window, or your own teeth." All we can know, on this sort of view, are facts about the way things seem or appear to us -- facts about our own sensations and other mental phenomena.

This sort of skepticism is quite a surprising view. Most of us think that we know that we have teeth, and, indeed, that we know some stuff about them -- that they are more or less white and used to chew food, for example. So skepticism about the external world is the sort of view that we should only accept if we are given a plausible argument.

And that is what Nagel aims to give us. He gives two distinct, though related, lines of argument in favor of skepticism about the external world. One of them is well-captured by the following passage:

Would things seem any different to you if in fact all these things existed *only* in your mind—if everything you took to be the real world outside was just a giant dream or hallucination, from which you will never wake up? If it were like that, then of course you *couldn't* wake up, as you can from a dream, because it would mean there was no “real” world to wake up into. So it wouldn't be exactly like a normal dream or hallucination. As we usually think of dreams, they go on in the minds of people who are actually lying in a real bed in a real house, even if in the dream they are running away from a homicidal lawnmower through the streets of Kansas City. We also assume that normal dreams depend on what is happening in the dreamer's brain while he sleeps.

But couldn't all your experiences be like a giant dream with *no* external world outside it? How can you know that isn't what's going on? If all your experience were a dream with *nothing* outside, then any evidence you tried to use to prove to yourself that there was an outside world would just be part of the dream.

This is, you might notice, reminiscent of an argument from the selection from Descartes that we read the first week of the semester. Descartes asked how we could know that we are not dreaming; and this is much the same question that Nagel asks us to consider in this passage.

But why, you might ask, is this supposed to be such a hard question to answer?

Nagel's answer is evident from question which begins the passage we just read. He asks: *Would things seem any different to you if in fact you were having a very long and detailed dream?* His implied answer is: *No. We can imagine having a dream which was, in every way, indistinguishable from our present experiences.*

The question we now want to ask is: if Nagel is right about this, why would this indicate that we have no knowledge of the existence of an external world?

Here is one way to turn the sort of example he discusses into an argument for skepticism. Let's use REALITY and DREAMWORLD as names for the following two possible scenarios: REALITY is a world in which I have just the experiences I have had, and these more or less accurately reflect a world of external objects; DREAMWORLD is a world in which I have just the experiences I have had, but that these are a part of a long and detailed dream, and so do not reflect the existence of any world of external objects.

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)

Let's call this the *dream* argument.

Given the definitions of REALITY and DREAMWORLD, the dream argument appears to be valid. Nagel, in effect, gives us an intuitive argument for premise (1). If we accept this premise, then premise (2) is the only other premise in the skeptical argument open to question.

How plausible is premise (2)? Can you think of any examples which either support premise (2), or seem to be counterexamples to it?

Nagel also says some things which suggest a different argument for skepticism about the external world. Here is one relevant passage:

It seems that you are stuck with nothing you can be sure of except the contents of your own mind at the present moment. And it seems that anything you try to do to argue your way out of this predicament will fail, because the argument will have to assume what you are trying to prove—the existence of the external world beyond your mind.

Here Nagel is interested less in questions about indistinguishability than in whether or not we can provide a certain kind of *argument* for the existence of the external world.

The idea seems to be this: we have some beliefs that are secure, that we can be sure of, such as our beliefs about what we are experiencing right now. Let's call these the *basic beliefs*. Then, if we want to see whether some other belief is knowledge, we see whether we can give an argument for that belief which uses only basic beliefs as premises.

Let's consider the case of the external world. We have some basic beliefs about our experiences, like my belief that I am now having a visual experience of what appears to be many freshman philosophy students. How might we argue from beliefs of this sort to claims about external objects?

Nagel suggests that any attempt to construct such an argument will involve some assumptions about the external world. Here is what he says:

If you try to argue that there must be an external physical world, because you wouldn't see buildings, people, or stars unless there were things out there that reflected or shed light into your eyes and caused your visual experiences, the reply is obvious: How do you know *that*? It's just another claim about the external world and your relation to it, and it has to be based on the evidence of your senses. *But you can rely on that specific evidence about how visual experiences are caused only if you can already rely in general on the contents of your mind to tell you about the external world. And that is exactly what has been called into question.* If you try to prove the reliability of your impressions by appealing to your impressions, you're arguing in a circle and won't get anywhere.

Nagel is imagining us trying to provide an argument for the existence of an external world, using as premises only claims about our own sensations. It is pretty clear, he thinks, that we would not get anywhere; at some point, we would need to use a premise which connects sensations to external objects, like

Sensations of red are typically caused by external objects which are red.

But, as he points out, this is not one of our basic beliefs; so it can't be used in an argument which uses only basic beliefs as premises for the conclusion that there is an external world.

We might put this sort of skeptical argument like this:

1. Every proposition we know must either be a basic belief -- a belief that we are sure can't be mistaken -- or the conclusion of a valid argument which has only basic beliefs as premises.
2. The existence of the external world is not a basic belief.
3. There is no argument for the existence of an external world which uses only basic beliefs.

C. We cannot know of the existence of an external world. (1,2,3)

Let's call this the *No-argument argument*. Like the dream argument, this appears to be valid. How plausible are the argument's premises?

So we have two kinds of arguments for the conclusion that we have no knowledge of the external world -- the dream argument, and the no-argument argument.

In the second half of the reading, Nagel argues that even if we just take for granted the existence of an external world, a further very serious skeptical problem arises. This is the problem of explaining how we could have knowledge of the existence of other minds -- the existence of sensations and other mental events belonging to creatures other than myself.

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?