

Royal Institute of Philosophy

A-Level Philosophy
Epistemology Resources*

Duncan Pritchard
University of Edinburgh

* I am grateful to James Garvey, Anthony O'Hear, and Mat Carmody for helpful feedback on earlier versions of this material.

*Table of Contents**Introduction**Topic 1: Perception*

- ❑ *The Problem of Perception*
- ❑ *Indirect Realism*
- ❑ *Idealism*
- ❑ *Solipsism*
- ❑ *Naïve Realism*

- ❑ *Topic Summary*
- ❑ *Study Questions*
- ❑ *Introductory Reading*
- ❑ *Further Reading*
- ❑ *Free Internet Resources*

Topic 2: The Definition of Knowledge

- ❑ *Knowledge, Truth and Belief*
- ❑ *Knowledge versus Mere True Belief*
- ❑ *The Classical Account of Knowledge*
- ❑ *The Gettier Problem*
- ❑ *Responding to the Gettier Problem*

- ❑ *Topic Summary*
- ❑ *Study Questions*
- ❑ *Introductory Reading*
- ❑ *Further Reading*
- ❑ *Free Internet Resources*

Topic 3: The Origins of Concepts and the Nature of Knowledge

- ❑ *Two Distinctions: A Priori/A Posteriori & Analytic/Synthetic*
- ❑ *Knowledge Empiricism*
- ❑ *Knowledge Innatism*
- ❑ *Concept Empiricism*
- ❑ *Concept Innatism*

- ❑ *Topic Summary*
- ❑ *Study Questions*
- ❑ *Introductory Reading*
- ❑ *Further Reading*
- ❑ *Free Internet Resources*

General Resources

- ❑ *Textbooks*
- ❑ *Edited Collections*
- ❑ *Free Internet Resources*

Glossary

INTRODUCTION

Although **epistemology** is one of the core areas of philosophy, the term itself only came into existence relatively recently (it was coined by the 19th Century Scottish philosopher, James Frederick Ferrier (1808-64)). Nonetheless, even though the term is quite new, the cluster of philosophical questions that it describes go right back to antiquity, to the birth of philosophy itself.

So what is epistemology? Well, it is essentially the theory of knowledge. Epistemologists are concerned with such questions as ‘What is the nature of knowledge?’, and ‘How is knowledge acquired?’ They are also interested in a range of concepts which are closely related to knowledge, such as truth, rationality, understanding, and wisdom.

Our epistemological focus will be on three general topics. The first, covered in part one, will be the nature of perception and its bearing on perceptual knowledge. The second, covered in part two, will be the definition of knowledge. Finally, the third topic, covered in part three, will be concerned with where our ideas and knowledge come from.

Each part closes with a summary of the points covered, a list of study questions, and an annotated catalogue of further readings (where these are in turn split up in terms of introductory readings, further readings, and free internet resources). At the end there is also a list of general resources in epistemology and a glossary of key terms (you’ll know when a term appears in the glossary, since it will be in **bold** the first time you come across it).

TOPIC 1:
PERCEPTION

- ❑ *The Problem of Perception*
- ❑ *Indirect Realism*
- ❑ *Idealism*
- ❑ *Solipsism*
- ❑ *Naïve Realism*

The Problem of Perception

A great deal of our knowledge of the world is gained via perception—that is, via our sensory faculties such as our sense of sight, hearing, touch, and so forth. My knowledge, if that’s what it is, that I am presently at my desk writing these words is itself largely perceptually gained. I can see the computer before me, and I can feel the hard touch of the computer keyboard on my fingers as I type. If we know much of what we think we know, then we must have a great deal of perceptual knowledge. As we will see, however, it is far from obvious that we do have widespread perceptual knowledge of the world around us, at least as that knowledge is usually understood.

Part of the problem is that the way things look isn’t always the way things are; appearances can be deceptive. There are familiar examples of this sort of deception, such as the way a straight stick will look bent when placed underwater, or the mirages that result from wandering dehydrated through a barren desert. In these cases, if one were not suitably refining one’s responses to one’s sensory experiences, then one would be led into forming a false belief. If one did not know about light refraction, for example, then one would think that the stick really is bending as it enters the water; if one did not know that one was experiencing a mirage, then one would really believe that there was an unexpected oasis on the horizon.

There are also less mundane cases of perceptual error where the illusion is more widespread. One could imagine, for example, an environment in which one’s sensory experiences are a completely unreliable guide as to the nature of the environment. This could be achieved by hiding the real colours of the objects in the environment by employing fluorescent lights, or by using visual tricks to distort one’s sense of perspective in order to give the impression that objects are closer or farther away than they really are. The existence of perceptual error of this sort reminds us that,

whilst we must depend upon our perceptual faculties for much of our knowledge of the world, the possibility always remains that these faculties can lead us into forming false beliefs if left unchecked.

Given that we can usually correct for misleading perceptual impressions when they occur—as when we make use of our knowledge of light refraction to account for why straight sticks appear bent when placed in water—the mere possibility of perceptual error is not that worrying. The problem posed by perception is not, then, that it is a fallible way of gaining knowledge of the world; instead, it is its apparent indirectness.

Consider the visual impression caused by a genuine sighting of an oasis on the horizon, and contrast it with the corresponding visual impression of an illusory sighting of an oasis on the horizon formed by one who is hallucinating. Here is the crux: *these two visual impressions could be exactly the same*. The problem, however, is that it seems that if this is the case then what we experience in perception is not the world itself, but something that falls short of the world, something that is common to both the ‘good’ case in which one’s senses are not being deceived (and one is actually looking at an oasis) and the ‘bad’ case in which one’s senses are being deceived (and one is the victim of an hallucination). This line of reasoning which makes use of undetectable error in perception in order to highlight the indirectness of perceptual experience is known as the **argument from illusion**.

The argument from illusion suggests an ‘indirect’ model of perceptual knowledge, such that what we are immediately aware of when we gain such knowledge is a sensory impression—a seeming—on the basis of which we then make an inference regarding how the world is. That is, in both the deceived and non-deceived ‘oasis’ case just considered, what is common is a sensory impression of an oasis on the horizon, which leads one to infer something about the world: that there really is an oasis on the horizon. The difference between the two cases is that whilst the inference generates a true belief in the non-deceived case, it generates a false belief in the deceived case. In the former case, one is thus in a position, all other things being equal at least, to have perceptual knowledge that there is an oasis before one; whilst in the latter case perceptual knowledge is out of the question because one’s visual impressions are deceiving oneself.

But why is the indirectness, in this sense, of perceptual knowledge a problem? Well, the worry is that on this model of our perceptual interactions with the world, it seems that we are never actually perceiving a world external to our senses at all, strictly speaking, since our experiences are forever falling short of the world and requiring supplementation from reason. But isn’t this

conclusion more than just a little odd? Think of your perceptual experiences just now as you read this book. Aren't you *directly* experiencing the book in your hands?

Moreover, notice that this picture of the way we perceive the world, and thus gain perceptual knowledge, seems to have the result that our perceptual knowledge is far less secure than we might have otherwise thought. We normally regard our perceptual knowledge as the most secure of all. We often say, for example, that seeing is believing, and if we do indeed see something in clear daylight with our own two eyes, then this will tend to trump any counter-evidence we might have. For example, suppose that those around you assure you that your brother is out of town, and yet you see him walking towards you in the high street. Surely the testimony of your peers would be quickly disregarded and you would immediately believe that he is in town. According to our ordinary conception of perceptual knowledge, then, it is epistemically privileged relative to (at least some) other types of knowledge. But if perceptual experience does not put us in direct contact with the world, as the argument from illusion suggests—such that perceptual knowledge rests in part on an inference—then it appears that our perceptual knowledge is no more privileged than other 'indirect' knowledge that we have of the world (e.g., through testimony). In short, our knowledge of the world when we see that things are so is no better than it is when, say, we are merely told that things are so. But why, then, are we so confident in our perception-based judgements about the world?

Indirect Realism

The way of understanding perceptual knowledge which embraces the apparent indirectness of perceptual experience that we just noted is known as **indirect realism**. It holds that we gain knowledge of an objective world indirectly by making inferences from our sense impressions. The main argument for indirect realism is, in essence, the argument from illusion just given. The general idea is that the phenomenon of perceptual illusion highlights that what is presented to us in perceptual experience is not the world itself but merely an impression of the world from which we must draw inferences about how the world really is.

There is also another type of consideration in favour of indirect realism which concerns the between **primary/secondary quality** distinction that was drawn (in modern times) by the philosopher **John Locke (1632–1704)**, himself a proponent of a version of indirect realism. A primary quality is a feature of an object that the object has independently of anyone perceiving the

object. A secondary quality of an object, in contrast, is a feature of the object which is dependent upon the perception of an agent.

A good example of a primary quality is shape, in that the shape of an object does not seem to be in any way dependent upon anyone perceiving that object (or, e.g., on how they perceive it). Compare shape in this respect with colour. The colour of an object is a secondary quality in that it depends upon a perceiver. If human beings were kitted-out with different perceptual faculties, then colours would be discriminated very differently. Indeed, think of the animal kingdom in this respect, where there are creatures that can see colours that we can't see, and also creatures that are unable to see colours that we can see.

Note that this is not to suggest that colour is in some way an unreal or illusory feature of an object, since it is certainly a stable fact about, say, the UK's Royal Mail postboxes that they will generate a visual impression of redness to any person with the standard visual faculties who is looking at the postbox in normal lighting conditions. It is thus a real feature of the world that there are objects that generate visual impressions in this way. The point is rather that the colour of an object is not intrinsic to the object in the way that its shape is, but instead depends upon there being perceivers who respond to the object with the appropriate visual impressions.

The indirect realist is clearly in a good position to accommodate the primary/secondary quality distinction. After all, there is, on this view, a distinction between the world as it is perceived and the world as it really is, independently of being perceived. This distinction maps neatly onto the primary/secondary quality distinction, with the secondary qualities of an object belonging to the former realm, and the primary qualities of an object belonging to the latter realm.

The chief problem with indirect realism is that by making our perceptual knowledge of the world inferential, it threatens to dislocate us from the world altogether. Intuitively, what I am aware of when I open my eyes is the world itself, not mere sensory impressions of the world—'sense-data', as **Bertrand Russell (1872-1970)** characterises it—from which I infer specific beliefs about the world. Indeed, once one has departed down the road of indirect realism, it is not difficult to see the attraction of a widespread **scepticism** about our knowledge of the world, where this is the view that it is impossible to know anything about the world. After all, if what I am immediately aware of when I perceive is only an impression of the world from which I must then make an inference about the way the world is that could be either right or wrong, then why should I think that I have *any* knowledge at all of how the world really is?

This point is exacerbated once one considers the possibility that the way the world appears and the way that it really is could be drastically different in this view. Suppose, for example, that I am being radically deceived in my sensory impressions by some mischievous super-being who is ‘feeding’ me sensory impressions that are entirely misleading (this is what is known as a **sceptical hypothesis**). If all that I am directly aware of in perceptual experience is the way the world appears, then it seems that I could never be in a position to detect that this deception was going on. If it were taking place, however, then the way the world appeared would be no guide at all to how the world is, and thus the inferences I would be making about the nature of the world on the basis of my visual impressions would be dubious at best. Given this problem, it seems that all that I am entitled to take myself to know on this view is how the world appears, and not how it really is.

This difficulty is known as the **problem of the external world**, and whilst this problem is one that must be dealt with, in some form, by all theories of perceptual knowledge, it does seem as if indirect realism aggravates this difficulty by offering an account of perceptual knowledge which makes our knowledge of the external world shakily inferential rather than direct. Indeed, some have responded to indirect realism by arguing that, if this is how we are to understand perceptual knowledge, then we lose any grounds for thinking that there is a world that is independent of our experience of it (i.e., a world which is ‘external’ in the relevant sense).

Idealism

The view that denies that there is an external world that is independent of our experience of it is known as **idealism**. Perhaps the most famous exponent of a version of this position is **George Berkeley (1685–1753)**. Idealists respond to the problem of the external world by claiming that perceptual knowledge is not knowledge of a world that is independent of our perception of it, but rather knowledge of a world that is constituted by our perception of it. On this view the world is, so to speak, ‘constructed’ out of appearances rather than being that which gives rise to such appearances, and thus it is not ‘external’ in the relevant sense at all. (Another way of putting the point is that for the idealist there are only secondary qualities.) As Berkeley famously put it in his book, *A Treatise Concerning the Principles of Human Knowledge*, ‘To be is to be perceived’. This is a very dramatic conclusion to draw, and appears to call much of our ordinary conception of the world and our relation to it into question.

If the view is not qualified in some way, then it will end up maintaining that the world ceases to exist when no one perceives it. For example, one can't say that a tree fell in the forest if there was no one around to see or hear (or otherwise sense) it fall; if no one experienced the falling of the tree, then in the idealist view the event didn't happen. This is clearly a very radical claim to make! Indeed, it is hard to distinguish a simple-minded idealism of this sort from plain scepticism about our perceptual knowledge. Although, unlike the sceptic, the idealist claims that we do know a great deal about the world, she achieves this feat by making what we mean by the 'world' so different from what we usually take it to mean that it feels as if the idealist is agreeing with the sceptic after all.

Berkeley's way of lessening some of the more outlandish consequences of a simple idealism was to appeal to the idea of an ever-present God. With God in the picture, we now no longer need to worry about what to make of unobserved events, since an all-seeing God will observe all events. Accordingly, we aren't forced to say that events that aren't observed by us mere mortals therefore don't happen. Berkeley was a Christian—a bishop, in fact—so this appeal to God is unsurprising. This sort of refinement to idealism would clearly offer little comfort to an idealist who was also an atheist though!

Solipsism

One worry about idealism is that it might collapse into a radical form of idealism known as **solipsism**. Solipsism is the view that the only thing that exists is one's own mind. On this view, then, both the external world and other minds are to be understood as lacking any independent existence, as they exist only as elements of one's own mind. It should be easy to see how idealism might lead one in this direction. As we noted above, if one removes God from Berkeley's idealism, then it would follow that the objects that we perceive do not exist when no-one is perceiving them. But of course, from one's own subjective point of view, how can one even be sure that there other people that exist when one is not perceiving them? Why isn't their existence dependent upon one's perception of them just as the objects we perceive are? Indeed, once we have begun to countenance this line of thought, then one might wonder whether the only thing that exists are one's *current* mental states—i.e., what one is experiencing right now.

Naïve Realism

All this talk of idealism and solipsism can make one wonder whether something didn't simply go wrong in our reasoning right at the start of our thinking about this topic. How could it be that reflecting on the nature of our perceptual experience of the world has led us to think that perhaps there is no external world to have knowledge of in the first place (or at least no external world that we can know through experience), and perhaps even no other minds but our own? With this in mind, it is worth considering the prospects for the simple-minded *direct realism*—or **naïve realism**, as it is often known. This is the view that we set aside earlier in order to opt for the indirect realism, on the grounds that only the latter appeared to be able to resolve the difficulties posed by the argument from illusion whilst also accounting for our intuitive distinction between primary and secondary qualities. But perhaps we were wrong to abandon this proposal so quickly?

In its simplest form, naïve realism takes our perceptual experiences at face value and argues that, at least in non-deceived cases, what we are aware of in perceptual experience is the external world itself. That is, if I am genuinely looking at an oasis on the horizon right now, then I am directly aware of the oasis itself, and thus I can have perceptual knowledge that there is an oasis before me without needing to make an inference from the way the world seems to how it is.

The motivation behind naïve realism, besides the obvious attraction that of all the views it most accords with common sense, is that other theories of perceptual knowledge, such as indirect realism and idealism, are far too quick to infer from the fact that our perceptual experience could be undetectably misleading that we are only directly aware of the way the world seems to us rather than the way the world is. The idea is that although it is true that in deceived cases, such as the scenario in which I am visually presented with a mirage of an oasis, that I am not directly aware of the world but only with the way the world appears, this should not be thought to entail that in non-deceived cases, such as where I am actually looking at an oasis in the distance, I am not directly acquainted with objects in the world. On this view, the fact that I am not always able to distinguish between deceived and non-deceived cases is neither here nor there. This is because it is not held to be a precondition of perceptual knowledge that one can tell the genuine cases of perceptual knowledge apart from the merely apparent cases.

Of course, the naïve realist cannot leave matters there, since she needs to go on to explain how such a view is to function. For one thing, she needs to develop a theory of knowledge that can allow us to have perceptual knowledge directly via perceptual experience even in cases where one is

unable to distinguish genuine from apparent perception. Moreover, she also needs to offer an explanation of the primary/secondary quality distinction. Nevertheless, given the unattractiveness of indirect realism and the versions of idealism that are suggested by the move to indirect realism, naïve realism needs to be taken very seriously indeed.

Summary

- ❑ A great deal of our knowledge of the world is gained via perception (i.e., via our senses). Our senses are sometimes prone to deceive us, though, as we noted, this is not a problem in itself, since we can often tell when they are not to be trusted (as when we see a stick ‘bend’ as it enters water). What is problematic about perceptual experience is brought out via the argument from illusion. In essence, this states that since a situation in which we are deceived about the world could be one in which we have, it seems, exactly the same experiences as we would have in a corresponding undeceived case, so we don’t directly experience the world at all.
- ❑ The conception of perceptual knowledge suggested by the argument from illusion is that of indirect realism. This holds that there is an objective world out there, one that is independent of our experience of it—this is the ‘realism’ part—but that we can only know this world indirectly through experience. In particular, what we directly experience is only how the world appears to us, and not how it is. On this basis, we can then make inferences to how the world really is.
- ❑ Indirect realism can also easily account for the primary/secondary quality distinction—the distinction between those (primary) properties or qualities of an object that are inherent in the object, such as its shape, and those (secondary) properties or qualities of an object that are dependent upon the perceiver, such as its colour.
- ❑ On the indirect realist view, we don’t have any direct experience of the external world, and this has prompted some to argue for a view known as idealism, which maintains that there is no external world that exists independently of experience. In particular, idealism maintains that the world is constructed out of appearances and does not extend beyond it—that is, there is no mind-independent world. We saw that Berkeley attempted to limit the radical implications of this thesis by arguing that God is always present to perceive the world, and hence the world doesn’t simply cease to exist when we are not perceiving it.

- ❑ One particularly radical form of idealism is known as solipsism. On this view, not only is there no mind-independent world, but there are also no other people. That is, the only thing that exists is one's own mind (or even, on a radical rendering of solipsism, just one's current experiences).
- ❑ Finally, we considered a common-sense view of perceptual experience called naïve (or direct) realism. This view holds that we *can* directly experience the world, and so rejects the conclusion usually derived from the argument from illusion that direct experience of the world is impossible.

Study Questions

- ❑ Think of two examples of when your experiences have been a misleading guide as to the way the world is.
- ❑ What is the argument from illusion? What is indirect realism? Explain, in your own words, why the argument from illusion offers support for indirect realism.
- ❑ Explain, in your own words, what the primary/secondary quality distinction is. Pick an object, and give an example of a primary quality that this object has and a secondary quality that this object has.
- ❑ Explain, in your own words, what idealism is. Do you find this position plausible? If not, say why. If so, then try to think why others might find it implausible, and try to see if you can offer any considerations in defence of the view in light of these concerns.
- ❑ How does Berkeley employ God to try to limit some of the counterintuitive implications for idealism? Does this help to make idealism more palatable, do you think?
- ❑ What is solipsism, and why might one think that solipsism is the natural consequence of endorsing idealism? Is there any way of preventing idealism from collapsing into solipsism?
- ❑ What is naïve realism? Do you find this position plausible? If not, say why. If so, then try to think why others might find it implausible, and try to see if you can offer any considerations in defence of the view in light of these concerns.

Introductory Reading

- ❑ Dancy, J. (1987) *Berkeley: An Introduction*, (Oxford: Blackwell). [The best introduction to Berkeley's philosophy in recent years].

- ❑ Dunn, J. (2003). *Locke: A Very Short Introduction*, (Oxford: Oxford University Press). [A nice introduction to Locke's philosophy].
- ❑ Martin, R. M. (2010). *Epistemology*, (London: Oneworld). [See chapter 7 for a very accessible discussion of perceptual knowledge].
- ❑ Pritchard, D. H. (2013). *What is this Thing Called Knowledge?*, (London: Routledge, 3rd Edn.). [See chapter 7 for an overview of issues regarding perceptual knowledge].
- ❑ Sosa, D. (2010). 'Perceptual Knowledge', *The Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, chapter 27, (New York: Routledge). [A sophisticated, yet reasonably accessible, overview of the epistemological issues raised by perception].

Further Reading

- ❑ Richmond, A. (2009). *Berkeley's Principles of Human Knowledge: A Reader's Guide*, (London: Bloomsbury). [An excellent exposition of the original text].
- ❑ Robinson, H. (1994). *Perception*, (London: Routledge). [A good recent discussion of the central issues in this area. Not for the beginner].
- ❑ Schwartz, R. (ed.) (2003). *Perception*, (Oxford: Blackwell). [A nice collection of articles on the philosophy of perception, including both historical texts and contemporary readings].

Free Internet Resources

- ❑ Bonjour, L. (2007). 'Epistemological Problems of Perception', *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/perception-episprob/>>. [An excellent overview of the central issues by one of the leading figures in the field].
- ❑ Downing, L. (2011) 'Berkeley', *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/berkeley/>>. [A good and completely up-to-date introduction to the work of Berkeley].
- ❑ Flage, D. (2005). 'George Berkeley', *Internet Encyclopaedia of Philosophy*, <<http://www.iep.utm.edu/berkeley/>>. [A helpful overview of Berkeley's life and works].
- ❑ O'Brien, D. (2004). 'The Epistemology of Perception', *Internet Encyclopaedia of Philosophy*, <<http://www.iep.utm.edu/e/epis-per.htm>>. [A first-rate introduction to the main issues regarding the epistemology of perception].

- Thornton, S. (2012). 'Solipsism and the Problem of Other Minds', *Internet Encyclopaedia of Philosophy*, <<http://www.iep.utm.edu/solipsis/>>. [A useful overview of issues surrounding solipsism, though do be aware that this takes in further considerations in this regard over and above what we have covered here].
- Uzgalis, W. (2012). 'Locke', *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/locke/>>. [An excellent and completely up-to-date overview of the life and works of Locke].

TOPIC 2:
THE DEFINITION OF KNOWLEDGE

- ❑ *Knowledge, Truth and Belief*
- ❑ *Knowledge versus Mere True Belief*
- ❑ *The Classical Account of Knowledge*
- ❑ *The Gettier Problem*
- ❑ *Responding to the Gettier Problem*

Knowledge, Truth and Belief

Think of all the things that you know, or at least think you know, right now. You know, for example, that the earth is round and that Paris is the capital of France. You know that you can speak (or at least read) English, and that two plus two is equal to four. You know, presumably, that all bachelors are unmarried men, that it is wrong to hurt people just for fun, that *The Godfather II* is a wonderful film, and that the moon is not made of cheese. And so on.

But what is it that all these cases of knowledge have in common? Think again of the examples just given, which include geographical, linguistic, mathematical, aesthetic, ethical, and scientific knowledge. Given these myriad types of knowledge, what, if anything, ties them all together?

In all the examples of knowledge just given, the type of knowledge in question is what is called **propositional knowledge**, in that it is knowledge of a **proposition**. A proposition is what is asserted by a sentence which says that something is the case—e.g., that the earth is flat, that bachelors are unmarried men, that two plus two is four, and so on. Propositional knowledge will be our focus here, but we should also recognise from the outset that it is not the only sort of knowledge that we possess.

There is, for example, **ability knowledge**, or *know-how*. Ability knowledge is clearly different from propositional knowledge; I know how to swim, for example, but I do not thereby know a set of propositions about how to swim. Indeed, I'm not altogether sure that I could tell you how to swim, but I do know how to swim nonetheless (and I could prove it by manifesting this ability—by jumping into a swimming pool and doing the breaststroke, say).

Ability knowledge is certainly an important type of knowledge to have. We want lots of know-how, such as to know how to ride a bicycle, to drive a car, or to operate a personal computer. Notice, however, that while only relatively sophisticated creatures like humans possess propositional knowledge, ability knowledge is far more common. An ant might plausibly be said to know how to navigate its terrain, but would we want to say that an ant has propositional knowledge; that there are facts which the ant knows? Could the ant know, for example, that the terrain it is presently crossing is someone's porch? Intuitively not, and this marks out the importance of propositional knowledge over other types of knowledge like ability knowledge, which is that such knowledge presupposes the sort of relatively sophisticated intellectual abilities possessed by (mature) humans.

Henceforth, when we talk about knowledge, we will have propositional knowledge in mind. Two things that just about every epistemologist agrees on are that a prerequisite for possessing knowledge is that one has a belief in the relevant proposition, and that that belief must be true. So if you know that Paris is the capital of France, then you must believe that this is the case, and your belief must also be true.

Take the truth requirement first. In order to assess this claim, consider what would follow if we dropped this requirement. In particular, is it plausible to suppose that one could know a false proposition? Of course, we often *think* that we know something and then it turns out that we were wrong, but that's just to say that we didn't really know it in the first place. But could one genuinely know a false proposition? Could I know, for example, that the moon is made of cheese, even though it manifestly isn't? I take it that when we talk of someone having knowledge, we mean to exclude such a possibility. This is because to ascribe knowledge to someone is to credit that person with having gotten things right, and that means that what we regard that person as knowing had better not be false, but true.

Next, consider the belief requirement. It is sometimes the case that we explicitly *contrast* belief and knowledge, as when we say things like, 'I don't merely believe that he was innocent, I know it', which might on the face of it be thought to imply that knowledge does not require belief after all. If you think about these sorts of assertions in a little more detail, however, then it becomes clear that the contrast between belief and knowledge is being used here simply to emphasise the fact that one *not only* believes the proposition in question, but *also* knows it. In this way, these assertions actually lend support to the claim that knowledge requires belief, rather than undermining it.

As with the truth requirement, we will assess the plausibility of the belief requirement for knowledge by imagining for a moment that it doesn't hold, which would mean that one could have knowledge of a proposition that one did not even believe. Suppose, for example, that someone claimed to have known a quiz answer, even though it was clear from that person's behaviour at the time that she didn't believe the proposition in question (perhaps she put forward a different answer to the question, or no answer at all). Clearly we would not agree that this person did have knowledge in this case. Again, the reason for this relates to the fact that to say that someone has knowledge is to credit that person with a certain kind of success. But for it to be your success, then belief in the proposition in question is essential, since otherwise this success is not creditable to you at all.

Knowledge versus Mere True Belief

It is often noted that belief *aims* at the truth, in the sense that when we believe a proposition, we believe it to be the case (i.e., to be true). When what we believe is true, then there is a match between what we think is the case and what is the case. We have got things right. If mere true belief suffices for 'getting things right', however, then one might wonder as to why epistemologists do not end their quest for an account of knowledge right there and simply hold that knowledge is nothing more than true belief (i.e., 'getting things right').

There is in fact a very good reason why epistemologists do not rest content with mere true belief as an account of knowledge, and that is that one can gain true belief entirely by *accident*, in which case it would be of no credit to you at all that you got things right. Consider Harry, who forms his belief that the horse Lucky Lass will win the next race purely on the basis of the fact that the name of the horse appeals to him. Clearly this is not a good basis on which to form one's belief about the winner of the next horse race, since whether or not a horse's name appeals to you has no bearing on its performance.

Suppose, however, that Harry's belief turns out to be true, in that Lucky Lass *does* win the next race. Is this knowledge? Intuitively not, since it is just a matter of *luck* that his belief was true in this case. Remember that knowledge involves a kind of success that is creditable to the agent. Crucially, however, successes that are merely down to luck are never creditable to the agent.

In order to emphasise this point, think for a moment about successes in another realm, such as archery. Notice that if one genuinely is a skilled archer, then if one tries to hit the bull's-eye, and the conditions are right (e.g., the wind is not gusting), then one usually *will* hit the bull's-eye. That's

just what it means to be a skilled archer. The word ‘usually’ is important here, since someone who isn’t a skilled archer might, as it happens, hit the bull’s eye on a particular occasion, but she wouldn’t *usually* hit the bull’s-eye in these conditions. Perhaps, for example, she aims her arrow and, by luck, it hits the centre of the target. Does the mere fact that she is successful on this one occasion mean that she is a skilled archer? No, and the reason is that she would not be able to repeat this success. If she tried again, for example, her arrow would in all likelihood sail off into the heavens.

Having knowledge is just like this. Imagine that one’s belief is an arrow, which is aimed at the centre of the target, truth. Hitting the bull’s-eye and forming a true belief suffices for getting things right, since all this means is that one was successful on that occasion. It does not suffice, however, for having knowledge any more than hitting the bull’s-eye purely by chance indicates that you are skilled in archery. To have knowledge, one’s success must genuinely be the result of one’s efforts, rather than merely being by chance. Only then is that success creditable to one. And this means that forming one’s belief in the way that one does ought usually, in those circumstances, to lead to a true belief.

Harry, who forms his true belief that Lucky Lass will win the race simply because he likes the name, is like the person who happens to hit the bull’s-eye, but who is not a skilled archer. Usually, forming one’s belief about whether a horse will win a race simply by considering whether the name of the horse appeals to you will lead you to form a false belief.

Contrast Harry with someone who genuinely knows that Lucky Lass will win the race. Perhaps, for example, this person is a ‘Mr Big’, a gangster who has fixed the race by drugging the other animals so that his horse, Lucky Lass, will win. He knows that the race will be won by Lucky Lass because the way he has formed his belief, by basing it on the special grounds he has for thinking that Lucky Lass cannot lose, would normally lead him to have a true belief. It is not a matter of luck that Mr Big hits the target of truth.

The challenge for epistemologists is thus to explain what needs to be added to mere true belief in order to get knowledge. In particular, epistemologists need to explain what needs to be added to true belief to capture this idea that knowledge, unlike mere true belief, involves a success that is creditable to the agent, where this means, for example, that the agent’s true belief was not simply a matter of luck.

The Classical Account of Knowledge

So it seems that there must be more to knowledge than just true belief. But what could this additional component be? The natural answer to this question, one that is often ascribed to the ancient Greek philosopher **Plato (c. 427-c. 347 BC)**, is that what is needed is a *justification* for one's belief, where this is understood as being in possession of good reasons for thinking that what one believes is true. This proposal is known as the **classical account of knowledge**. (It also sometimes referred to as the 'tripartite'—i.e., three-part—account of knowledge).

Consider again the case of Harry, who believes that Lucky Lass will win the race because he likes the name, and Mr. Big, who forms the same belief on the grounds that he has fixed the race. As we noted, although both of these agents believe truly, only Mr. Big intuitively has knowledge of what he believes. The claim that it is justification that marks the difference between knowledge and mere true belief accords with this assessment of our two agents' beliefs. Mr. Big, after all, has excellent reasons in support of his true belief, since he is aware that the other horses have been drugged and so don't have a hope of winning (unlike the undrugged Lucky Lass). Harry, in contrast, can't offer any good reasons in support of his belief. That he happens to like the name of a horse is hardly a good reason for thinking that this horse will win a race!

Plausibly, then, the missing ingredient in our account of knowledge is justification, such that knowledge is justified true belief. Indeed, until relatively recently most epistemologists thought that this theory of knowledge was correct. Unfortunately, as we will now see, the classical account of knowledge cannot be right, even despite its surface plausibility.

The Gettier Problem

The person who demonstrated that the classical account of knowledge is untenable was a philosopher named Edmund Gettier (b. 1927). In a very short article—just two-and-a-half pages in length—he offered a devastating set of counterexamples to the classical account: what are now known as **Gettier cases**. In essence, what Gettier showed was that you could have a justified true belief and yet still lack knowledge of what you believe because your true belief was ultimately gained via luck in much the same way as Harry's belief was gained by luck.

We will use a different example from the ones cited by Gettier, though one that has the same general structure. Imagine a man, let's call him John, who comes downstairs one morning and sees that the time on the grandfather clock in the hall says '8.20'. On this basis John comes to believe that

it is 8.20 a.m., and this belief is true, since it *is* 8.20 a.m. Moreover, John's belief is justified in that it is based on excellent grounds. For example, John usually comes downstairs in the morning about this time, so he knows that the time is about right. Moreover, this clock has been very reliable at telling the time for many years and John has no reason to think that it is faulty now. He thus has good reasons for thinking that the time on the clock is correct.

Suppose, however, that the clock had, unbeknownst to him, stopped 24 hours earlier, so that John is now forming his justified true belief by looking at a stopped clock. Intuitively, if this were so then John would lack knowledge even though he has met the conditions laid down by the classical account of knowledge. After all, that John has a true belief in this case is, ultimately, a matter of luck, just like Harry's belief that Lucky Lass would win the 4.20 at Kempton. If John had come downstairs a moment earlier or a moment later—or if the clock had stopped at a slightly different time—then he would have formed a false belief about the time by looking at this clock. Thus we can conclude that knowledge is not simply justified true belief.

There is a general form to all Gettier cases, and once we know this we can use it to construct an unlimited number of them. To begin with, we need to note that you can have a justified false belief, since this is crucial to the Gettier cases. For example, suppose you formed a false belief by looking at a clock that you have every reason for thinking is working properly (e.g., the clock is normally reliable, it is showing what appears to be the right time, and so on), but which is, in fact, and unbeknownst to you, broken. This belief would clearly be justified, even though it is false, since you have excellent reasons to offer in support of what you believe. With this point in mind, there are three stages to constructing your own Gettier case.

First, you take an agent who forms her belief in a way that would usually lead her to have a false belief. In the example just cited, we took the case of someone looking at a stopped clock in order to find out the time. Clearly, using a stopped clock to find out the time would usually result in a false belief. Second, you add some detail to the example to ensure that the agent's belief is justified nonetheless. In the example just cited, the detail we added was that the agent had every reason for thinking that the clock was working properly, thus ensuring that her belief is entirely justified. Finally, you make the case such that while the way in which the agent formed her belief would normally have resulted in a justified false belief, in this case it so happened that the belief was true. In the stopped clock case, stipulating that the stopped clock just happens to be 'telling' the right time does this.

Putting all this together, we can construct an entirely new Gettier case from scratch. As an example of someone forming a belief in a way that would normally result in a false belief, let's take someone who forms her belief that Beyoncé is across the street by looking at a life-sized cardboard cut-out of Beyoncé which is advertising her forthcoming tour, and which is posted just across the street. Forming one's belief about whether someone is across the street by looking at a life-sized cut-out of that person would not normally result in a true belief. Next, we add some detail to the example to ensure that the belief is justified. In this case we can just stipulate that the cut-out is very authentic-looking, and that there is nothing about it which would obviously give away the fact that it is a cardboard cut-out—e.g., it does not depict Beyoncé in an outrageous costume that she wouldn't plausibly wear on a normal street. The agent's belief is thus justified. Finally, we make the scenario such that the belief is true. In this case, for instance, all we need to do is stipulate that, as it happens, Beyoncé *is* across the street, doing some window-shopping out of view of our agent. *Voilà*, we have constructed our very own Gettier case!

Responding to the Gettier Problem

There is no easy way to respond to the Gettier cases, and since Gettier's article back in 1963 a plethora of different theories of knowledge have been developed in order to offer an account of knowledge that is Gettier-proof. Initially, it was thought that all one needed to do to deal with these cases is simply tweak the classical account of knowledge. For instance, one proposal was that in order to have knowledge, one's true belief must be justified and also not in any way based on false presuppositions, such as, in the case of the stopped clock described above, the false presupposition that the clock is working. This proposal is known as the **no false lemmas** account, where a lemma is roughly a premise that one is relying upon when reasoning to a conclusion.

There is a pretty devastating problem with this sort of proposal, however, which is that it is difficult to spell out this idea of a 'presupposition' such that it is strong enough to deal with Gettier cases and yet not so strong that it prevents us from having most of the knowledge that we think we have.

For example, suppose that our agent in the stopped clock example (John) has a sister across town—let's call her Sally—who is in fact at this moment finding out what the time is by looking at a working clock. Intuitively, Sally *does* gain knowledge of what the time is by looking at the time on the clock. Notice, however, that Sally may believe all sorts of other related propositions, some of which

may be false—for example, she may believe that the clock is regularly maintained, when in fact no one is taking care of it. Is this belief a presupposition of her belief in what the time is? If it is (i.e., if we understand the notion of a ‘presupposition’ liberally) then this false presupposition will prevent her from having knowledge of the time, even though we would normally think that looking at a reliable working clock is a great way of coming to know what the time is.

Alternatively, suppose we understand the notion of a ‘presupposition’ in a more restrictive way such that this belief isn’t a presupposition of Sally’s belief in the time. The problem now is to explain why John’s false belief that he’s looking at a working clock counts as a presupposition of his belief in the time (and so prevents him from counting as knowing what the time is) if Sally’s false belief that the clock is regularly maintained is not also treated as a presupposition. Why don’t they *both* lack knowledge of what the time is?

If this problem weren’t bad enough, there is also a second objection to this line of response to the Gettier cases, which is that it is not clear that the agent in a Gettier case need presuppose *anything* at all. Consider a different Gettier case in this regard, due to Roderick Chisholm (1916-1999). In this example, we have a farmer—let’s call her Gayle—who forms her belief that there is a sheep in the field by looking at a shaggy dog, which happens to look just like a sheep. As it turns out, however, there is a sheep in the field (standing behind the dog), and hence Gayle’s belief is true. Moreover, her belief is also justified because she has great evidence for thinking that there is a sheep in the field (she can see what looks to be a sheep in the field, for example).

Given the immediacy of Gayle’s belief in this case, however, it is hard to see that it really presupposes any further beliefs at all, at least unless we are to understand the notion of a presupposition *very* liberally. And notice that if we do understand the notion of a presupposition so liberally that Gayle counts as illicitly making a presupposition, the problem then re-emerges of how to account for apparently genuine cases of knowledge, such as that intuitively possessed by Sally.

The general dilemma for proponents of this sort of response to the Gettier cases is thus to explain how we should understand the notion of a presupposition broadly enough so that it applies to the Gettier cases while at the same time understanding it narrowly enough so that it doesn’t apply to other non-Gettier cases in which, intuitively, we would regard the agent concerned as having knowledge. In short, we want a response to the problem that explains why John lacks knowledge in such a way that it doesn’t thereby deprive Sally of knowledge.

Once it was recognised that there was no easy answer to the problem posed to the classical account of knowledge by the Gettier cases, the race was on to find a radically new way of analysing knowledge which was Gettier-proof. One feature that all such accounts share is that they understand the conditions for knowledge such that they demand more in the way of co-operation from the world than simply that the belief in question is true. That is, on the classical account of knowledge there is one condition which relates to the world—the truth condition—and two conditions that relate to us as agents—the belief and justification conditions. These last two conditions, at least as they are usually understood in any case, don't demand anything from the world in the sense that they could obtain regardless of how the world is. If I were the victim of an hallucination, for example, then I might have a whole range of wholly deceptive experiences, experiences that, nonetheless, lead me to believe something and, moreover, to justifiably believe it. (For example, if I seem to see that, say, there is a glass in front of me, then this is surely a good, and thus justifying, reason for believing that there is a glass in front of me, even if the appearance of the glass is an illusion.) The moral of the Gettier cases is, however, that you need to demand more from the world than simply that one's justified belief is true if you are to have knowledge.

In the stopped-clock Gettier case, for example, the problem came about because, although John had excellent grounds for believing what he did, it nevertheless remained that he did not know what he believed because of some oddity in the world—in this case that the normally reliable clock had not only stopped but had stopped in such a way that John still formed a true belief. It thus appears that we need an account of knowledge, which imposes a further requirement on the world over and above the truth of the target belief—that, for example, the agent is *in fact* forming his belief in the right kind of way. But specifying exactly what this requirement involves is far from easy.

For example, one way that could deal with the Gettier problem would be to demand that one's justified belief is **infallible**. That is, one might demand that when one knows one has formed one's belief in such a way that one could not have been wrong. This kind of position is known as **infallibilism**. Clearly this would resolve the Gettier problem, as all Gettier-style cases trade on the idea that one has a justified true belief that could have easily been false. But this solution to the problem simply exchanges one epistemological difficulty for another. For while we now have an account of knowledge that is Gettier-proof, it is also now far from obvious that any of our everyday beliefs would count as knowledge by the lights of this proposal.

Another way to go is to reject the justification condition and instead opt for something else. One proposal that has been popular in this regard is to demand that one's true belief should have been formed as a result of one's exercise of **epistemic virtue**. This proposal is known as **virtue epistemology**. What is meant by epistemic virtue here is usually a cluster of reliable cognitive skills that we have, ranging from innate cognitive faculties, such as our perceptual faculties, to quite sophisticated intellectual traits known as intellectual virtues, such as conscientiousness. The thinking behind virtue epistemology is that what is important when it comes to knowledge is not so much the ability to offer good reasons in support of one's beliefs, but rather how one's formation of a true belief manifests one's epistemic character. So, for example, someone who skilfully forms a true belief via a manifestation of their cognitive skills ought to count as having knowledge even if they are unable to offer a justification for their belief. Moreover, even when one can offer a justification for one's true belief, the reason why this is important to knowledge is that in offering justifications one is also thereby manifesting one's good cognitive character.

Although there are lots of attractive features of virtue epistemology, it doesn't obviously give us a handle on the Gettier problem. After all, in standard Gettier-style cases the agent concerned is forming a true belief that manifests their exercise of epistemic virtue. Our Gettierised farmer that we described above, for example, is manifesting her reliable perceptual faculties in picking out a sheep-shaped object and forming the true belief that there is a sheep in the field. It seems, then, that she ought to count as having knowledge according to virtue epistemology and yet, as we have seen, it doesn't seem right to credit agents in Gettier-style cases with knowledge.

Summary

- ❑ Epistemology is (essentially) the theory of knowledge. One of the characteristic questions of epistemology concerns what all the myriad kinds of knowledge we ascribe to ourselves have in common: *What is knowledge?*
- ❑ We can distinguish between knowledge of propositions, or propositional knowledge, and know-how, or ability knowledge. Intuitively, the former demands a greater degree of intellectual sophistication on the part of the knower than the latter.
- ❑ In order to have knowledge of a proposition, that proposition must be true, and one must believe it.

- ❑ Mere true belief does not suffice for knowledge, however, since one can gain mere true belief purely by luck, and yet you cannot gain knowledge purely by luck.
- ❑ According to the classical (or tripartite) account of knowledge, knowledge is understood as justified true belief, where a justification for one's belief consists of good reasons for thinking that the belief in question is true.
- ❑ Gettier cases are cases in which one forms a true justified belief and yet lacks knowledge because the truth of the belief is largely a matter of luck. (The example we gave of this was that of someone forming a true belief about what the time is by looking at a stopped clock, which just so happens to be displaying the right time.) Gettier cases show that the classical account of knowledge in terms of justified true belief is unsustainable.
- ❑ There is no easy answer to the Gettier cases; no simple way of supplementing the classical account of knowledge so that it can deal with these cases. We looked at the no false lemmas response to the Gettier problem, and saw that it faced the problem of finding a principled way of identifying which false lemmas are the relevant ones. We saw that infallibilism would deal with the problem, but would do so at the expense of making much of our knowledge unobtainable. And we saw that virtue epistemology seems to struggle with the Gettier problem just as much as the classical account.

Study Questions

- 1 Give examples of your own of the following types of knowledge:
 - ❑ scientific knowledge;
 - ❑ geographical knowledge;
 - ❑ historical knowledge;
 - ❑ religious knowledge.
- 2 Explain, in your own words, the distinction between ability knowledge and propositional knowledge. Give two examples of each kind of knowledge.
- 3 Why is mere true belief not sufficient for knowledge? Give an example of your own of a case in which an agent truly believes something, but does not know it.
- 4 What is the classical account of knowledge? How does the classical account of knowledge explain why a lucky true belief doesn't count as knowledge?

- 5 What is a Gettier case, and what do such cases show? Try to formulate a Gettier case of your own.
- 6 In what way might it be said that the problem with Gettier cases is that they involve a justified true belief that is based on a false presupposition? Explain, with an example, why one cannot straightforwardly deal with the Gettier cases by advancing a theory of knowledge that demands justified true belief that does not rest on any false presuppositions.
- 7 What is infallibilism, and why would embracing this view deal with Gettier cases? Given that it can deal with Gettier cases, is there any reason why we shouldn't adopt this position?
- 8 What is virtue epistemology, and why do virtue epistemologists claim that we should focus on epistemic virtues rather than justification? Can this proposal help us resolve the Gettier problem?

Introductory Further Reading

- ❑ Martin, R. M. (2010). *Epistemology*, (London: Oneworld). [See chapters 1-4 for a very accessible discussion of some of the issues surrounding the definition of knowledge].
- ❑ Pritchard, D. H. (2013). *What is this Thing Called Knowledge?*, (3rd Edn.), (London: Routledge). [See chapters 3-6 for further introductory-level discussion of the nature of knowledge].

Advanced Further Reading

- ❑ Hetherington, S. (2010). 'The Gettier Problem', *The Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, chapter 12, (London: Routledge). [A very useful and completely up-to-date survey of the main issues raised by Gettier-style examples].
- ❑ Pritchard, D. H. (2016). *Epistemology*, (London: Palgrave Macmillan). [This is an advanced textbook in epistemology. Chapters 1-4 offer a critical overview of some of the main analyses of knowledge in the contemporary literature].
- ❑ Shope, R. K. (2002). 'Conditions and Analyses of Knowing', *The Oxford Handbook to Epistemology*, (ed.) P. K. Moser, pp. 25–70, (Oxford: Oxford University Press). [A comprehensive treatment of the problem posed by Gettier cases and the various contemporary responses to that problem in the literature. The discussion that starts on page 29 is most relevant to this chapter. Note that as this chapter develops it becomes increasingly more demanding].

- Steup, M., Turri, J., & Sosa, E. (eds.) (2013) *Contemporary Debates in Epistemology* (2nd Edn.), (Oxford: Blackwell). [See the exchange between Duncan Pritchard and Stephen Hetherington on whether there can be lucky knowledge (§7)].
- Zagzebski, L. (1999). ‘What is Knowledge?’ *The Blackwell Companion to Epistemology*, (eds.) J. Greco & E. Sosa, pp. 92–116 (Oxford: Blackwell). [A very thorough overview of the issues surrounding the project of defining knowledge, especially in the light of the Gettier cases. Note that Zagzebski is a prominent virtue epistemologist, and so she also introduces this position].

Free Internet Resources

- Gettier, E. (1963). ‘Is Justified True Belief Knowledge?’, *Analysis* 23, 121–3 (freely available on-line here: <http://www.ditext.com/gettier/gettier.html>). [The article which started the contemporary debate about how best to define knowledge and which contains, by definition, the first official Gettier cases].
- Greco, J., & Turri, J. (2011). ‘Virtue Epistemology’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/epistemology-virtue/>>. [A superb overview of virtue epistemology, written by two experts in the field].
- Hetherington, S. (2005). ‘Gettier Problems’, *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden & J. Fieser, <http://www.iep.utm.edu/g/gettier.htm>. [An excellent overview of the Gettier problem, and the main responses to it, by one of the leading epistemologists].
- Ichikawa, J., & Steup, M. (2012). ‘The Analysis of Knowledge’, *Stanford Encyclopaedia of Philosophy*, <http://plato.stanford.edu/entries/knowledge-analysis/>. [An excellent and comprehensive overview of the issues regarding the project of defining knowledge].
- Kraut, R. (2013). ‘Plato’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/plato/>>. [An excellent overview of the main themes in Plato’s philosophy].
- Truncellito, D. (2007). ‘Epistemology’, *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden & J. Fieser, <http://www.iep.utm.edu/e/epistemo.htm>. [Read up to the end of §2.b for more on the basic requirements for knowledge].

TOPIC 3:
THE ORIGINS OF CONCEPTS
AND THE NATURE OF KNOWLEDGE

- ❑ *Two Distinctions: A Priori/A Posteriori & Analytic/Synthetic*
- ❑ *Knowledge Empiricism*
- ❑ *Knowledge Innatism*
- ❑ *Concept Empiricism*
- ❑ *Concept Innatism*

Two Distinctions: A Priori/A Posteriori & Analytic/Synthetic

For this topic, we will be exploring an historically important philosophical debate between empiricism and rationalism. In particular, we will be looking at **knowledge empiricism** (along with **knowledge innatism**, which it is in conflict with), and **concept empiricism** (along with its opposing view, **concept innatism**). Before we can do that, however, we first need to introduce two very important distinctions.

The first is a distinction between ***a priori*** and ***a posteriori*** knowledge. This concerns two ways of knowing a proposition. *A posteriori* knowledge—or *empirical knowledge*, as it is sometimes called—is knowledge that is gained via sensory experience. So, for example, one might know that one’s eyes are blue because one saw them in a mirror. In contrast, *a priori* knowledge is knowledge that was gained independently of sensory experience. So, for example, one might know that all triangles have three sides purely in virtue of understanding what a triangle is.

One thing that we need to be careful about here is that sometimes the *a priori/a posteriori* distinction is applied to how one in fact acquired one’s knowledge, while other times it is applied to a particular category of proposition. So, for example, when it is claimed that it is *a priori* that triangles have three sides, what is meant is that this is a proposition which can be known *a priori*. In contrast, it is not *a priori* that one’s eyes are blue, since this is only something that can be known by someone looking to check. That triangles have three sides is *a priori* does not mean, however, that one can’t come to know this proposition in an *a posteriori* fashion. Perhaps, for example, one knows this proposition not via one’s understanding of what a triangle is—i.e., by recognising that this must be so—but only because someone authoritative, such as a schoolteacher, told you so. In this case you

knowledge would be *a posteriori*, since listening to the testimony of someone involves sensory experience.

This brings us to a second important point about *a priori* and *a posteriori* knowledge. In determining whether an instance of knowledge falls into one of these two categories it is important to focus on what is your current reason for believing it rather than what initially prompted you to believe it. So, for example, suppose that one initially comes to believe that triangles have three sides because one trusts the testimony of one's schoolteacher. At this point, the knowledge is thus *a posteriori*. But suppose that one subsequently gains a better understanding of what a triangle is, and so one comes to recognise that it is in the very nature of a triangle that it has three sides. One's knowledge is now *a priori*, since the basis for the knowledge—one's recognition of the nature of triangles—is independent of sensory experience.

This brings us to a third point that we need to bear in mind when talking about *a priori* and *a posteriori* knowledge, which is what is meant by the claim that *a priori* knowledge is gained *independently* of sensory experience. One might be puzzled by this, since in order to know something one must be alive, right? And if one is alive, isn't one inevitably having some sensory experiences? So doesn't it follow that all knowledge is dependent on sensory experience? What this line of argument indicates is that when it is said that *a priori* knowledge is independent of sensory experience, the idea is not that it is knowledge that can be gained even while having no sensory experiences at all (which is almost certainly impossible). Indeed, as we saw a moment ago, one's knowledge of a particular proposition might be initially gained in an *a posteriori* fashion and then subsequently become *a priori*, so the 'independence' of *a priori* knowledge from sensory experience obviously cannot mean that one has no sensory experiences at all.

I think the best way to get a handle on what is meant by 'independence' in this regard is to ask the question of whether one's knowledge is dependent upon a particular investigation of the world. If that answer is 'yes'—as applies when one looks to see what colour one's eyes are, or one listens to what one's schoolteacher say—then the knowledge is *a posteriori*. But if the answer is 'no'—as when one's basis for the knowledge is purely the recognition of the fact that triangles must have three sides—then the knowledge is *a priori*.

A second distinction that we need to consider is that between **analytic and synthetic propositions**. An analytic proposition is one that is true solely in virtue of its meaning. A common example that is used to illustrate this is the proposition that *bachelors are unmarried men*. This is true in

virtue of its meaning because ‘bachelor’ simply means ‘unmarried man’. Hence, simply by knowing what a bachelor is, one can know that this proposition must be true. In contrast, consider a proposition like *Paris is the capital of France*. What makes this true is not only its meaning, but rather some fact about the world—i.e., that a certain city in France, Paris, has been made its capital city. Simply knowing the meanings of the words involved in this proposition would not be enough to determine whether it is true. One would instead need some further information (e.g., by looking in a reliable atlas).

Although the analytic/synthetic distinction is primarily applied to propositions, we can also apply it to knowledge. Analytic knowledge is knowledge of a proposition that is gained purely from knowing the meanings of the terms used. In contrast, synthetic knowledge is knowledge of a proposition that is not gained in this way (but by, e.g., making an investigation).

As we will see, one of the most fundamental debates in philosophy concerns how these two distinctions—i.e., between *a priori* and *a posteriori* knowledge, and between analytic and synthetic propositions—relate to one another.

Knowledge Empiricism

We noted above that *A posteriori* knowledge is also known as empirical knowledge, and this is reflected in the fact that an important historical movement in philosophy, known as empiricism, puts particular emphasis on this kind of knowledge. According to knowledge empiricism, there is a very straightforward connection between the analytic/synthetic and *a priori/a posteriori* distinctions. This is that all knowledge of synthetic propositions is *a posteriori*, while the only propositions that one can know *a priori* are analytic propositions. In contrast, **knowledge rationalism** denies this, and claims that some *a priori* knowledge is of synthetic propositions.

Before we consider the case for knowledge rationalism, let’s reflect a little on the plausibility of knowledge empiricism. At first blush, at least, the case seems very strong. Since the truth of synthetic propositions is not solely to do with the meaning of the words used, what makes them true thus appears to be facts about the world around us (think, for example, of our example concerning Paris above). But if that’s right, then wouldn’t it be plausible to suppose that knowledge of synthetic propositions must involve an investigation of the world (an *empirical investigation*, as it is sometimes called), and hence must be *a posteriori*?

That analytic propositions can be known *a priori* is uncontroversial, as recognising that something is true in virtue of understanding the meanings of the terms involved clearly does not involve an empirical investigation. But are analytical propositions the only kind of propositions that are known in this way? The reason why one might think this is plausible relates to the point just made about synthetic propositions being made true (at least in part) by the world rather than by the meanings of the terms employed. If that's right, then it's hard to see how one could come to know a synthetic proposition *a priori*, since to know such a proposition would surely always involve some sort of empirical investigation of the world.

There is thus a strong *prima facie* case for knowledge empiricism. One famous defender of this view was the Scottish philosopher **David Hume (1711-1776)**. He claims that there are just two kinds of knowledge. The first is knowledge of the relations between ideas, the second is knowledge of matters of fact (this distinction is sometimes called 'Hume's fork'). In the former case, this is knowledge that one can gain just by thinking about the ideas in question and how they relate to one another. What makes these claims true thus has nothing to do with the world around us, but solely depends on the nature of the ideas themselves. These claims can thus be known *a priori*, since one can know them without undertaking any empirical investigation. In contrast, the second kind of knowledge is concerned with the world around us, and hence can't be acquired (according to Hume) in an *a priori* manner, but essentially depends on an empirical investigation. It is thus *a posteriori* knowledge. Although Hume doesn't employ the particular terminology of analytic and synthetic propositions himself, it is nonetheless clear that he subscribes to knowledge empiricism. Knowledge of the relations between ideas is knowledge of analytical propositions, and only knowledge of this kind can be acquired *a priori*. In contrast, knowledge of matters of fact concerns synthetic propositions, and for Hume this knowledge is always *a posteriori*.

Hume's knowledge empiricism led him to make the following famous remark in his book, *An Enquiry Concerning Human Understanding*:

“If we take in our hand any volume; of divinity or school metaphysics, for instance; let us ask, Does it contain any abstract reasoning concerning quantity or number? No. Does it contain any experimental reasoning concerning matter of fact and existence? No. Commit it then to the flames: for it can contain nothing but sophistry and illusion.”

Hume's point was that a lot of what passes for genuine knowledge does not fall on either side of the 'fork' that he describes. But if that's right, then as far as Hume is concerned this means that it isn't genuine knowledge at all, but mere 'sophistry' that should be cast aside.

In contrast to knowledge empiricism, knowledge rationalism argues that at least some synthetic truths—i.e., truths about the world around us—can be known *a priori*. A famous example of this is supplied by the French philosopher **René Descartes (1596-1650)**. In his masterpiece, the *Meditations*, Descartes explores the limits of radical doubt about one's beliefs. To this end he puts forward what are known as **sceptical hypotheses**. These are scenarios which are indistinguishable from normal life, but where one is radically deceived. So, for example, one of the sceptical scenarios that he puts forward is that there could be an evil demon who is tricking him into forming false beliefs, not only about the world around him but also regarding such domains as mathematics (i.e., the demon tricks him into believing that $2 + 2 = 5$, say). Descartes concludes that just about any proposition is thus open to doubt. There is, however, one proposition that he claims is immune to doubt (i.e., indubitable), which is that he exists. This is where we get the famous remark 'cogito ergo sum' (I think therefore I am). For Descartes' point is that in doubting one's own existence one is thereby affirming that one exists. It is thus certain that one exists, and no sceptical hypothesis can call this into question.

Notice, however, that the proposition that one exists is a synthetic proposition. That is, it is clearly not something that is true purely in virtue of the meanings of the words in play, but rather concerns a fact about the world. And yet Descartes seems to have established the truth of this proposition via a completely *a priori* route (i.e., by reasoning about the inherent limitations of sceptical hypotheses). Hence, if Descartes's reasoning is correct anyway, it seems that one can have *a priori* knowledge of a synthetic proposition after all, contrary to what Hume has claimed with his knowledge empiricism.

Knowledge Innatism

The form of knowledge rationalism that we just saw Descartes defending claimed that one could, through reason, come to have *a priori* knowledge of a particular synthetic proposition. A different way of motivating knowledge rationalism is by claiming that some knowledge of synthetic propositions is innate. After all, if there is innate knowledge, then it must be *a priori*, since it obviously cannot have been acquired via an empirical investigation of the world. This view is known as **knowledge innatism**, and insofar as it is applied to synthetic propositions it is in conflict with knowledge empiricism.

Knowledge innatism is famously defended by **Plato (c. 427-347)** in the *Meno*. Through the mouthpiece of Socrates, a puzzle is posed, what is sometimes called the *paradox of inquiry*. Suppose we are trying to answer a certain question, such as what the nature of virtue is (one of the main concerns of the *Meno*). How would we recognise that we had the correct answer if we came across it? On the one hand, if we don't know what virtue is, then it is hard to see how we would recognise that a particular explanation of the nature of virtue is the correct one. Alternatively, if we are able to recognise that a particular explanation of the nature of virtue is the correct one, then doesn't that mean that we already knew what virtue was, since how else are we to be confident that this is the correct explanation?

Plato offers a radical resolution to this puzzle, which is to claim that learning is in fact a form of remembering. That is, the knowledge is already there—it is innate—it is just that we have to learn how to 'unlock' it. To illustrate this point, Plato has Socrates asking a boy a series of questions about geometry and in the process enabling him to come to recognise a geometrical truth. Plato's point is that the boy already had the knowledge in question, and what Socrates is doing by asking these questions—what is known as the method of *elenchus*—is merely bring out that existing knowledge (the *elenchus* is sometimes described as a form of midwifery for this very reason). So although an empirical process (the *elenchus*) might be required to draw out the knowledge, it is nonetheless *a priori* knowledge, since it was acquired independently of experience in virtue of being innate. (Moreover, notice that even after being drawn out in this way, the knowledge continues to be *a priori*, since it is based on the boy's recognition of the geometrical truth and not, for instance, on the testimony of Socrates).

Knowledge empiricists reject this line of thinking. Instead they argue that what is going on in these cases is not the extraction of existing innate *a priori* knowledge, but rather the acquisition of new *a priori* knowledge by coming to know analytical truths. As regards geometrical truths, this might initially seem fairly plausible, in that one might hold that they are analytically true—i.e., that they are simply concerned with relations within ideas, rather than dependent on the nature of the world. As it happens, however, it is not now widely held that geometrical truths are analytic—indeed, the standard view is that geometrical truths are about the nature of space itself, which is of course part of the nature of world around us. (Indeed, it turns out that the properties of space are very different to what many thought them to be in the ancient world). If that is right, and if Plato

has indeed shown that we can have innate geometrical knowledge, then there is a case for thinking that there can be innate, and thus *a priori*, knowledge of synthetic propositions.

One empiricist who was opposed to the idea of innate knowledge was the English philosopher **John Locke (1632-1704)**. In his book, *An Essay Concerning Human Understanding*, Locke argues that if there were such a thing as innate knowledge, then it would be universal (i.e., everyone would have it). But he maintains that there aren't any claims that are agreed upon by everyone. Moreover, there are lots of people (e.g., small children) who do not know the kinds of claims that defenders of innate knowledge cite as instances of knowledge of this kind (e.g., children do not know geometrical truths, for example). Instead, Locke maintains that all knowledge is acquired, either through reason or through experience (and in the latter case the knowledge will be *a posteriori* knowledge).

One problem for Locke's account is that it doesn't really engage with knowledge innatism as someone like Plato understands the view. After all, as we saw above, Plato doesn't think that small children are always able to articulate the innate knowledge that we have. Indeed, as with the case of geometrical knowledge that he offers, it may require the skill of someone like Socrates to draw this knowledge out. Thus Plato's understanding of innate knowledge is entirely compatible with what Locke says about there being no claims that are universally agreed upon.

Concept Empiricism

A different way of thinking about the empiricism/rationalism distinction is by considering how we acquire our concepts. **Concept empiricism** is the view that all concepts are derived from experience. In contrast, a kind of concept rationalism, known as **concept innatism**, contends that some concepts are innate (and hence not acquired via experience). The status of concept innatism is very important to the question of whether knowledge innatism is plausible, since innate knowledge is only possible if there can be innate concepts. Thus, if one can reject concept innatism then one can also reject knowledge innatism.

Locke was a defender of concept empiricism. He famously talked of the mind, at birth, as being a 'blank slate' (or 'tabula rasa'). His idea is that we come into the world with no innate concepts, but rather acquire those concepts over time through experience of the world and through reflection. Although he expresses his views slightly differently, Hume also endorses concept empiricism. According to Hume, all ideas derive from impressions (i.e., from sensory experience).

(Note, by the way, that by ‘ideas’ here Hume means both concepts and sensations). This means that if one lacks a particular kind of experience, then one will lack the ability to form the idea of that experience, and thus will lack certain concepts. A blind person, for example, will lack the idea of what colour is, and hence will lack colour concepts.

One objection to this proposal—an objection that Hume considers before dismissing as unimportant—is known as the *missing shade of blue problem*. If someone is presented with a spectrum of shades of blue with one shade missing, then it seems that one can, simply by using one’s imagination, form an idea of that shade. But if Hume’s concept empiricism is right, then how is that possible? After all, the person may have had no experience at all of that shade of blue.

Locke and Hume also maintain that complex concepts can be reducible to simple ideas or impressions, where the latter are effectively the conceptual ‘building blocks’ of our ideas. So we can form the idea of a *black dog* by bringing together the simpler ideas of blackness and dogs, ideas that we have acquired through having relevant experiences. This is how we are able to form ideas of things that we have had no experience of. Consider the concept of a unicorn. While we have had no experience of unicorns—since they don’t exist—we have had experience of the constituent parts of unicorns. Unicorns, after all, are white horses with horns, and since one has had experiences of these constituent parts, there need be nothing puzzling about the fact that one can form the complex concept of a unicorn. (Note that this doesn’t obviously help us with the missing shade of blue problem though, as this doesn’t seem to be a complex concept at all).

Concept Innatism

Locke’s arguments against concept innatism are—perhaps unsurprisingly—very similar to his arguments against knowledge innatism. For example, he emphasises the fact that small children do not appear to have any concepts at all, just as he emphasises their lack of knowledge. As with his arguments against knowledge innatism, however, this type of objection struggles to gain a purchase on concept innatism as it is usually understood. For just as defenders of knowledge innatism grant that innate knowledge often needs to be ‘extracted’ from the subject (e.g., through the elenctic method), so defenders of concept innatism often claim that certain experiences can be required to trigger our development of a concept, even though one’s possession of the concept is innate. This is, for example, essentially what the German philosopher **Gottfried Leibniz (1646-1716)** argued in his *New Essays Concerning Human Understanding*. He argues that to say that a concept is innate is to say

that we have an innate disposition to form the concept. It is thus entirely consistent with the notion of an innate concept that it is only manifested at a certain developmental stage.

Relatedly, Leibniz also notes, *contra* Locke, that we shouldn't confuse having an innate disposition to form a concept with having a particular word in mind. So, for example, it could be true that everyone is born with an innate disposition to form the concept of God while never developing mastery of a word to describe this concept. Leibniz's thought is that there might be a universal receptivity to God that only later, through experience, leads people to form the full concept of God (or something approximating it anyway). Other concepts which Leibniz claims are innate include unity, substance, being, action, and pleasure. These are concepts that we develop not through experience of the world, but rather through reflection, which for Leibniz involves attending to that which is already within us.

Summary

- ❑ We began with two distinctions that are crucial to this topic. The first was between *a priori* and *a posteriori* knowledge. This concerns two ways of knowing a proposition. One knows a proposition *a posteriori* if one's basis for that knowledge is via sensory experience. In contrast, if one's basis for knowing a proposition is not via sensory experience—e.g., if it is exclusively based on reasoning—then it is *a priori* knowledge.
- ❑ The second distinction was between analytic and synthetic propositions. An analytic proposition is a proposition that is true simply in virtue of its meaning, such as that bachelors are unmarried men. In contrast, a synthetic proposition is not true simply in virtue of its meaning, but rather depends upon further facts (e.g., that John is a bachelor).
- ❑ With these two distinctions in play, we were able to characterise some important philosophical positions. First, there is knowledge empiricism. This is the view that all knowledge of synthetic propositions is *a posteriori*, while all *a priori* knowledge is of analytic propositions. This is opposed to knowledge rationalism, which holds that some synthetic propositions can be known *a priori*.
- ❑ We looked at Hume's case for knowledge empiricism, and in particular the idea that since the truth of synthetic propositions depends on the nature of the world around us, so they can only be known *a posteriori*. Opposed to this, we also considered Descartes's argument

that one can know that one must exist. If this argument is right, then it seems that one can know a synthetic proposition purely via reason, and hence know it *a priori*.

- ❑ Next we looked at knowledge innatism, which is the idea that some of our knowledge is possessed from birth. If this is right, then this knowledge will be independent of experience and hence *a priori*. It thus potentially creates a problem for knowledge empiricism. We examined Plato's argument for knowledge innatism, via appeal to a puzzle known as the paradox of inquiry. We also looked at Locke's empiricist objections to knowledge innatism.
- ❑ Finally we examined a different dispute between empiricists and rationalists. Whereas concept empiricists hold that all our concepts are derived from experience, concept rationalists reject this and contend that some concepts are innate. We saw that this debate was very relevant to the debate about knowledge empiricism and knowledge rationalism, since if there can be no innate concepts, then it is hard to see how there could be any innate knowledge either (as one would need the relevant concepts in order to have the knowledge). We saw that Locke and Hume defended versions of concept empiricism, with the former contending that we are born with a 'blank slate' on which experience charts concepts. In contrast Leibniz defended concept innatism, claiming that, for instance, we all have an innate disposition to form the concept of God.

Study Questions

- 1 Try to describe, in your own words, the distinction between *a priori* and *a posteriori* knowledge.
- 2 Classify the following cases in terms of whether they are instances of *a priori* or *a posteriori* knowledge:
 - ❑ Reasoning one's way to a mathematical conclusion.
 - ❑ Perceiving that the apple is red, and forming a belief on this basis.
 - ❑ Believing a true mathematical claim by trusting the word of an expert.
 - ❑ Reflecting on the nature of a square and realising that all squares have four sides.
- 3 Try to describe, in your own words, the distinction between analytic and synthetic propositions.
- 4 Classify the following cases in terms of whether they are analytic or synthetic propositions:
 - ❑ Paris is the capital of France.

- Triangles have three sides.
 - Red is a colour.
 - The apple is red.
- 5 What is knowledge empiricism, and why might one hold such a view?
- 6 What is knowledge rationalism, and what grounds can be offered in its support?
- 7 What is knowledge innatism? What is the paradox of inquiry, and how might knowledge innatism be a way of resolving that paradox? Why might knowledge innatism be incompatible with knowledge empiricism?
- 8 What is concept empiricism, and why might one hold such a view?
- 9 What is concept innatism? Why might it be important to defenders of knowledge innatism that concept innatism is also true? Is concept innatism defensible, do you think?

Introductory Reading

- Warburton, N. (2014). *Philosophy: The Classics* (4th Edn.), (London: Routledge). [It is difficult to assign introductory readings on empiricism and rationalism because a lot of the literature on this debate is written at a very high level. That said, this book provides a very helpful commentary on a number of classic philosophical texts, including some of the texts that we have covered as part of this topic, such as Descartes' *Meditations*, Locke's *An Essay Concerning Human Understanding*, and Hume's *An Inquiry Concerning Human Understanding*].

Further Reading

- Thomas, J. (2009). *The Minds of the Moderns: Rationalism, Empiricism and Philosophy of Mind*, (London: Routledge). [An excellent book about the debates between empiricists and rationalists during the early modern period].

Free Internet Resources

- Baehr, J. (2015). '*A Priori* and *A Posteriori*', *Internet Encyclopaedia of Philosophy*, <http://www.iep.utm.edu/apriori/> >. [A very helpful discussion of the *a priori/a posteriori* distinction].

- Look, B. C. (2013). ‘Gottfried Wilhelm Leibniz’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/leibniz/>>. [A useful overview of Leibniz’s philosophical works, including a sub-section on his defence of innate concepts/ideas].
- Markie, P. (2013). ‘Rationalism vs. Empiricism’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/rationalism-empiricism/>>. [A comprehensive account of the debate between empiricism and rationalism, including sub-sections on both knowledge innatism and concept innatism].
- Morris, W. E., & Brown, C. R. (2014). ‘David Hume’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/hume/>>. [A wide-ranging account of Hume’s life and philosophical works, including his defence of empiricism].
- Russell, B. (2014). ‘*A Priori* Justification and Knowledge’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/apriori/>>. [An excellent overview of the issues surrounding *a priori* knowledge. Note, though that this can be rather hard-going in places].
- Rey, G. (2013). ‘The Analytic/Synthetic Distinction’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/analytic-synthetic/>>. [A very thorough discussion of the analytic/synthetic distinction].
- Smith, K. (2014). ‘Descartes’s Life and Works’, *Stanford Encyclopaedia of Philosophy*, <<http://plato.stanford.edu/entries/descartes-works/>>. [A very helpful overview of Descartes’s writings, with a sub-section devoted to his *Meditations*].

GENERAL RESOURCES

Textbooks

- ❑ Audi, R. (2010). *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (3rd Edn.) (London: Routledge). [An excellent and comprehensive textbook, though perhaps a little advanced in places for students new to epistemology].
- ❑ Bonjour, L., & Sosa, E. (2003). *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues* (Oxford: Blackwell). [It is not quite true to say that this is a textbook, since it in fact features two opposing essays from the main contributors, along with a critique and response from each contributor to the other. Nevertheless, it is an excellent way of getting an overview of some of the key issues in contemporary epistemology].
- ❑ Blaauw, M., & Pritchard, D. H. (2006). *Epistemology A-Z*, (Edinburgh: Edinburgh University Press). [A short and inexpensive dictionary of key terms, ideas and thinkers in epistemology.].
- ❑ Fumerton, R. (2006). *Epistemology*, (Oxford: Blackwell). [A short and very readable overview of the field].
- ❑ Martin, R. M. (2010). *Epistemology: A Beginner's Guide*, (London: Oneworld). [An ideal introduction to epistemology for those completely new to the field].
- ❑ Pritchard, D. H. (2015). *What is this Thing Called Knowledge?*, (London: Routledge, 3rd Edn.). [This comprehensive introduction to epistemology presupposes no previous knowledge of philosophy, and so is well-suited for students studying this topic for the first time].
- ❑ Pritchard, D. H. (2016). *Epistemology*, (London: Palgrave Macmillan). [This book is written for advanced undergraduates and aims to offer an opinionated overview of some of the main themes in contemporary epistemology. (Note that this is the retitled second edition of an earlier textbook entitled *Knowledge*)].

Edited Collections

- ❑ Alcock, L. M. (ed.) (1998). *Epistemology: The Big Questions*, (Oxford: Blackwell). [A good selection of articles, with more breadth than most collections, but as a consequence not quite so much depth].

- ❑ Bernecker, S. (ed.) (2006). *Reading Epistemology*, (Oxford: Blackwell). [A nice collection of articles, each of which is accompanied by a very useful commentary from the editor].
- ❑ Bernecker, S., & Dretske, F. (eds.) (2000) *Knowledge: Readings in Contemporary Epistemology*, (Oxford: Oxford University Press). [An excellent and well-priced anthology of articles in epistemology].
- ❑ Bernecker, S., & Pritchard, D. H. (eds.) (2010). *The Routledge Companion to Epistemology*, (London: Routledge). [A mammoth, completely up-to-date collection of articles on all the key areas of epistemology].
- ❑ Dancy, J., Sosa, E., & Steup, M. (eds.) (2010). *A Companion to Epistemology*, (Oxford: Blackwell, 2nd Edn.). [A very full list of entries. The second edition also includes 20 self-profiles from prominent epistemologists and 10 new review essays on central topics in epistemology. Very useful to have to hand].
- ❑ Greco, J., & Sosa, E. (eds.) (1999) *The Blackwell Guide to Epistemology*, (Oxford: Blackwell). [A series of introductory articles on the main topics in epistemology. A very good collection of papers, if a little dated now].
- ❑ Hetherington, S. (ed.) (2012) *Epistemology: The Key Thinkers*, (London: Continuum). Helpful profiles of some of the main figures in epistemology, going right back to the ancients and extending up to the present day.
- ❑ Moser, P. K. (ed.) (2002) *The Oxford Handbook of Epistemology*, (Oxford: Oxford University Press). [Contains lots of essays on the main topics in the area, written by the key figures involved].
- ❑ Neta, R., & Pritchard, D. H. (eds.) (2008). *Arguing About Knowledge*, (London: Routledge). [This collection aims to cover the main themes in epistemology by offering a selection of articles that present a ‘for and against’ treatment of the relevant positions.].
- ❑ Sosa, E., Kim, J., Fantl, J., & McGrath, M. (eds.) (2008). *Epistemology: An Anthology*, (Oxford: Blackwell, 2nd Edn.). [An excellent and well-priced anthology of articles that has recently been updated to include a number of new articles on recent developments in the epistemological literature].
- ❑ Steup, M., Turri, J., & Sosa, E. (eds.) (2013). *Contemporary Debates in Epistemology*, (Oxford: Blackwell, 2nd Edn.). [An excellent idea: the main figures in the literature offer alternative perspectives on a key issue, and then respond to each other’s articles. The second edition

includes several newly commissioned exchanges, and covers all the core topics in contemporary epistemology. Very up-to-date].

Free Internet Resources

- ❑ *Internet Encyclopaedia of Philosophy*, (eds.) J. Fieser & B. Dowden, <http://www.iep.utm.edu>. [The second-best completely free internet encyclopedia of philosophy available. It's not quite as comprehensive or authoritative as the *Stanford Encyclopedia of Philosophy* (see below), but still contains some very good entries on epistemology].
- ❑ *Stanford Encyclopaedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu>. [The best completely free internet encyclopedia of philosophy available. It's continually being updated, and has many great articles on epistemology].

GLOSSARY

Ability knowledge

This is often referred to as ‘know-how’, since it involves knowing how to do something, such as ride a bike or swim. It is usually contrasted with propositional knowledge, which is knowledge of a proposition. The two types of knowledge are treated differently because, intuitively at least, one might know how to do something (e.g., swim) without having any relevant propositional knowledge (e.g., without knowing that you can swim, perhaps because you forgot that you could until you fell in the water). *See also* **propositional knowledge**.

Analytic/synthetic propositions

A proposition that is analytic is true by virtue of its meaning alone. For example, that bachelors are unmarried men is simply true by definition. In contrast, a synthetic proposition is one that not true by virtue of its meaning along. That John is a bachelor, for example, is not true in virtue of its meaning alone, but rather depends rather on further facts about John (e.g., that he isn’t married).

***A priori/a posteriori* knowledge**

The distinction between *a priori* and *a posteriori* knowledge—note that the latter is sometimes known as *empirical knowledge*—relates to whether the knowledge in question was gained independently of an investigation of the world (what is known as an empirical inquiry). If it was, it is *a priori* knowledge; if it wasn’t, it is *a posteriori* knowledge. For example, my knowledge that Minsk is the capital of Belarus is empirical knowledge because I gained it by making an investigation of the world (e.g., I looked it up in an atlas). In contrast, my knowledge that all bachelors are unmarried is *a priori* knowledge, because I gained it by reflecting on what the words mean and so no investigation of the world was required (though note that I could have gained this knowledge empirically, by asking someone, for example).

Argument from illusion

Consider the visual impression caused by a genuine sighting of an oasis on the horizon and contrast it with the corresponding visual impression of an illusory sighting of an oasis on the horizon formed by one who is hallucinating. Here is the crux: these two visual impressions could be exactly the same. The problem, however, is that it seems that if this is the case then what we experience in perception is not the world itself, but something that falls short of the world, something that is common to both the ‘good’ case in which one’s senses are not being deceived (and one is actually looking at an oasis) and the ‘bad’ case in which one’s senses are being deceived (and one is the victim of an hallucination). This line of reasoning, which makes use of undetectable error in perception in order to highlight the indirectness of perceptual experience, is known as the argument from illusion. It suggests an ‘indirect’ model of perceptual knowledge, such that what we are immediately aware of when we gain such knowledge is a sensory impression—a *seeming*—on the basis of which we then make an inference regarding how the world is. *See also* **indirect realism**.

Berkeley, George (1685-1753)

George Berkeley, otherwise known as Bishop Berkeley (he was the Bishop of Cloyne in what is now the Republic of Ireland), was, like John Locke (1632–1704) and David Hume (1711–1776), an empiricist. An empiricist is someone who believes that all knowledge of substance is ultimately derivable from experience. (Locke, Berkeley and Hume are collectively known as the *British empiricists*.) Unlike Locke and Hume, however, he famously saw in empiricism a motivation for idealism—the view that there is no mind-independent world. Berkeley led a very interesting life, including a spell living in Bermuda. He also has the unusual distinction of having a city (and a university) named after him, the city of Berkeley in California. *See also idealism; knowledge empiricism.*

Classical account of knowledge

According to classical account of knowledge, knowledge is defined as justified true belief. This view is often credited to Plato, and is sometimes referred to as the ‘tripartite’ (i.e., three-part) account of knowledge. *See also Gettier cases.*

Concept empiricism

Concept empiricism is the view that all of our concepts are derived from experience. It is opposed to concept innatism, which holds that some of our concepts are innate. *See also concept innatism; knowledge empiricism; knowledge innatism.*

Concept innatism

Concept innatism is the view that some of our concepts are innate. It is thus opposed to concept empiricism, which claims that all of our concepts are derived from experience. *See also concept empiricism; knowledge empiricism; knowledge innatism.*

Descartes, René (1596–1650)

The French philosopher and mathematician René Descartes is one of the founding fathers of modern philosophy. His most famous work is his *Meditations on First Philosophy* in which, amongst other things, he offers a radically new way of approaching epistemology. Descartes’ idea is that in order to put our knowledge on a secure foundation, it is necessary to first subject it to what he called the ‘method of doubt’. This involves doubting as much as can be doubted amongst one’s beliefs until one finds the indubitable, and thus epistemologically secure, foundation on which one’s knowledge can be built. In the service of this end, Descartes put forward a number of radical sceptical hypotheses—scenarios which are indistinguishable from normal experience, but in which one is radically in error, such as that one’s experiences are a product of a dream—in order to discover which of his beliefs were immune to doubt. By applying the method of doubt, Descartes was led to the conclusion that the indubitable foundation of our knowledge is our belief in our own existence, since in doubting our existence we thereby prove that we exist (since how else could we be able to doubt?). Hence the famous claim, ‘I think, therefore I am’. Ironically, the powerful sceptical arguments that Descartes invented have held more sway than his subsequent anti-sceptical arguments. Accordingly, although it was not his aim to make us sceptical about the possibility of knowledge, this is in fact what his epistemological investigations seem to have achieved. Aside from his work in epistemology, Descartes made important contributions to just about every other area of philosophy as well. In addition, he also conducted research

on scientific and mathematical questions, making a long-standing contribution to, for example, geometry.

Epistemology

This is the name given for the theory of knowledge. Those who study epistemology—known as *epistemologists*—are also interested in those notions closely associated with knowledge, such as truth, justification and rationality.

Epistemic virtues

An epistemic virtue is a reliable cognitive ability that makes you better suited to gaining true beliefs. As epistemic virtues are usually understood—as part of a wider virtue epistemology—they can include both innate cognitive faculties, such as one’s perceptual faculties, and also quite sophisticated intellectual traits, such as conscientiousness. That one is a reliable perceiver will obviously enable one to reliably form true beliefs about one’s environment. But that one is a conscientious inquirer, in that one is careful to avoid error and takes all the available evidence into account, will also help one to form more accurate beliefs too. *See also* **virtue epistemology**.

Gettier cases

Gettier cases are scenarios in which an agent has a justified true belief and yet lacks knowledge because it is substantially due to luck that the belief in question is true. Imagine someone who forms her belief about what the time is by looking at a stopped clock that she has every reason for thinking is working. Crucially, however, she happens to look at the clock at the one time in the day when it is showing the right time, and so forms a true belief as a result. Her belief is thus both true and justified, yet it isn’t a case of knowledge since it is just luck that her belief is true given that the clock is not working. Gettier cases show that the classical account of knowledge that analyses knowledge into justified true belief is unsustainable. *See also* **classical account of knowledge**.

Hume, David (1711-1776)

David Hume is one of Scotland’s most important philosophers, and arguably the greatest ever philosopher to write in the English language. Born in Edinburgh, he led an interesting and varied life, writing a celebrated history of England as well as a number of central works in philosophy. Possibly his greatest work, *A Treatise on Human Nature*, was completed by the time he was 26. Hume’s intellectual achievements made him a key figure in a period of history known as the Age of Enlightenment, a time of great intellectual ferment. Hume’s philosophy is characterised by his empiricism, which is the belief that all knowledge is ultimately traceable back to the senses. Hume’s empiricism led him to be sceptical (see scepticism) about a lot of things that his contemporaries took for granted, particularly when it came to religious belief. His scepticism about religious belief created lots of personal obstacles for him for example, it was what prevented him from taking up a Chair in Philosophy at the University of Edinburgh, despite being by far the most deserving candidate but he pressed on regardless. Because of his ardent empiricism, Hume is often described, along with George Berkeley (1685–1753) and John Locke (1632–1704), as one of the *British empiricists*. *See also* **knowledge empiricism**.

Idealism

Idealism is the view that there is no external world (i.e., no world that is independent of our experience). In its simplest form, the view is not very appealing since it entails that the world ceases to exist when it is not being experienced. (For example, in order for a tree to fall in a forest, it is essential that there be someone present to perceive it fall.) In order to make the view more appealing, philosophers have supplemented the view in various ways. For example, George Berkeley (1685–1753) gets around some of the more counter-intuitive aspects of the view by arguing that God is always present and perceives everything, and thus the world does not cease to exist when it is not being experienced.

Indirect Realism

According to the argument from illusion, one's experiences when one is perceiving normally could be exactly the same as the experiences one would have were one to be deceived in some way (e.g., if one were having an hallucination). Indirect realists embrace the conclusion of this argument by claiming, in opposition to direct realists, that one never directly experiences the world in perception. Instead, one experiences only how the world seems to one, and on this basis one must make inferences regarding how the world is independently of how it appears. *See also* **argument from illusion; naïve realism.**

Infallibilism

Infallibilism is the view that in order to have knowledge one must have a belief that is infallible. *See also* **infallibility.**

Infallibility

If one's belief is infallible, then it could not be in error. *See also* **infallibilism.**

Knowledge empiricism

Knowledge empiricism is the view that all knowledge of synthetic propositions is *a posteriori* knowledge, while all *a priori* knowledge is of analytic propositions. It is opposed to knowledge rationalism, which is the view that some synthetic propositions can be known a priori. *See also* **analytic/synthetic propositions; a priori/a posteriori knowledge; knowledge rationalism; knowledge innatism.**

Knowledge innatism

Knowledge innatism is the view that some knowledge is innate, in the sense that it is not acquired via experience but is rather in some sense possessed from birth. If there is innate knowledge, then it must be *a priori* knowledge. *See also* **a priori/a posteriori knowledge; knowledge empiricism.**

Knowledge rationalism

Knowledge rationalism is the view that some synthetic propositions can be known *a priori*. It is opposed to knowledge empiricism, which is the view that all knowledge of synthetic propositions is a posteriori knowledge, while all *a priori* knowledge is of analytic propositions. *See also* **analytic/synthetic propositions; a priori/a posteriori knowledge; knowledge empiricism; knowledge innatism.**

Leibniz, Gottfried (1646-1716)

Gottfried Wilhelm von Leibniz was a German philosopher who not only made important contributions to philosophy, but advanced a number of other intellectual fields too, most notably mathematics. He is associated with the rationalist school of philosophy through his use of reason to determine important philosophical truths. Leibniz famously argued in his book, the *Theodicy*, that the world in which we live, though apparently imperfect in various ways (e.g., because of natural disasters and so forth), is in fact the best of all possible worlds that God could have created. *See also* **knowledge rationalism**.

Locke, John (1632–1704)

The English philosopher John Locke is perhaps most noted for his work on political theory, especially regarding the limits of the power of the state. Indeed, Locke’s broadly liberal conception of the role of the state was very influential on the establishment of the US constitution. In his philosophy more generally, Locke belongs to a school of thought known as *empiricism*, which traces all knowledge of any substance back to sensory experience. Along with George Berkeley (1685–1753) and David Hume (1711–1776), Locke is often referred to as one of the British *empiricists*. This commitment to empiricism is reflected in his famous claim that the mind at birth is like a “tabula rasa”—that is, like a ‘blank slate’ on which nothing is written. What Locke means by this is that there are no innate ideas. Instead, all our ideas, and thus our knowledge, are derived via experience of the world. *See also* **knowledge empiricism**.

Naïve realism

Naïve realism is a thesis about perceptual experience that has ramifications for perceptual knowledge. It holds that, at least in non-deceived cases, what we are aware of in perceptual experience is the external world itself. That is, if I am genuinely looking at an oasis on the horizon right now, then I am directly aware of the oasis itself, and thus I can have perceptual knowledge that there is an oasis before me without needing to make an inference from the way the world seems to how it is. *See also* **argument from illusion; indirect realism**.

No false lemmas

The no false lemmas account is a way of responding of supplementing the classical account of knowledge so that it can avoid Gettier-style cases. The general idea is that knowledge is justified true belief that is not based on any relevant presuppositions. The chief problem facing the account is to find a principled way of formulating this condition such that it is not so strong that it rules out *bona fide* cases of knowledge or so weak that it fails to exclude all Gettier-style cases. *See also* **classical account of knowledge; Gettier cases**.

Plato (c. 427–c. 347 BC)

Plato is one of the most influential philosophers who ever lived. He resided for most of his life in the Greek city of Athens, in Greece, which is where he came under the influence of Socrates (470–399 BC) and where he in turn influenced the philosophical development of Aristotle (384–322 BC). After Socrates’ death—an account of which is offered in Plato’s book, *The Apology*—Plato founded ‘The Academy’, a kind of early university in which a range of topics, but principally philosophy, was taught. Plato’s writing was often in the style of a dialogue between Socrates, the mouthpiece of Plato, and an imagined adversary (or

adversaries) on topics of vital philosophical importance. In *The Republic*, for example (perhaps his most famous work), he examines the question, central to political philosophy, of what the ideal political state is. Of more interest for our purposes, however, is his book the *Theaetetus*, in which he discusses the nature of knowledge.

Primary/secondary qualities

This is a distinction that was drawn (in modern times) by the philosopher John Locke (1632–1704). A primary quality is a feature of an object that the object has independently of anyone perceiving the object, whilst an object's secondary qualities are dependent upon the perception of an agent. A good example of a primary quality is shape, in that the shape of an object is not in any way dependent upon anyone perceiving that object. Compare shape in this respect with colour. The colour of an object is a secondary quality in that it depends upon a perceiver. If human beings were kitted-out with different perceptual faculties, then colours would be discriminated very differently. *See also indirect realism.*

Problem of the external world

According to the argument from illusion, all that I am directly aware of in perceptual experience is how the world appears, not how it is independently of how it appears. If all that I am directly aware of in perceptual experience is the way the world appears, however, then this opens up the possibility that the way the world appears might be no guide at all to how the world is (there is nothing about my experiences that would indicate that this is not the case after all). This is the problem of the external world (i.e., a world 'external' to our experience of it), which is a standard way of motivating radical scepticism about our knowledge of the external world. *See also argument from illusion; indirect realism; naïve realism; scepticism.*

Proposition

A proposition is what is stated by a declarative sentence. For example, the sentence 'The cat is on the mat' states that something is the case; namely, that the cat is on the mat, and this is the proposition expressed by this sentence. Notice that the same proposition will be expressed by an analogue declarative sentence, which is in a different language, such as French, just so long as what is stated by that sentence is the same.

Propositional knowledge

This is knowledge that something (i.e., a proposition) is the case. It is typically contrasted with ability knowledge, or *know-how*. The two types of knowledge are treated differently because, intuitively at least, one might know how to do something (e.g., swim) without having any relevant propositional knowledge (e.g., without knowing that you can swim, perhaps because you forgot that you could until you fell in the water). *See also ability knowledge.*

Russell, Bertrand (1872-1970)

Bertrand Russell was one of the most influential British philosophers of the last century. As a young man, he rebelled against idealism, which then dominated British philosophy. Instead, he wrote, with Alfred Lord Whitehead (1861-1947), one of the founding books of analytical philosophy, *Principia Mathematica*. This was an ambitious attempt to provide a

philosophical basis for mathematics. Although it was in fundamental ways unsuccessful, it led to many important developments, both within philosophy and beyond (e.g., in mathematics, computer science, and cognitive science). In his later life his focus became more on politics than philosophy, and he was a prominent proponent of nuclear disarmament. In 1950 he was awarded the Nobel Prize for Literature, largely in recognition of his masterly book, *A History of Western Philosophy*.

Sceptical hypotheses

A sceptical hypothesis is a scenario in which you are radically deceived about the world and yet your experience of the world is exactly as it would be had you not been deceived.

Consider, for example, the fate of the protagonist in the film *The Matrix*, who comes to realise that his previous experiences of the world were in fact being ‘fed’ into his brain whilst his body was confined to a large vat. Accordingly, whilst he seemed to be experiencing a world rich with interaction between himself and other people, in fact he was not interacting with anybody or any thing at all (at least over and above the tubes in the vat that were ‘feeding’ him his experiences), but was instead simply floating motionlessly. The problem posed by sceptical hypotheses is that we seem unable to know that they are false. After all, if our experience of the world could be exactly as it is and yet we are the victims of a sceptical hypothesis, then on what basis could we ever hope to distinguish a genuine experience of the world from an illusory one? Sceptical hypotheses are thus used to motivate scepticism. *See also scepticism.*

Scepticism

To advance scepticism about a certain subject matter is to argue that it is impossible to have any knowledge of that subject matter. For example, scepticism about the existence of other minds would be the view that it is impossible to know that there exist other minds. Radical scepticism is a form of scepticism that targets a very broad subject matter. For example, one form of radical scepticism argues that we are unable to know anything at all about the external world (i.e., a world that is ‘external’ to our experience of it). Although it is natural to speak of radical scepticism as being a philosophical position, it is not usually advanced in this way but is rather put forward as a challenge to existing theories of knowledge to show why they don’t generate the type of radical scepticism in question. *See also problem of the external world.*

Solipsism

Solipsism is the view that the only thing that exists is one’s own mind. On this view, then, both the external world and other minds are to be understood as lacking any independent existence, as they exist only as elements of one’s own mind. As such, solipsism is a particularly radical form of idealism. Indeed, one kind of objection against idealism is that it collapses into solipsism. *See also idealism.*

Virtue epistemology

A virtue epistemology is any theory of knowledge, which holds that knowledge is true belief that is gained as a result of the operation of reliable epistemic virtues or cognitive faculties. One version of this thesis is simply a refinement of a simple form of reliabilism. Whereas

reliabilism in its most basic form holds that one can gain knowledge through any reliable belief-forming process, the virtue epistemologist of this sort claims that only certain reliable belief-forming processes are knowledge-conducive (i.e., those which are epistemic virtues or cognitive faculties of the agent). In common with reliabilism, this form of virtue epistemology is a form of epistemic externalism, in that it holds that an agent can have knowledge simply by forming a true belief via one of her reliable cognitive faculties, even if she lacks good grounds to back up that belief. In contrast, there are versions of virtue epistemology, which are allied to epistemic internalism rather than epistemic externalism, and so claim that it is essential that a knowing agent is able to offer good grounds in favour of what she believes. This form of virtue epistemology holds that it is essential that one gains one's true belief via one's epistemic virtues, the thinking being that one cannot correctly employ one's epistemic virtues without thereby acquiring good grounds in favour of what one believes. *See also* **epistemic virtues**.