

# Knowledge

*The Internet Encyclopedia of Philosophy* (IEP) <https://www.iep.utm.edu/knowledg/>

Philosophy's history of reflection upon knowledge is a history of theses and theories; but no less of questions, concepts, distinctions, syntheses, and taxonomies. All of these will appear in this article. They generate, colour, and refine these philosophical theses and theories about knowledge. The results are *epistemological* — philosophical attempts to understand whatever is most fundamentally understandable about the nature and availability of knowledge. We will gain a sense of what philosophers have thought knowledge is and might be, along with why some philosophers have thought knowledge both does not and could not exist.

Thus, we will examine some of the general kinds or forms of knowledge that epistemologists have thought it important to highlight (section 1), followed by the idea of knowledge as a kind or phenomenon at all (section 2). Knowledge seems to be something we gain as we live; how do we gain it, though? That will be our next question (section 3), before we ask whether our apparently gaining knowledge is an illusion: might no one ever really gain knowledge (section 4)?

## 1. Kinds of Knowledge

We talk of knowledge: all of us do; philosophers do. But what *is* knowledge? We can best answer that potentially complex question in several stages. Let us begin by considering whether there are different kinds of knowledge. Epistemologists have contemplated at least the following general possibilities.

### a. Knowing by Acquaintance

Your knowing a person, it seems, involves direct interaction with him or her. Otherwise, at most, you should claim only that it is *almost as if* you know him or her: 'I've seen and heard so much about her that I feel like I know her. I wonder whether I'll ever meet her — whether I will ever actually know her.' Without that meeting, you could well know facts about the person (this being a kind of knowledge to be discussed in section 1.b). Nonetheless, could you know facts about a person without ever meeting him or her? If so, there could well be a kind of knowledge which is different to knowing a fact; maybe knowing a thing or entity (such as a person) is distinct from knowing a fact about that thing or entity.

Bertrand Russell (1959 [1912]: ch. 5) famously distinguished between knowledge by description and a quite particular kind of knowledge by acquaintance. He allowed there to be a form of acquaintance that was immediate and unquestionable, linking one with such things as abstract properties and momentary sensory items passing before one's mind: you can be acquainted with the abstract property of redness, as well as with a specific patch of redness briefly in your visual field. Knowledge by description was the means by which, in Russell's view, a person could proceed to know about what he or she had not experienced directly. We formulate definite descriptions ('the third man listed in the current Sydney residential phonebook') and indefinite ones ('a man listed in the current Sydney residential phonebook'). With these, we can designate individuals with whom we have not interacted. Then we can formulate claims using such descriptions. Some of these claims could be knowledge. Thus, we may open up for ourselves a world of knowledge beyond what is revealed by our immediate experiences.

### b. Knowledge-That

Most philosophical discussion of knowledge is directed at knowledge-that — such as knowledge that kangaroos hop, knowledge that koalas sleep most of the time, knowledge that kookaburras cackle, and the like. This is generally called propositional knowledge (a proposition *that such-and-such is so* is the object of the knowledge), declarative knowledge (the knowledge's object is represented by a

declarative sentence: ‘Such-and-such is so’), or knowledge-that (the knowledge is represented in the form ‘*that* such-and-such is so’). Knowledge by description (mentioned in section 1.a) would be one form that could be taken by knowledge-that: some known propositions include descriptions; but not all do. In principle, knowledge-that is the kind of knowledge present whenever there is knowledge of a fact or truth — no matter what type of fact or truth is involved: knowledge that  $2 + 2 = 4$ ; knowledge that rape is cruel; knowledge that there is gravity; and so on. When philosophers use the term ‘know’ unqualifiedly, knowledge-that is standardly what they mean to be designating. (It will therefore be the intended sense throughout most of this article.)

### **c. Knowledge-Wh**

But should knowledge-that receive such sustained and uninterrupted focus by philosophers? After all, there is a far wider range of ways in which we talk and think, using the term ‘know’. Here are some of them (collectively referred to as knowledge-wh): knowing *whether* it is 2 p.m.; knowing *who* is due to visit; knowing *why* a visit is needed; knowing *what* the visit is meant to accomplish; knowing *how* that outcome is best accomplished; and so forth.

How should these be understood? The usual view among epistemologists is that these are specific sorts of knowledge-that. For example, knowing whether it is 2 p.m. is knowing that it is 2 p.m., if it is; and knowing that it is not 2 p.m., if it is not. Knowing who is due to visit is knowing, for some specified person, that it is he or she who is due to visit. Knowing what the visit is meant to accomplish is knowing, for some specified outcome, that it is what the visit is meant to accomplish. Knowing how that outcome is best accomplished is knowing, for some specified description of how that outcome could be accomplished, that this describes the best way of accomplishing that outcome. And so on.

Still, not everyone will assess these examples in quite that way. Note a variation on this theme that is currently being developed. Called *contrastivism*, its basic idea is that (perhaps always; at least sometimes) to know is to know this *rather than* that. (For different versions, see Schaffer 2005; 2007; Morton 2011.) One’s knowing, understood contrastively, is explicitly one’s knowing one from among some understood or presumed bunch of possible alternatives. The word ‘explicitly’ is used here because one would know while acknowledging those alternatives. Consider the example of knowing-who. On contrastivism, you could know that it is Fred *rather than* Arjuna and Diego who is due to visit; and this might be the only way in which you know that Fred is due. ‘Who is due?’ ‘Fred, as against Arjuna or Diego.’ Your knowing-who would not be simply your knowing, of Fred, that *it is he* who is due to visit. Your knowing-who would be your knowing that *it is Fred as against Arjuna or Diego* who is due to visit. This remains propositional knowledge, nonetheless.

### **d. Knowing-How**

Gilbert Ryle (1971 [1946]; 1949) made apparent to other philosophers the potential importance of distinguishing knowledge-that from knowledge-how. The latter is not (thought Ryle) one’s knowing how *it is that* something is so; this, we noted in section 1.c, is quite likely a form of knowledge-that. What Ryle meant by ‘knowing how’ was one’s knowing how *to do* something: knowing how to read the time on a clock, knowing how to call a friend, knowing how to cook a particular meal, and so forth. These seem to be skills or at least abilities. Are they not simply another form of knowledge-that? Ryle argued for their distinctness from knowledge-that; and often knowledge-how is termed ‘practical knowledge’. Is one’s knowing how to cook a particular meal really only one’s knowing a lot of truths — having much knowledge-that — bearing upon ingredients, combinations, timing, and the like? If Ryle was right, knowing-how is somehow distinct: even *if* it involves having relevant knowledge-that, it is also something more — so that what makes it *knowledge-how* need not be knowledge-*that*. [For more on this issue, see, for example, Bengson and Moffett 2012). Might

knowledge-that even be a kind of knowledge-how itself, so that all instances of knowledge-that themselves are skills or abilities (for example, Hetherington 2011a: ch. 2)?]

## 2. Knowledge as a Kind

Section 1 shows how there might be different kinds of knowledge. We will now focus on one of them — knowledge-that. What kind of thing is such knowledge? In particular, is it a natural kind — a naturally occurring element in the scientifically describable world? Alternatively, is knowledge at least partly a conventional or artifactual kind — a part of our practices of judging and evaluating, possessing a socially describable nature?

The former idea portrays knowledge as an identifiable and explanatory aspect of what it is for beings relevantly like us to function as a natural component of a natural world. We have beliefs, some of which help us to achieve our aims by telling us how not to ‘bump into’ the world around us. We can ‘fit into’ — by ‘finding our way within’ — the world by using beliefs. Is that because these beliefs are knowledge? Is that part of why humans as a natural kind (if this is what we are) have prospered so markedly? In introducing epistemologists to the idea of what he called a naturalized epistemology, W. V. Quine (1969) recommended that philosophy conceive of us in psychological terms, so that when it seeks to understand us as reasoning, as believing, and as rational, it does not do this in terms distinct from those scientific ways of describing our psychological and physical features. Hilary Kornblith (2002) continues that theme: in effect, we know as other animals do — limitedly but reliably, thanks to our roles as sensing and believing beings operating within the world’s natural order. There would be natural laws, say, or at least natural regularities — scientifically formulable ones, we may hope — about how we know.

In contrast, we may feel that knowing is a distinctively *conventional* accomplishment. It might consist of socially constituted and approved patterns — not thereby natural laws or regularities admitting of scientific description — in aspects of how we interact with other people. Perhaps we can collectively choose what to count as knowledge. Perhaps that is all there is to knowing. Such a view could even say that this is how knowledge differs from belief: beliefs happen to or within us; knowledge we shape from beliefs. And we might do this deliberately, subjecting ourselves and others to social norms of inquiry, responding to other people and their concepts, aims, and values. As civilizations expand and mutate, could knowing change not only its content (that is, what is known), but its basic nature (for example, how the knowing occurs and even what in general is required for it to occur)? Different social arrangements would bring into being different ways of thinking and acting, new aims and values. In that sense, possibly knowledge is an artefact, created by us in social groupings, used by us in those same groupings — often wittingly and deliberately so. In short, maybe knowing is a matter of functioning in *socially apt* ways. Barry Allen (2004) is one who argues for an artifactual interpretation of knowing’s nature.

The rest of this article will remain neutral between these two broad ideas. Some of the suggestions to be considered will be more appropriate (and clearly so) for one than the other of the two. But in general the article’s aim will be to display, not to favour.

## 3. Ways of Knowing

To say the least, not everyone knows everything, not even everything that in principle is knowable. Individual instances of knowledge come to individual people at individual times, remaining in place for varying — individual — lengths of time. So it is right to ask how it is that individual cases of knowledge reach, or are acquired by, people; along with how it is that these cases of knowledge are then retained by people. In what broadly characterisable ways do people gain and maintain their knowledge? In practice, philosophers do not treat that as a question about the ineliminable specificities of each person, each moment, and each particular piece of knowledge. It is treated as a question about *general* ways and means of coming to know a specific fact or truth.

Over the centuries, these have been some of the more philosophically pondered forms of answer to that question:

Some or all knowledge is innate. (And then it is remembered later, during life.)

Some or all knowledge is observational. Some or all knowledge is non-observational, attained by thought alone.

Some or all knowledge is partly observational and partly not — attained at once by observing *and* thinking.

The rest of this section will consider these in turn.

### **a. Innate Knowledge**

If some instances of knowledge accompany a person into life, how will they reveal themselves within his or her life? How would the person, or indeed anyone else, know that he or she has this innate knowledge? It could depend on *what* is being known innately — the subject matter of this knowledge with which the person has been born.

For example, if people begin life already knowing some grammatical rules (an idea famously due to Noam Chomsky: see Stich 1975, ch. 4), this innate knowledge would be shown in subsequent speedy, widespread, and reliable language-learning by those involved. These instances of people learning so readily and predictably would be actions expressing some knowledge-*how*. But (as section 1.d acknowledged) such manifestations of knowledge-how might actually reflect the presence within of knowledge-*that*.

Or consider another possible example: knowledge of some mathematics and some logical principles. Seemingly, Plato (in the *Meno*, one of his dialogues) accorded people this sort of innate knowledge; as did Leibniz, in his *New Essays*. (For excerpts from Plato and from Leibniz, see Stich 1975, ch. 2.) Plato presented us with a story of a slaveboy, lacking education, whom Socrates brought, via minimal questioning, to a state of remembering some geometrical knowledge.

Naturally, it could be difficult to ascertain that any particular knowledge is genuinely innate. Knowledge which is not innate, but which is acquired especially easily, seemingly effortlessly, might nonetheless feel innate. And (as section 1.d also acknowledged) even when an action, such as of language-learning, is manifesting knowledge-how, there remains a philosophical question as to whether that action is reflecting knowledge-that already existing within, dormant until activated. The answer to that question might be that there is only knowledge-how present — without owing its existence to some related prior knowledge-that. (As ever throughout this article these possibilities are suggested for continued consideration, not as manifestly decisive refutations.)

### **b. Observational Knowledge**

One of epistemology's perennially central topics has been that of observational knowledge. Let us consider a few of the vast number of philosophical questions that have arisen about such knowledge. *Can there be purely or directly observational knowledge?* When you observe a cat sleeping in front of you, do you know observationally — and only observationally — that the cat is sleeping there? Observation is occurring; and you do not consciously 'construct' the knowledge. Still, is there a perceptual experience present, *along with* some conceptual or even theoretical knowledge (for example, that cats are thus-and-so, that to sleep is to do this-and-not-that, and so forth)? Otherwise, how could your experience constitute your *knowing* this-content-rather-than-another? Is conceptual knowledge what gives knowledgeable content to your observational experience? Is this so, even for experiences that are as simple as you can imagine having?

*Can there be foundational observational knowledge?* Wilfrid Sellars (1963) engaged famously with this question, confronting what he called the myth of the given. Part of the traditional epistemological appeal of the idea of there being purely or directly observational knowledge was the idea that such knowledge could be foundational knowledge. It would be knowledge given to us in experiences which would be cases of knowledge, yet which would be conceptually simple. Sellars argued, however, that they would not be conceptually so simple.

For example, imagine knowing observationally that *here is something white*. This would possibly be as simple, in conceptual terms, as observational knowledge could be for you. Nevertheless, even here

the question remains of whether you are applying concepts (such as of being here, of being something, and of being white); and if you are doing so, of whether you must be able to know that you are using them correctly. Would you need to find even simpler observational experiences, via which you could know what these concepts involve? If so, the other experience — knowing observationally that *here is something white* — would not have been foundational. That is, it would not have amounted to a basic piece of knowledge, upon which other pieces of knowledge can be based and which need not itself be based upon other pieces of knowledge.

*How much observation is needed for observational knowledge?* When you look at what appears to be a cat, for how long must you maintain your gaze if you are to know that you are seeing a cat? Do you need also to walk around it, still looking at it, scrutinising it from different angles, if you are to know that you are seeing a cat? And what of your other senses? Could the animal's sounding or smelling like a cat, for example, be needed if the knowledge in question is to be yours? There is a more general question behind those ones: What *standard* must observational knowledge meet? You are using, it seems, observational evidence; what standard must it meet, if it is to be giving you observational knowledge? (And that sort of question will arise about all evidence and all knowledge. That will become apparent as this article proceeds.)

### **c. Knowing Purely by Thinking**

When philosophers ask about the possibility of some knowledge's being gained purely by thinking — by reflection rather than observation — they are wondering whether a priori knowledge is possible. Historically, those who believe that some such knowledge is possible are called *rationalists* about knowledge. (*Empiricists*, in contrast, believe that all knowledge is observational in its underlying nature, even when it might not seem so. This is the belief that all knowledge is *a posteriori* — present only after some suitably supportive observations are made.) As was done for observational knowledge in section 3.b, this section mentions a few of the multitude of questions that have arisen about *a priori* knowledge — knowledge which would be present, if it ever is, purely by thinking, maybe through an accompanying rational insight.

*How would there be a priori knowledge?* It is difficult, to say the least, for us ever to know that a piece of putative knowledge would not be *at all* observational, so that it would be gained *purely* by thought or reflection. We talk of pure mathematics, for example, and our knowledge of it. Consider the content of the sentence, ' $2 + 2 = 4$ .' It could be applied to physical objects; nonetheless, we might deny that it is at all *about* such objects. But then we must explain how we know that we are using thought alone in knowing that  $2 + 2 = 4$ , rather than knowing this mathematical truth in a way which is simply much less *directly* observational. Would we know it, for instance, partly by knowing how to interpret various physical representations which we would observe — numerals ('2' and '4') and function signs ('+' and '=')? If this is even part of how we know that  $2 + 2 = 4$ , is the knowledge at least not *purely* a result of thought rather than observation?

[On related issues, see Quine's 'Two Dogmas of Empiricism', in Moser 1987, a collection with many readings relevant to this section.]

*Could a priori knowledge be substantive?* It might be thought that pure reflection — and hence *a priori* knowledge — is possible when the truths being known are especially simple, even trivial. 'All bachelors are unmarried' is true, yet trivial: it is uninformative for anyone who understands at all the concept of a bachelor. 'There is more than one infinity' is true yet not trivial: it is informative for some who understand at all the concept of an infinitude. If 'There is more than one infinity' is knowable by thought alone, that would be substantive *a priori* knowledge. But if only truths like 'All bachelors are unmarried' are knowable purely by thinking, maybe there cannot be substantive *a priori* knowledge. So, which is it to be? (If we reply that it depends upon what a particular *a priori* known truth is about, we return to the previous paragraph's question about knowledge gained purely by thinking. Alternatively, if we reply that it depends upon which standard is being met — such as when understanding a specific concept like that of bachelorhood or of infinitude, so as to gain knowledge from it — this takes us to the next paragraph's question.)

[Classically, the issue of whether there can be substantive *a priori* knowledge was posed by Immanuel Kant, in his eighteenth-century *Critique of Pure Reason* (2007 [1781/1787] — as the question of whether there can be synthetic *a priori* knowledge.)

*What standard would a priori knowledge have to satisfy?* If there could be *a priori* knowledge, is it clear what standard it would need to have satisfied? There have long been philosophers for whom part of the appeal in the idea of *a priori* knowledge is the presumption that it would be infallible. That is, it would satisfy a conclusive — in effect, a perfect — evidential standard. It would do this because a capacity for pure thought, undistracted by observed contingencies within this world, would be what has provided the *a priori* knowledge. However, some recent epistemologists (for example, Bonjour 1998) regard that picture as overly optimistic. The one person is both observing and thinking; and if we expect fallibility to be part of how she observes, maybe we should expect fallibility likewise when she is thinking. Is it simply *obvious* that when we are not observing, only thinking, we are more — let alone perfectly — reliable or trustworthy in our views? Or do we also think only imperfectly? Perhaps we need observations as ‘checks’ on what could otherwise become thoughts ‘floating free’ in our minds. Yet maybe, even so, these ‘checks’ remain imperfect. To think without observing might not be to improve dramatically, if at all, the use of one’s mind.

#### **d. Knowing by Thinking-Plus-Observing**

And so again we meet the question of the extent to which, in one way or another, we are vulnerable when trying to gain whatever knowledge we can. Of course, we might claim that we are only vulnerable when focussing just on observation or on reflection — ignoring the other. Surely (it will be suggested), much or even all of our knowledge is a mixture — both observational and reasoned. Is that how we will stride forward as knowers?

Optimism replies, ‘Yes. Possibly there are philosophical limits upon the effectiveness of observation by itself and of reason by itself. Still, to combine them is to overcome those limits, or at least enough of them.’ In response to which, less-than-optimism counsels, ‘Maybe not. If each of observation and reflection has limitations of its own, a combination of them might compound those weaknesses. The result could be a blurring of the two, so that we would never know whether, on a particular occasion, weakness in one — in the observing or in the reflecting — is weakening the whole.’ Which of those alternatives is right? Optimism? Less-than-optimism?

That depends. We should now consider an epistemologically classic doubt about people’s abilities ever to gain knowledge.

#### **4. Sceptical Doubts about Knowing**

From the outset of philosophical thinking about knowledge, doubts have never been far away: do we really know what we think we know? And that question was not meant merely to ask whether sometimes we are mistaken in claiming a particular piece of knowledge. The philosophical concern was more pressing: do we *ever* know what we think we know? Even when lacking all views on whether we know, could we always fail to know? Is knowledge an attainment forever beyond us — all of us, everyone, all of the time?

That question confronts us with a *radical sceptical* possibility. Possibilities that are less radical but still possibly disturbing, and less widely sceptical but still sceptical, have also been discussed. Is there no knowledge of a physical world? Is there no scientific knowledge? Is there no knowledge of moral truths? Is there no knowledge of the future? And so it goes. Let us now examine one of these. It is one of philosophy’s most famous non-radical sceptical arguments — a scepticism about external world knowledge. (It is sceptical, partly because it denies something otherwise accepted by almost everyone: sceptical denials are surprising in that sense.) Here is how it unfolds.

If there is observational knowledge (section 3.b), it is knowledge of what philosophers generally call the external world. By this, they mean to designate the physical world, regarded as something with an existence and nature distinct from (and perhaps, or perhaps not, represented accurately in) any individual’s beliefs as to its existence and nature. Those beliefs could be true because there is a

physical world with a nature matching what the beliefs attribute to it. Equally, however, the beliefs could be false because there is no physical world quite, or even at all, as the beliefs claim it to be. And *if* the beliefs are false, the usual philosophical moral to be drawn would be that they are not knowledge. (Knowledge is only of truths or facts: see section 6.f.)

Still, do we ever have reason to regard all of our beliefs about the physical world as actually false? Perhaps not consciously so, while ever in fact we have the beliefs; for part of having a belief is some sort of acceptance of its content as true, not false. Nevertheless, maybe one can have a belief while accepting that one cannot know quite how one has *gained* that belief. And this is significant because there are ways of having a belief which — even without guaranteeing the belief's being false — would be incompatible with the belief's being knowledge. For instance, even if one feels as though a particular belief has been formed via careful reasoning, perhaps ultimately that belief is present largely because one *wants* it to be. And one might concede this, even if reluctantly, as a possibility about oneself. More generally, therefore, maybe one could have a belief while also accepting one's not quite being able to know that one has not gained it in a way which is wholly unsuitable for its being knowledge.

In theory, there are many possible knowledge-precluding ways of gaining a particular belief. Here are a few generically described ways:

Sometimes, your individual sensing or thinking might be only yours, in the worrying sense that it could be misleading on the particular topic of your belief, more so than other people's sensing or thinking would be on that same topic.

Sometimes, anyone's sensing is only human, in the sense that it could be misleading about aspects of the world which other animals sense more accurately.

Human reasoning is also only ever human in the sense that (as Christopher Cherniak has explained: 1986) even some seemingly simple assessments could be computationally beyond our capacities. There is only so much that any person's brain can do with so much data. Even checking for something as familiar as consistency between many of one's beliefs is an extremely complex task. This is not necessarily because consistency in itself is always complex. It is because there is too *much* checking to do, given the need to evaluate every possible combination from among one's beliefs.

Sceptical arguments could be generated from those and from comparable possibilities. One historically prominent suggestion — philosophers usually attribute its most influential form to Descartes (1911 [1641]), in his 'Meditation I' — directs us to the phenomenon of *dreaming*. Suppose that you feel as though you are sensing, in a normal way, a cat's sitting in front of you. But suppose that this experience is actually present as part of your dreaming, not as part of using your senses in a normal way. There seems to you to be a cat; the circumstance feels normal to you; even so, in fact you are asleep, dreaming. Presumably, therefore, your feeling or experience at this time is not providing you with knowledge right now of the cat's presence.

Now, *could* that be how it is on every occasion of your feeling there to be a cat in front of you? Indeed, we can generalise that question, to this philosophical challenge: *Whenever* you seem to be having a sensory experience about the world around you, can you know that you are not dreaming at that time? And this question is a challenge, not only a question, because it might not be clear how you *could* have that knowledge of not dreaming at that time. Any evidence you mention in support of the contention that you are not dreaming will be the *same* sort of evidence as that which has just been questioned. Imagine thinking to yourself, 'I remember waking up this morning. I feel awake still. I feel *so* awake.' You thereby feel as though you are mentioning some good evidence, reflecting decisive non-dreaming experiences. But your having that feeling could *itself* be present as part of your dreaming; and if it is, then it is not knowledge. So, any such experience on your part of reaching for apparently good evidence, of bringing to mind how awake you feel, will merely be more of the same. That is, it will be just another instance of the same sort of experience as was being questioned in the first place; and it will be no less vulnerable to the possibility of merely being part of a more or less extended moment of dreaming by you. Your citing these further

experiences thus provides no new form of evidence which is somehow above suspicion in this context of questioning the apparently observational evidence (the suspicion, remember, of possibly being an experience produced as part of a dreaming experience).

Then the sceptical conclusion follows swiftly. If you never know that your apparent experiences of the physical world around you are not present as part of your dreaming while asleep, you never know that what feels to you like a normally produced belief about the world is not present as part of an experience which *precludes* that you are thereby having a belief at this time which is knowledge. Accordingly, for all that you do know about yourself at that time, you fail to have knowledge of your surroundings. In that sense, you might not have knowledge of the physical world around you. *Do* your apparent beliefs about the world fail in that way to be knowledge? Indeed so, concludes the sceptical reasoning: if (for all that you do otherwise know about them) they might not be knowledge, then they are not sufficiently well supported by you to actually be knowledge. (...)

Stephen Hetherington

Email: [s.hetherington@unsw.edu.au](mailto:s.hetherington@unsw.edu.au)

University of New South Wales Australia