Epistemology

Course notes by Richard Baron

This document is available at www.rbphilo.com/coursenotes

Contents

Version of March 2011

Introduction to epistemology

The name comes from the Greek word *episteme*, knowledge.

The questions in the contents list indicate the scope of the subject, although they do not cover everything in the subject.

There are links with other areas of philosophy. For example:

ethics: how do we acquire moral knowledge, and what sort of knowledge it is?

philosophy of science: what is special about scientific knowledge?

philosophy of language: how words get meaning affects whether beliefs can be knowledge.

History

Many of the biggest names in philosophy have worked on epistemology, including:

Plato (427-347 BC): discussed what distinguishes knowledge from true belief

René Descartes (1596-1650): formulated an extreme scepticism and gave a way out of it

David Hume (1711-1776): limited our knowledge to what experience and logic tell us

Immanuel Kant (1724-1804): separated what is real to us from things in themselves

Reading

You do not need to do extra reading, but if you would like to do so, there are plenty of textbooks on epistemology. You might for example like to try either one of these two:

Noah Lemos, *An Introduction to the Theory of Knowledge*. Cambridge University Press, 2007.

Duncan Pritchard, *What is This Thing Called Knowledge?* Routledge, second edition, 2009.
(This is not the same as another book by Duncan Pritchard, which is just called *Knowledge*.)

You can also find good articles on epistemological topics in:

The Stanford Encyclopedia of Philosophy, http://plato.stanford.edu

The Internet Encyclopedia of Philosophy, http://www.iep.utm.edu

What is knowledge and why is it valuable?

The traditional analysis

Justified true belief.

Could we do without any of those elements?

Plato, *Theaetetus*

This dialogue was probably written around 370 BC. It considers and rejects three definitions of knowledge, but does not give a definition to accept. The three definitions are:

> perception – many objections, including the relativity of perception, knowledge of the future, our use of concepts not acquired through perception;
>
> true belief – but a lawyer can persuade a jury to a true belief without their having knowledge; also technical problems about false belief;
>
> true belief with an account – but then we need knowledge of the account.

"True belief with an account" comes quite close to "justified true belief".

Gettier cases

Edmund Gettier, "Is Justified True Belief Knowledge?", *Analysis*, volume 23, number 6, June 1963, pages 121-123.

Smith and Jones are in your office. Smith has often mentioned how much he likes his Ford motor car. And you have seen him in it. But unknown to you, he sold it yesterday. Jones owns a Ford, but he has never mentioned that. So you believe that someone in your office owns a Ford, and you have justification for that belief, and the belief is true.

> Response: no false lemmas allowed.

You are walking down a road, and you see a black sheep lying down in a field, clearly and in broad daylight. So you have a justified true belief that there is a black sheep in that field. But in 90% of the fields, there are rocks that look like black sheep, and in 10% there are sheep (one rock or one sheep, and not both, per field).

> The "no false lemmas" response does not seem to work.

Tracking

Robert Nozick (1938-2002, professor at Harvard), *Philosophical Explanations* (1981), chapter 3, section I.

You only know something if you satisfy the tracking condition: if it were false, you would not believe it (in addition to the conditions that it is in fact true and that you in fact believe it).

This seems to deal with both the Smith-Jones-Ford example and the sheep-rocks example.

There is a problem with tracking. If what you believe were false, things would be different. How different? Some differences would lead you to reject the belief, but others would mislead you into accepting it. Do you have to be able to cope with every possible way in which the belief might be false?

I see a table in front of me. In most circumstances in which there was no table there, I would not believe that there was a table there. So I would mostly satisfy the tracking condition. But some very special trick of the light might create the illusion of a table. Then I would still believe that there was a table in front of me. So do I fail the tracking condition? Do I really not know that there is a table in front of me?

Knowledge, certainty and the impossibility of error

Can you know something even though you are not certain of it?

Can you know something even though you might be wrong?

If you know something, must you know that you know it?

Is knowledge better than true belief?

Plato, *Meno* (probably written around 385 BC), 97-98. True opinion is as good as knowledge for directing your actions, but it is unstable. If you have knowledge, you will retain it.

What we make of Plato's view depends on what we think knowledge is. In the *Meno*, he has a theory of knowledge as recollection from a past life elsewhere. And he works with a picture of knowledge by direct experience. Someone who has been to Larissa will know the way there. Someone who has never been there may have a true opinion as to which is the right road, and will then be just as good a guide in practice.

Knowledge has resilience: someone who knows is less likely than someone who happens to have a true belief to abandon the belief when they are given misleading counter-evidence (Miranda Fricker, "The Value of Knowledge and the Test of Time", chapter 7 of Anthony O'Hear (ed.), *Epistemology*, Cambridge University Press, 2009). This makes sense even without a theory of knowledge as recollection. It works if you think that knowledge is true belief plus your having thought through the evidence (and possibly plus something else too).

What kinds of knowledge are there?

Classification by subject matter

Logic and mathematics

The natural sciences: physics, chemistry, biology, some parts of psychology

The human sciences: other parts of psychology, sociology, economics, political science

The humanities: history, literary criticism, philosophy

Moral knowledge (is it knowledge of facts?)

Knowing how to live, common sense

Can we put all or some of these types of knowledge on a scale, with links between different types?

Classification by role in the knower's life

Knowing that (facts), knowing how (skills) and knowing a person (acquaintance)

Basic assumptions about individual people or things: "My attitude towards him is an attitude towards a soul. I am not of the *opinion* that he has a soul" (Wittgenstein, *Philosophical Investigations*, part 2, section 4).

Basic ways in which we approach the world:

> we can change the future but not the past;
> every physical event has a physical cause.

Classification of propositions by their relationship to concepts and to the world

Analytic propositions

These are true by virtue of meaning. Different people take this to amount to different things. It certainly includes "all bachelors are unmarried". Does it include the whole of logic and mathematics?

Do analytic propositions tell us anything about what the world is like, or do they just map out the relationships between concepts?

Synthetic propositions

The propositions that are not analytic. An example is, “human beings need oxygen”. You cannot work that out from the definition of “oxygen”, or from the definition of “human being”.

Quine’s attack on the distinction

Willard Van Orman Quine (1908-2000), wrote “Two Dogmas of Empiricism” (1951).

First dogma: there is a basic distinction between analytic and synthetic. We must abandon this dogma because we would need to understand sameness of meaning, synonymy, in order to identify analytic statements. But we cannot get a proper grip on synonymy. We cannot just say that terms which apply to the same things are synonymous, because “has a heart” and “has kidneys” apply to the same animals. And we cannot rely on dictionaries, because they just report how people happen to use their languages.

Second dogma: reductionism. We can cash out statements about the world in terms of what observations they imply. “All pillar boxes are red” implies that whenever you see a pillar box, you will see it as red. Then statements with the same implications for observations would be synonymous.

But statements about the world do not directly correspond to particular sets of observations. If you saw a pillar box that appeared to be yellow, you could reject the claim that all pillar boxes were red. Or you could put it down to a trick of the light, or an eye disease. Or … or … or … .

You would probably just reject the claim that all pillar boxes were red. But in other cases, you could have more of a choice, at least until you had collected a lot more evidence. For example, when your observations of the rotation of galaxies do not make sense given the observed matter, do you assume that there is dark matter or do you re-work your theory of how galaxies behave?

Quine claims that we should not take our beliefs one by one, but should see them all as woven together in a fabric. Some are near the edge, and make close contact with experience. “All pillar boxes are red” would be like that. When you get contrary data, the easiest and most natural thing to do is to change the belief. Others are further in, like theories of chemistry and physics. When the data come out wrong, we may have a choice of ways to make adjustments to relieve the tension in the fabric. Even supposedly analytic truths, and the laws of logic, right at the centre, might be changed if we felt that doing so was the best way to relieve tensions in the fabric. So nearness to the centre replaces a neat analytic/synthetic split.

What forms of justification are there?

Direct justifications for specific beliefs

"Henry VIII was not very nice": justify by listing executions that he ordered (beliefs justifying beliefs)

"There is a tree in the park": justify by saying that you saw it (methods justifying beliefs)

Justifications for using beliefs to justify beliefs

"I believe that when William Rufus was killed by an arrow in 1100, it was an accident, because the chronicler William of Malmesbury says so. I am justified in using this justification because William of Malmesbury wrote only 20 years or so after the event."

"I believe that there will be economic growth next year because interest rates are currently low. Interest rates affect economic growth because they affect companies' willingness to invest."

Justifications for using methods to justify beliefs

"I checked that there was a tree by the method of looking": justify use of the method by pointing out that it was broad daylight.

"I confirmed that chimpanzees are our cousins with a common ancestor by reading papers on the subject in scholarly journals": justify use of the method by pointing out that papers only get published in scholarly journals if they are written by experts and reviewed by other experts.

Internalism about justification

You need to be aware of justifications for your beliefs, if they are to count as knowledge.

Arguments for internalism

If someone claims to know something but cannot defend his or her belief at all, we are reluctant to count that belief as knowledge.

Being able to produce justifications for a belief is a good sign that it was acquired rationally, not by accident or because someone used propaganda on the believer.

Arguments against internalism

It does not allow non-human animals to know things. (We do not expect animals to be able to defend their beliefs, because they do not claim to know things.)

If you have to be aware of justification J for belief B, and be aware that J is a justification for B, then you have to be aware of some justification K for J, and be aware that K is a justification for J, and so on. (You need K to justify J, because if you had no justification for J, you could not think that it was a justification for B. You could only think that if J were true, it would justify B.)

Externalism about justification

You do not need to be aware of justifications for your beliefs, in order for them to count as knowledge.

Arguments for externalism

It fits well with a definition of knowledge that is based on tracking.

We do not need to develop a notion of justification.

It captures what matters in practice: someone who has knowledge should get the right results, for example by directing us correctly to Larissa.

It accommodates the knowledge that non-human animals have.

Arguments against externalism

It does not give much of a role to your awareness of evidence for your beliefs, or to your rational thought about the evidence.

If you are a reliable clairvoyant, so that your guesses are always right but they still feel like guesses to you, externalists will say that you have knowledge. Is that acceptable?

Can we defeat Agrippan scepticism?

Agrippa's Trilemma

The trilemma is named after Agrippa the Sceptic (first century AD), but we do not have his own works. We have information from *Outlines of Pyrrhonism* by Sextus Empiricus (c.160-210 AD) and from the life of Pyrrho (c.360-c.270 BC) in *Lives and Opinions of Eminent Philosophers* by Diogenes Laertius (third century AD). Hans Albert (1921-, former professor in Mannheim) brought it into current philosophical debate and called it the Münchhausen Trilemma.

I have a belief. I have justifications for that belief (level 1 justifications). I may also have justifications for believing the level 1 justifications (level 2 justifications). And so on. The trilemma is that there are three options:

1. Infinite chains of justifications

   We cannot complete the chains, so we have no idea whether our beliefs have solid or shaky support; and it would only take one shaky stage in a chain to make the whole chain shaky.

2. Basic justifications that can bring chains to an end

   What sort of thing could bring a chain to an end? The evidence of our senses would be the obvious candidate for beliefs about the physical world, but we can be misled by our senses. When we get on to mathematics, we just have to accept axioms. Should we accept them? And mathematical physics, which looks like our highest-quality knowledge about the physical world, seems to depend on our acceptance both of sensory evidence and of axioms. Do we really want to take such important foundations of our knowledge on trust?

3. Circles in chains of justifications

   We don't normally think much of circular arguments. A circular argument might offer some support for a belief, if the circle was large so that it involved lots of different thoughts and observations. And it would help if there were lots of criss-cross linkages, creating a complex web of beliefs. But we could not be sure that we did not have one big internally consistent web of beliefs that did not reflect reality, one big complex dream.

The trilemma is only a decisive argument against justification that guarantees truth

If we think that we can have knowledge without having justification that is strong enough to guarantee truth, then the trilemma will not prevent us from having knowledge. We can have limited justification for our beliefs.

Critical rationalists like Karl Popper (1902-1994, professor at the London School of Economics) would go further. For them, we do not get positive justification for a belief. We can collect data and perform experiments, and if a belief survives a lot of tests, that is comforting. But we might at any time get new information that would force us to drop the belief. Reason is the tool of criticism, not of justification. We can have knowledge, but without the justification component.

Structures of justification

Foundationalism

Work down to basic beliefs, which do not need justification. This tends to go with a correspondence theory of truth: statements are true if they correspond to reality.

But which basic beliefs do not need justification? The evidence of our senses can mislead, so perhaps we need justification for accepting a given piece of sensory evidence. Axioms just have to be accepted, and it would seem better if they were justified.

Can we justify our acceptance of basic beliefs by the results that we get? Our senses allow us to cope with the world, and our mathematical axioms are amazingly productive. But do we need to justify our belief that we are successful, and our belief that success is good evidence of correctness?

Coherentism

The coherence theory of justification: beliefs are justified by cohering with our existing set of beliefs. There is also the coherence theory of truth: beliefs are true if they are part of a coherent set.

What is the nature of the coherence that justifies? Logical consistency would be the minimum, but that would be too undemanding, allowing too much to be justified. Logical implication would be too demanding, allowing very little to be justified. Is explaining many things by using a few basic facts a mark of coherence?

What about the danger of having a free-floating coherent set of beliefs that we take to be justified, but that has no contact with reality?

> Response 1: we can expect that there will be enough contact with reality to keep us on track. Observations and the results of experiments form part of our set of beliefs. If they conflict with existing beliefs, the results is incoherence. That then needs to be removed by changing our beliefs.
>
> Response 2: there is no definite form of reality beyond our system of beliefs. This looks like a rather extreme response, but anti-realism is a respectable philosophical position.

Foundherentism

Susan Haack (1945-), professor at the University of Miami: *Evidence and Inquiry* (1993), chapter 4.

Justification is a compound of foundationalism and coherentism. It is like a crossword puzzle. Each answer must be founded on its clue, but it must also cohere with other answers by having the same letters at the places where they cross.

Horses for courses?

Perhaps we should say that the form of justification depends on the nature of the subject.

We have a scale: logic and mathematics – the natural sciences – the human sciences – the humanities. At the early points, axioms and very straightforward objective observations predominate. As we move up the scale, these become less important, and coherence and circular arguments come to predominate.

At the mathematics and physics end, our acceptance of axioms and of basic observations is justified by our immense success. We construct tight logical structures that explain a huge amount on the basis of a very modest set of assumptions.

Objection: there could be a completely different mathematics and physics that would be just as successful.

At the humanities end, our use of circular arguments is justified by the fact that the circles are large, bringing many considerations to bear, and that the results really do deepen our understanding of ourselves and of the world.

Objection: if we accept circular arguments as justifications, we may find that we have equally good justifications for two contradictory positions.

The two shifts as we move along the scale

First shift: the acceptance of axioms, and their use in constructing tight logical structures, become less significant, and large circles in argument become more significant.

Second shift: in the sources of justification for the use of our methods of argument, whether axiomatic, circular or mixed. We shift from the measurable empirical success of comprehensive theories, through the measurable empirical success of particular results, to the deepening of understanding.

These two shifts bear a complex relationship to the balance between foundations and coherence. As we move up the scale of disciplines, we do not simply move from the overwhelming importance of foundations in justification to the overwhelming importance of coherence.

> Axiomatic structures, such as those that we have in mathematics and in fundamental physics, are a key source of coherence, and the success of theories in making sense of individual situations, a foundational type of justification, is important even in the least axiomatised disciplines.
>
> Well-established fundamental principles of a natural science can provide support for particular beliefs within that science in a foundational way, and beliefs in the humanities can be tested by asking whether the overall picture that they give us makes sense, a coherentist criterion.

What ways of knowing are there?

Propositions that are knowable a priori and a posteriori

A priori: propositions that could not be contradicted by experience, so you do not need to go and check the facts of the world. You know a priori that bears are mammals, because "bear" is defined as "mammal of the family ursidae". There is no need to go and look at actual bears. If you meet a large furry animal that is not a mammal, this does not show that some bears are not mammals. It shows that the animal is not a bear.

A posteriori: propositions that could be contradicted by experience, so you do need to go and check the facts. You only know a posteriori that bears like honey. If you found lots of bears who ignored honey that was put in front of them, or who spat it out, and did not find bears who gobbled it up, that would show that bears did not, on the whole, like honey.

	Analytic	Synthetic
A priori	Yes	?
A posteriori	No	Yes

Rationalism

René Descartes (1596-1650), wrote *Discourse on the Method* (1637) and *Meditations* (1641).

Descartes got quite a lot by pure reason:

> I exist. This is obvious because I think (the *Cogito*).
>
> Mind-body dualism. Descartes can imagine himself without a body, but not without a mind. He has a clear and distinct idea of the mind as a thinking and non-extended thing, and of the body as a non-thinking and extended thing.

But we have to worry about the Cartesian circle. He can rely on his clear and distinct ideas because there is a benevolent God. At the same time, he cannot doubt that there is a God because he has a clear and distinct idea of an infinite and perfect being, which cannot have come from within himself, a finite and imperfect being.

And we still need observation to get the details of science.

Empiricism

David Hume (1711-1776), wrote *A Treatise of Human Nature* (1739-40) and *An Enquiry Concerning Human Understanding* (1748).

Reason will only get you relations of ideas (like “all bachelors are unmarried”, or “3 x 5 = 15”). You need experience to get matters of fact. And these are the only two sorts of knowledge.

> We cannot get to matters of fact simply by reasoning, because the contrary of any matter of fact is conceivable. We can easily imagine that the sun will not rise again.
>
> We cannot work out the causes and effects of things by reason alone. We can always imagine some different effect following a cause. When a billiard ball strikes another, the second one moves, but we can imagine it just staying where it is.
>
> Even our idea of causal necessity arises simply from seeing the same causes followed by the same effects over and over again.

“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” (*An Enquiry Concerning Human Understanding*, section 12, last paragraph).

Kant

Immanuel Kant (1724-1804), wrote *Critique of Pure Reason* (first edition 1781, second 1787).

He gave us the analytic/synthetic distinction. For him, a proposition was analytic if the predicate was contained in the subject: all bachelors are unmarried (subject: bachelor, predicate: unmarried), all bodies are extended (= take up space). Otherwise it was synthetic: all bachelors are badly fed, all bodies are heavy.

Synthetic propositions that we know a priori:

> A straight line between two points is the shortest line between them
> In all changes of the material world, the quantity of matter remains unchanged

How can we know synthetic propositions a priori?

> There are things as they appear to us (the phenomena) and things as they are in themselves (the noumena). These need not be two sets of things. They could be the same things considered in two different ways.
>
> We can only have knowledge of the phenomena. We cannot know about the noumena.
>
> To grasp the phenomena, we have to arrange them in space and time. We also have to use categories like reality and causation, in order to think about the phenomena.
>
> But then the phenomena have to comply with the properties that our forms of space and time, and our categories, imply. Our geometry, and our notions of things, properties and causes, have to apply because that is how we see the world.

Mill on mathematics

John Stuart Mill (1806-1873), wrote *A System of Logic* (1843), *On Liberty* (1859) and *Utilitarianism* (1861 as articles, 1863 as a book).

We get our notions of geometrical objects, such as triangles and circles, from real objects. We omit the uninteresting details, so that we just have the shapes. We then formalise our geometry using axioms, like the axiom that if two lines cross, they cannot both be parallel to a third line. But the axioms are generalisations from real objects. It is no use claiming that they have to be true because we cannot conceive of their being false. There are plenty of things we used to be unable to conceive, but now we can. People used to be unable to conceive that if a stone was thrown and no force at all acted on it, it would go on at the same speed for ever. (*A System of Logic*, book 2, chapter 5)

“2 + 1 = 3” is likewise a generalisation from experience. But propositions of arithmetic apply to absolutely everything. (*A System of Logic*, book 2, chapter 6)

Mill wanted to ensure that propositions of mathematics were real, and not merely verbal. But did he put their certainty at risk?

Can we defeat Cartesian scepticism?

Scepticism: we don't know anything, or not nearly as much as non-philosophers think we know.

Dreams, demons and brains

Descartes, *Meditation* 1: I might be dreaming that I am sitting by the fire with a piece of paper in my hand. Worse, there might be a demon who systematically deceives me about everything.

(Descartes was not himself a sceptic. He got out of the difficulty, as mentioned above, but in a way that we might not think satisfactory.)

The brain in a vat: your brain might be sitting in a vat of nutrients, wired up so that it seems to you as if you have a body and live in the world in the ordinary way. How could you tell that anything was wrong?

First route to scepticism: any belief might be mistaken

We cannot be absolutely certain of anything.

Response: so what? Knowledge does not require certainty.

Response: we can have a lot of confidence in some of our beliefs, even without certainty.

Second route to scepticism: the closure argument

I only know something if I can also know anything that is implied by it and that is not too difficult to work out. (Example: I only know that there are 50 states in the USA if I can also know that there are more than 13 states in the USA.)

1. If I am at my desk, then I am not a brain in a vat.
2. I only know that I am at my desk if I can know that I am not a brain in a vat.
3. I cannot know that I am not a brain in a vat.
4. Therefore, I do not know that I am at my desk.

Response: epistemically relevant worlds

We do not need to consider crazy possibilities, like the possibility that we are brains in vats. We need only consider epistemically relevant worlds. These are worlds that are near enough to the world as we think it actually is (Keith DeRose, "Solving the Skeptical Problem", *Philosophical Review*, volume 104, number 1, January 1995, pages 1-52.sections 11 and 12).

I only know that I am at my desk if I can know that I have not gone to the cinema. (I might well have gone to the cinema: the world in which I am, at this time, at the cinema, is epistemically relevant.) But I can know that I have not gone to the cinema. So I know that I am at my desk. I do not need to be able to know that I am not a brain in a vat, because that is not something that would be at all likely, given how we think the world actually is.

Semantic externalism

Hilary Putnam (1926-, was a professor at Harvard)

Meanings just ain't in the head

What your words mean depends on how you came to use them, and in particular on what causes the sensations that lead you to use the words. Thus the meaning of the word "tomato" is a certain kind of fruit, because it is usually that kind of fruit which causes the red-blob-with-a-green-bit-on-top sensations that lead you to use the word.

"Tomato" would have the same meaning if you occasionally hallucinated tomatoes. You can sometimes be caused to use the word by something else that triggers the appropriate disturbance in your brain cells, rather than by a fruit.

But if the only thing that ever caused you to use the word was some such trigger of disturbances in your brain cells, there would be no causal connection between the fruit and your use of the word.

Even then, you could say that the meaning of the word was a kind of fruit, because the meaning would be fixed by how most people used the word. They would still see real pieces of fruit.

All in the vat together

Now suppose that everybody only ever hallucinated tomatoes. Nobody ever saw a real piece of fruit (or at least, no-one had seen one for years). Then the word would not mean a kind of fruit, but whatever else it was that triggered appropriate disturbances in people's brain cells.

That thing could be a feature of the computer program that fed us all stimuli, if we were all brains in a vat. (Note that we all need to be in there, so no-one has a causal sensory relationship to pieces of fruit. The vat and the computer must run automatically.)

So what words mean depends on environment. If we are all brains in vats, then our words mean something different from what they mean if we are not brains in vats.

So if I am a brain in a vat (along with everyone else), and I say "I see a tomato", I am not mistaken, even though there is no fruit there, because "tomato" does not mean a piece of fruit of a certain kind. It means the feature of the computer program that stimulates some of my brain cells in a certain way. And "see" means "am being stimulated by a feature of the program". I say "I see a tomato" when that feature is active, so I speak the truth.

But that's not Putnam's main argument. He wants to show that …

I am not a brain in a vat

Either I am a brain in a vat, or I am not.

Suppose that I am a brain in a vat.

Then when I use the word “brain”, it does not mean a 1.5 kg chunk of grey matter. It means some feature of the computer program which stimulates my thoughts about mechanisms of human thought. And when I use the word “vat”, it does not mean a tank. It means some feature of the computer program which stimulates visions that are just like the ones that people outside vats would have if they stood in front of tanks.

So if I am a brain in a vat and I say “I am a brain in a vat”, I speak falsely, because I am not a bit of a computer program floating in another bit of a computer program.

And if I am not a brain in a vat and I say “I am a brain in a vat”, then I speak falsely, because I am not a brain in a vat.

So either way, all my utterances of “I am a brain in a vat” are false.

So all my utterances of “I am not a brain in a vat” are true.

So I am not a brain in a vat.

The problem with the last step

Suppose that I am not a brain in a vat. Then there is no problem with the last step. The words in “I am not a brain in a vat” have their ordinary meanings. So if the utterance is true, then I am not a brain in a vat.

But suppose that I am a brain in a vat. Then the words in “I am not a brain in a vat” do not have their ordinary meanings. They mean that I am not a bit of a computer program floating in another bit of a computer program. That could be true, even if I was a brain in a vat.

The limitation problem

Putnam’s set-up assumes that we are all brains in a vat, and always have been (or at least, have been for some time).

If there are lots of people outside the vat, then they can fix the meanings of words in the ordinary way. “Tomato” means a certain kind of fruit, “brain” means a 1.5 kg chunk of grey matter, “vat” means a tank, and so on.

If we were outside the vat until recently, our words would still hang on to their ordinary meanings.

So even if his argument works, it does not deal with all brain-in-a-vat worries. But it would still be pretty impressive to deal with some of them.

Is virtue epistemology the way forward?

Virtue reliabilists: epistemic virtues are stable faculties, like vision or the ability to reason logically, which tend to lead to the formation of true beliefs and the rejection of false beliefs. References to virtues can play a role in defining knowledge and justification.

> Big names: Ernest Sosa (Rutgers University), John Greco (St Louis University)

Virtue responsibilists: epistemic virtues are character traits, like fair-mindedness and a desire to search out all of the evidence. References to these virtues can be used in the definition of knowledge and of justification, but virtue responsibilists also study the social aspects of the pursuit of knowledge, and ideas of our responsibilities as people who pursue knowledge.

> Big names: Linda Zagzebski (University of Oklahoma), Lorraine Code (York University, Toronto), James Montmarquet (Tennessee State University)

Differences between virtue reliabilists and virtue responsibilists

The boundary is hazy, and there is overlap.

Possessing the virtues that interest the responsibilists might be said to make you a better person in a moral or quasi-moral sense. The virtues that interest the reliabilists seem to have less connection with morality.

If responsibilist virtues started to lead to false beliefs (perhaps because of an evil demon), we would not suddenly regard someone's possession of those qualities as a sign that he or she was wicked.

Virtue reliabilists can account for the passive acquisition of knowledge, as when you just see that it is raining. You do not need to exercise virtues of the responsibilist type in such cases. But how much knowledge can be acquired passively?

Defining knowledge and dealing with Gettier cases

Knowledge is true belief that results from an exercise of epistemic virtues.

> Which virtues? Vision and clear reasoning, or fair-mindedness and the like, or both?
>
> Does it mean their exercise in circumstances in which they reliably lead to true beliefs? If it does, then we have to define those circumstances.
>
> That definition of circumstances will probably involve saying that the belief must be formed through the exercise of the virtues, with luck not playing a crucial role. But are we then just re-stating the tracking condition?

Transcending the debate between foundationalism and coherentism

Justification for a belief comes from its having its source in an epistemic virtue (of the reliabilist sort).

This might answer the Agrippan sceptic. In order to have knowledge, we do not need to be conscious of grounds for thinking that our beliefs are true. We merely need to have acquired the beliefs in ways that are in fact likely to lead to true beliefs.

But this can at best yield animal knowledge, not reflective knowledge (knowledge that you can consciously defend).

And it assumes externalism, so it transcends the debate between foundationalism and coherentism because that is an internalist debate.

Social aspects of knowledge

We can move away from a focus on the definition of knowledge and the conditions for justification, and examine the conditions for being an intellectually flourishing individual in an intellectually flourishing society. These could include:

> accepting that we can have knowledge of the world;
>
> using honest argument, making all of our evidence available, not using tricks and admitting our mistakes;
>
> conducting debates properly, with openness and respect.

We can also be concerned with epistemic injustice, as described by Miranda Fricker (Birkbeck College). This can include:

> testimonial injustice, where someone has something useful to say but other people do not give his or her views due weight because of his or her race, sex, age, accent or whatever;
>
> hermeneutical injustice, where someone's social experience cannot be understood because prejudice means that the concepts needed to express it are not in use in the society. There might, for example, be no concept of disempowerment in use.

Epistemic virtue and moral responsibility

You did something with bad consequences because you did not know what would happen, or because your beliefs implied that the consequences would be good. Can you escape blame if you did not exercise virtues such as taking care to seek out relevant evidence and considering it in an open-minded way?