

Okay, so how might we begin to pin down what we mean by knowledge that *P*? What is it for somebody, let's call them *S*, the subject, to know a proposition *P*? Well, the standard traditional analysis is to say first of all, *P* has to be true. You can't know a proposition that's false. You might think you know it, but if actually it's false, you don't know it. Secondly, you have to believe it. You can't be said to know something you don't even believe. And thirdly, you have to be justified in believing it. So, those are the three standard conditions, sometimes called the JTB analysis (justified true belief). A.J. Ayer, who is particularly well known in this connection, gave the last two conditions slightly differently. He said that to know something, you have to be sure of it and you have to have the right to be sure. Then, you might think that's rather preferable. If I just vaguely believed something without any strong commitment, that might not be enough for knowledge. Suppose something to my mind has a 60% probability, "Oh, I reckon it's going to rain tomorrow," something like that. And maybe it's justified because I've seen the weather forecast. But is that knowledge? Probably, we'd say it isn't. That kind of weak belief isn't enough. You've got to be sure in order for it to count as knowledge. And you have to have the right to be sure. At any rate, that seems quite plausible. But let's ask some further questions about all this. If somebody knows that *P*, does it actually follow that *P* must be true? Notice that there are two slightly different claims that might be made here. When I say, if *S* knows that *P*, *P* must be true or if *S* knows that *P*, *P* is necessarily true, that's an ambiguous claim. On the one hand, I might be saying that if *S* knows that *P*, it follows that *P* is necessarily true. But *P* is a necessary truth. But it might be tempting to think that, but it's just wrong. I know that I exist, but that I exist is not a necessary truth. It's not like  $1 = 1$  or  $2 > 1$ . I could easily not have existed. Once I didn't, some time, I won't. But I do know that I exist right now, don't I? So, it's simply not true, and it can't follow from the meaning of know that the only things you can know are necessary truths. The second interpretation of this claim, however, looks much more plausible. Necessarily, if *S* knows that *P*, then *P* is true. Something can't count as a case of propositional knowledge unless the proposition in question is true, and that seems right. But do be ever so careful when you use words like 'necessarily' in philosophy. Always be careful to watch for the scope of the modal operator. So, distinguish between the two versions there. In the second case, you see 'necessarily' lies outside the brackets, whereas in the first case, if you were to put brackets, 'necessarily' lies inside. So talking about the scope of the modal operator sounds terribly technical, but be very aware that word orderings of this sort can matter a great deal in philosophy. Let's now ask, is it in fact the case that if somebody knows some proposition, that proposition must be true? Well, supposing I say, "I know that France is hexagonal. I know that France is hexagonal and Italy is shaped like a boot." But of course, France isn't hexagonal. So, there we are, I know a falsehood. Now, that seems to me to be just a confusion. There is a sense in which France is hexagonal, namely that it's roughly hexagonal. In that sense, I can know that France is hexagonal. There's another sense, a precise sense, in which France isn't anything like hexagonal. And in that sense, I can't know that it is. So, as long as we're clear about what we mean by hexagonal, exactly hexagonal or roughly hexagonal, that straightens out the problem. For a different case, if you want to get your philosophy tutors really cross, say something like this, "Well, that might be true for you, but it's not true for me." Philosophers hate that. Often, people say that kind of thing when all they really mean is, "I believe *P*, but you don't believe *P*." Some people will say, "It was true in the medieval period that the Sun orbited around the Earth." No, it wasn't. It wasn't true. It was universally believed, but it was never true. Probably, it was quite a reasonable belief at that time, but that didn't make it true. Take another example. It's true, we think, that the continents drift over centuries, millennia. Nobody believed that. But it was, in fact, true. Indeed, if you didn't think it was true way back in the time of the dinosaurs and so on, you're going to have a very difficult job explaining the distribution of fossils around the world.

That distribution only makes sense in the context that continents were moving, even when nobody believed it. So don't confuse "P is true" with "everyone believes that P." They're quite different. Let's move on. I'm going to take it for granted from now on that you can't know a falsehood. What about belief? In order to know that P, do you have to believe P? This isn't so clear. Suppose, for example, I'm in a quiz and I'm asked, let's say, to compare a list of capital cities and the countries, and I have to say which city is the capital city of each country. And suppose I say, "I don't know. I haven't a clue. I don't know any of these." And then somebody says, "Go on, have a go. Have a guess." Well, I really don't know. It's random. "Go on, guess." And I do, and I get them all right. Now, there might be some very plausible explanation of that. It might be that when I was at school, I learned all these things. Maybe I had a teacher who was really keen on capital cities. And fortunately, the test was confined to countries that existed when I was at school. And although I've completely forgotten all those lessons, I'm able to do the test accurately. I didn't believe that I knew, but actually, I gave the right answers. And you can imagine the quiz master saying afterwards, "Millikan didn't think he knew any of those, but in fact, he did know them." Now, one way of dealing with this kind of problem is to say, "I had unconscious knowledge of the capital cities." And in that same sense, you might want to say, "I had unconscious belief." We can flesh this out a bit. And suppose somebody's putting one of those quizzes. They claim to be guessing. They make their guesses, and all their guesses come out wrong. But then you put them in the quiz again. They guess again. They come out wrong again. But they come out wrong consistently. They're still getting the same thing. Well, clearly, they don't know the answers because they're getting them wrong. But you might want to say they've got a non-kind of unconscious belief because they're consistently giving the same answers. And that could give you a way of saying that knowledge does require belief, even if you accept that in the case of the quiz, I do have knowledge. You could say, "Yes, you can have unconscious knowledge, but only if you've got unconscious belief, only if your answers display a kind of consistency." The related case is the strange phenomenon called blindsight. This is where someone has damage to their brain in such a way that they have no conscious awareness of seeing anything. And yet, if you ask them to point to things, they can do it with reasonable reliability. So, if you take them into a room, spin them around and say, "Point to the desk." They'll say, "I can't see a thing. How do you expect me to point to the desk?" And you say, "Go on, yes." And they get it right. Much better than chance. So, again, with this sort of case, it might be tempting to say they know something, even though they don't believe it. Here, the idea of unconscious belief may seem a little bit less plausible because there's no long-term memory or consistent pattern underlying the behavior. A related problem: It's tempting to think that knowledge must be completely conscious. That if you know something, you must know that you know it. So, at one extreme, you've got the question, "Do you even need to believe it?" At the other extreme, you've got people who might want to say that if you know it, not only do you have to believe it, you have to know that you know. Well, that at least must be wrong. Suppose it's true that in order to know something, I have to know that I know it. In that case, if I know that P, I have to know that I know that P. And that means I have to know that I know that I know that P. Which means that I have to know that I know that I know that I know that P. And obviously, there's no stop to that. In order to know that P, I have to know that I know that I know, and so on. Even when there are Googleplex knows in that sentence, something I'm obviously completely incapable of understanding, let alone believing. There's no way that you can possibly believe an infinite number of those propositions. And I personally lose grip on what's being said once you've got five or six knows in there. So, insisting on consciousness all the way through, that all knowledge has to be self-reflective and known to be known, is just not going to work. Okay, I'm going to put the question of belief to one side now and move on to the key condition

that attracts most of the attention. Because the main point in distinguishing between just having a belief that P and knowing that P seems to be to focus on justification. The main reason we want a concept of knowledge is to distinguish between things that we really know and things that we think we know. So, let's just assume that P is true here. What is it that makes the difference between truly believing that P and knowing that P? Suppose I have a true belief that P. That alone surely isn't enough to imply that I know it. It might be just a lucky guess. It might be something I've been told by someone who's actually completely unreliable. Yet, on this occasion, they just happen to have told me a truth. There are all sorts of ways that I can have a true belief and yet it fails to be knowledge because it's not justified. So, in Ayer's terms, if I'm to know that P, I must have a right to believe it. Or a right to be sure of it. At least, that seems very plausible. Here, the shadow of skepticism can come back to haunt us. Suppose I believe that P. What is required for this to be knowledge? Well, plausibly, it has to be justified. No doubt, it will be justified in terms of other beliefs. If you asked, "How do you know that P?" I'm likely to have to appeal to other beliefs, call them Q and R. And then, the question can arise, "Okay, how can you justify Q and R?" And they, in turn, require further justification, perhaps in terms of S and T. And so, we go on. It looks like there's a threat of infinite regress here. Everything has to be justified by reference to something else. How can you stop it? Well, there are two different approaches that are traditionally taken. Indeed, it's difficult to see how else you could do it. One of them is coherentism, which simply says, "You have this web of interlocking beliefs, and if they all cohere together strongly enough, then they can justify each other." The other approach, foundationalism, says that ultimately, you hit rock bottom. Ultimately, you get to certain things that you can just know directly, without there having to be justified by anything else. Maybe "one equals one," for example. How do you know "one equals one?" Well, I just see it to be true. I don't have to justify that in terms of anything else. How do you know you're thinking? Well, I just see it to be true. Though Descartes was a foundationalist, he thought some beliefs were just totally secure in and of themselves. A more modern approach to halting the regress of justification, which moves away from the emphasis on conscious justification, which we've seen is probably not something that we want to insist on right the way through, is called externalism. Externalism features in a lot of modern approaches to philosophy. An internalist counter justification is one that requires that all the relevant factors that play a role in assessing a belief as worthy of being called knowledge must be cognitively accessible to the subject. Nothing can be hidden. An externalist doesn't require this. An externalist, by contrast, will say that some factors that are relevant to judging whether you're justified in believing something may be inaccessible to you, external to you. Now, one obvious advantage of this is the following. Suppose you want to say that dogs can know things or cats can know things. Well, that better not be in terms of some intellectual justification the dog or the cat is able to provide. So, how does the dog know that there's a cat nearby? Well, it smells it. Can the dog give an explanation of that? Of course it can't. What makes this knowledge is that there's a reliable causal connection between the smell and the existence of a cat nearby and the dog's detecting the smell and the smell actually being there. So, you've got a reliable causal link between the cat and the dog sensing the presence of a cat. Now, it's quite tempting to say that the key to knowledge in this sort of case is not intellectual justification, not reflection, but simply the existence of that reliable causal link and the dog having reliable faculties. Well, why not say the same about us? Why not say that we are capable of knowing things without necessarily knowing how we know? Because we've already seen, after all, that we don't want to say that you can only know that P if you can know that you know that P. Why not go the whole hog and say you don't even have to be able to give an account of how you know that you're justified?