

So, let's suppose that after all this discussion of what knowledge is, we end up saying, well, actually we do want to use the word knowledge in a way that reflects our use in ordinary language. We want to allow that I can know the train is time to leave at seventeen thirty six, even though I might be dreaming. Maybe we want to allow the contextualist to apply different levels of knowledge in different circumstances, and we certainly don't want to insist that only super-rational beings who reflect on all their beliefs and justifications can know things. Quite the reverse, we want to accept that dogs and cats can know things. We're happy to say that ordinary unreflective people can know things when there are reliable connections between their faculties and the truth, even if they can't explain what those connections are. So, this naturally leads us back to externalism as an attractive account of knowledge. This enables me to claim that I do know this is a hand, even if I can't prove it, and even if I can't know that I know it. Whether I do in fact know it depends on how things stand outside my mind - the various causal links between the world and my perceptions and so forth. As long as they're working fine, then I can have knowledge, just as the dog can have knowledge. I don't have to be an expert philosopher or an expert in human perception for my perceptual faculties to operate correctly and give me knowledge.

But the skeptic is still lurking in the wings. Let's suppose we accept all I've said. Suppose we accept that the word knowledge, as it's used in ordinary language, fits with this sort of externalist account and can quite properly be used in the various loose ways I've described. That doesn't actually defeat the skeptic because the skeptic can say, well look, if what you say is right, if your beliefs are, in fact, true, then I'll accept that you know all these things in this ordinary language sense. But I still challenge you that those beliefs might be false. So even if skepticism can be answered from a God's-eye externalist point of view - God can look down and say, Millikan's faculties are working fine, so he does, in fact, know that there's a hand there - the question whether they're actually true can still be asked from the internal perspective. I can't necessarily know that I know. So, I can still raise skeptical problems about things that, from a God's-eye point of view, I do supposedly know. So, is there any answer to this kind of skepticism? Well, surprisingly perhaps, one rather prominent answer, which aims to show that we can be confident of some of our basic perceptual beliefs, has come from the direction of the analysis of ordinary language.

So suppose I refer to this and I call it a hand, and I conclude that there are two hands. Are they really? Are they really hands, you might think? No, for all you know, you're a brain in a vat, you're dreaming, whatever. Well, okay, let's suppose I am, in fact, a brain in a vat. Okay, I'm a brain in a vat. I look at this, I look at this, I think there's something there. Let's not worry about where there is. I'm aware of something, and I call this a hand. And if these are actually hand images, then when I use the word hand to refer to them, I'm referring to hand images. Okay? But if when I say hand I mean a hand image, then this is a hand after all, even if I'm a brain in a vat. This, which I call a hand, is a hand. Good. Maybe from a God's-eye point of view, they're just hand images. From my point of view, that's what I call a hand. So, if even if I am a brain in a vat, I can say with G.E. Moore, "There's a hand! There's another!" If the meaning of hand is determined by what we're actually referring to, it looks like the skeptic can be defeated. Or at least I need not worry about whether this is really a hand. This must be a hand because that's what I mean by hand. And that's the kind of approach that Putnam suggests.

Now, you might well think that's a little bit too quick. It's not really a very satisfactory answer to the skeptic. I think that worry is right. Here's how I might spell it out. When I look at this thing, I think there's an object there which is actually moving in space, and whose movement is

systematically correlated with my perception in such a way that my perceptions give a directly reliable indicator of where it is. I have an idea of the sort of causal interaction which is responsible for these perceptions, in terms of light shining on my hand, bouncing off, I see it with my eyes, and so forth. And that's a very different picture from the picture of some mad scientist manipulating electrodes or running some computer program which is bringing it about that my perceptions correlate as though there were an object there. So maybe I can make some sense of a God's-eye point of view from which it would turn out that what I call a hand is really nothing like what I take to be a hand. In that case, the Putnam approach can be challenged. It won't follow that when I say there's a hand, I can be utterly sure that that's true. If what's really there is nothing that bears any systematic correlation to what I'm perceiving, except through the manipulation of some mad scientist. And there's another problem with Putnam's approach. Let's step back from the vat for a moment and return me to real life. Okay, I know what a hand is. There's a hand. I'm walking along in Oxford one day, on my way to a lecture, and I get kidnapped and invited. Some mad scientist extracts my brain and puts it in a vat. I forget about all this, of course. I'm given the illusion of coming to a lecture. I look at this and I say, "Here's a hand." But actually, it's just a hand image. And now, it looks like Putnam's approach isn't going to work because I learned the use of the word hand by referring to real hands. So, when I say hand, I mean a real hand. I don't mean a hand image. In which case, I can raise the skeptical worry. Maybe this isn't a real hand. Maybe I am a brain in a vat. So, needless to say, the bogey of skepticism comes back, as indeed it usually does. There's no magic bullet to defeat the skeptic, and at least Putnam hasn't given us one.

Finally, I want to go back to induction. The Putnam approach might lead you to the following thought. Leaving aside the worry about being kidnapped and invited, suppose we worried that our whole life is lived in a vat or in the matrix or something like that. It's tempting to think, well, suppose I am a brain in a vat, suppose I am in a matrix. Why should I care? I'm living my life perfectly well. Maybe it is in the matrix, but that doesn't stop me enjoying the things that I enjoy. It doesn't stop me getting satisfaction from the company of matrix people, eating nice matrix food, or whatever. You might, however, wonder where this leaves issues about moral obligation to those matrix people. But let's put that to one side. All these matrix things, they may not be real, but they bring me the same pleasure. So why worry about it? Why not just go on as before, even if I am a brain in a vat, even if I am living in a matrix? Now, seeming like this vertical skepticism, that is worrying about inference from one level - the level of perception - to some deeper level - the level of objects - can seem not so worrying after all. And there's a contrast here with horizontal skepticism, the kind of skepticism that you get in the problem of induction, and that will remain even if you're happy with living in the matrix or as a brain in a vat. Everything so far might have gone on fine, but how can I be confident that it will carry on going on fine? This problem, the problem of induction, arises whether I'm in the real world or in a matrix or a brain in a vat. So there's a sense in which horizontal skepticism, that is skepticism about inferring more of the same, though it seems less radical, is actually potentially more worrying than vertical skepticism.

And this gives me an excuse, having gone back to human induction, to look at the kind of response that Hume gave, and is commonly given now, to these kinds of skeptical worries. And that is to focus on the ethics of a belief. What should we believe? Descartes started out his skepticism saying that he shouldn't believe anything that is less than certain. The message of all the discussion about skepticism is that if we do determine ourselves not to believe anything that is less than certain, we might end up believing virtually nothing at all. But is that possible? Are we able to believe nothing at all, and should we even go along with it? Why should we condemn ourselves to believing nothing that is less than certain? Well, I think most philosophers would agree with Hume that suspension of all belief is just impossible for us. The way we're made, we

just cannot help believing certain things. And it's probably a good thing that we're made that way because if we weren't, then we'd be in serious trouble. Notice also that this approach goes well with contemporary externalism. The thought is that we shouldn't aim for all our beliefs to be such that we can justify them internally. We shouldn't expect to be able to work out internally the justification for everything we believe. Perhaps we have to rely on our animal nature that leads us inevitably to believe certain things and to trust, in general, that our faculties are thankfully more or less reliable. Of course, that doesn't mean that we should become indiscriminating, and there remain big questions about how to distinguish between things that remain justified and things that aren't. But if we want to hold out against the skeptic, we probably have to be prepared to accept standards that are less than absolute. Thank you.