

General Philosophy

Dr Peter Millican, Hertford College

Lecture 5: Knowledge and Scepticism



From Scepticism to Knowledge

- Sceptical arguments, such as those of Descartes, suggest that we know very little. But we still want to distinguish between things that we consider we have a right to believe (e.g. on the basis of experience or strong testimony), and other less secure beliefs (e.g. “superstitions”).
- If the sceptical arguments can't be answered, then it's tempting to attack the problem by (re-?) defining “knowledge”.

What is Knowledge?

- “What is X?” questions:
 - X might be “truth”, “perception”, “reason”, “the mind”, “personal identity”, “freedom”, etc.
 - Seen as important in Philosophy since Plato.
- But they are puzzling. Are we asking:
 - “When do we apply the word ‘X’?” or
 - “What is a *genuine* case of X?”
- The former seems merely linguistic; the latter – if different – can appear senseless.

What is Geography?

- “Geography” as a discipline:
 - Initially, perhaps, described the study of places in terms of location, physical characteristics, mineral resources, natural flora and fauna etc.
 - Then extended to cover land-use, farming, and other economic factors, even culture ...
 - Suppose one were now to ask “But is culture *really* part of the discipline of geography?”
 - Well, if “geography” as actually used does cover the study of culture, the answer is “Yes!”

The Concept of Knowledge

- Core *normative* concept, *versus* particular judgements:
 - The concept of “knowledge” plays a central role in distinguishing *reliable* beliefs from others.
 - This makes it *normative*: calling something “knowledge” does more than just categorising it as something we standardly *call* knowledge.
 - Hence it does seem to be possible to ask “Everyone *calls* this knowledge, but is it *really*?”
 - Compare the response to Strawson on induction: we *call* it reasonable, but is it *really* good evidence?

Intuitions, Puzzle Cases, and Conceptual Analysis

- Conceptual analysis can involve:
 - Appeal to linguistic “intuitions” (i.e. judgements that we are naturally inclined to make).
 - Puzzle cases (“intuition pumps”) that can put pressure on those intuitions.
 - Argument, in which we draw out implications of these plausible judgements and principles.
 - Systematisation, in which we try to clarify the concept coherently in the light of all this.

Three Kinds of Knowledge

■ Acquaintance

- “I know Oxford”, “Do you know John Smith?”.

■ Knowing How

- “I know how to drive”, “Do you know how to open this?”

■ *Knowing That, or Propositional Knowledge*

- “I know *that* this building is the Exam Schools”, “Do you know *that* it will rain?”
- Where P is the proposition concerned, this is often referred to as “Knowledge that P ”.

The Traditional Analysis of Knowledge that P

- A subject (i.e. a person) S knows that P if, and only if:
 - P is true
 - S believes that P
 - S is justified in believing that P
- A.J. Ayer gives the last two conditions as:
 - S is sure that P is true
 - S has the right to be sure that P is true

P is true

- If S knows that P , does it follow that P must be *true*? Distinguish two claims:
 - S knows that $P \rightarrow P$ is necessarily true
false: I know that I exist, but it doesn't follow that I exist necessarily.
 - Necessarily (S knows that $P \rightarrow P$ is true)
convincing: We wouldn't allow S 's belief that P to be counted as a case of *knowledge* unless the belief is, in fact, *true*. So it is a necessary truth that anything known is true.

Complications?

■ Knowing Falsehoods?

- “I know that France is hexagonal”

In a sense this can be considered true, because France is *roughly* hexagonal, but in that same sense, it is also true that France is hexagonal.

■ An Abomination

- *Never* confuse “ P is true” with “ P is *believed* to be true”. *Don't* say “ P is true for me, but P is false for him” when what you mean is simply “I believe P , but he does not”. It was never true than the Sun orbits the Earth, even when everyone thought so!

S believes that P

■ If S knows that P , does it follow that S believes that P ? Not so clear:

– Reliable guessing

Suppose that I am not aware of knowing anything about some topic, but my “guesses” in a quiz are always accurate. I might be reported as *knowing* P , even though I don’t believe P .

– Blindsight

Someone with blindsight has no conscious visual awareness, but can “guess” fairly reliably when asked to point towards objects.

Knowing that One Knows

- Suppose that knowledge must always be “conscious”. Then if I know that P , will it follow that I must know that *I know that P* ?
 - The principle is tempting, but we can iterate ...
 - I know that P
 - I know that *I know that P*
 - I know that *I know that I know that P*
 - I know that *I know that I know that I know that P ...*
 - It is clearly impossible to have conscious belief in all of this infinite sequence.

S is justified in believing that P

- Perhaps the central role of the concept of knowledge is to distinguish between beliefs that are “secure” and those that aren’t.
- So what makes the difference between:
 - *believing* that P (where P happens to be true)
 - *knowing* that P ?
- “Surely”, if a belief that P is to count as a case of *knowledge*, it must be a *justified* belief: one must have *the right* to believe it.

The Regress of Justification

- Suppose that I believe that P , and this belief is to be justified. Its justification will typically involve other beliefs. But then if P is to be justified, these other beliefs must be justified too, and so on ... ?
- How to prevent an infinite regress? We could take the whole web of interlocking beliefs as mutually justifying in some way (*coherentism*), or else some beliefs must be justified in a way that does not depend on any other belief. Descartes was a *foundationalist*, taking some beliefs to be totally secure. A more modern approach is *externalism*.

Internalism and Externalism

- An *internalist* account of justification requires all relevant factors to be *cognitively accessible* to *S*. We'll see that this faces difficulties ...
- An *externalist* account (e.g. Armstrong, Goldman) allows that some factors relevant to judging *S*'s justification (for belief that *P*) can be *inaccessible* to *S*; or *external* to *S*'s cognitive perspective.
- So justification could be a matter of a *reliable* causal link between facts and beliefs. I might know that *P* (because my belief reliably depends on *P*'s truth) without knowing how I know.

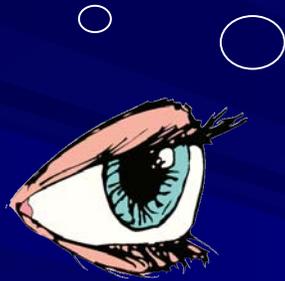
Gettier Cases

- Suppose that:
 - S is justified in believing that P .
 - P clearly implies Q .

Does it follow that S , after inferring Q from P , is justified in believing that Q ?

- On internalist interpretations of “justified”, this does seem to follow. But it leads to so-called “Gettier counterexamples” to the traditional analysis of knowledge.

A Gettier-style Counterexample



*“There’s an oasis
over there”*



Mirage



*Real Oasis
(out of sight)*

- S’s belief is true, and apparently justified, since he infers it from the (apparently justified) belief that he can see an oasis. But we would not say he *knew* that there’s an oasis there.

“No Dependence on False Beliefs”

- Should we add a fourth condition? For example, S knows that P if, and only if:
 - P is true
 - S believes that P
 - S is justified in believing that P ...
 - ... in a way that doesn't depend on any falsehood
- But this seems too strong. If you tell me “there were exactly 78 people there”, but you slightly miscounted (in fact there were 77), can't I *know* that there were more than 40 people there, even though I've inferred this from a falsehood?

The Lottery Paradox

- Another approach would be to understand justification as involving *very high probability* of truth (given the evidence available to *S*).
- But then consider a billion-ticket lottery:
 - I believe that ticket 000000000 won't win
 - I believe that ticket 000000001 won't win
 - ...
 - I believe that ticket 999999999 won't win
- Each of these is extremely probable, but we're reluctant to call any of them "knowledge". So it seems that no probability threshold will do.

Non-Accidental Truth

- To deal with the lottery paradox, it's plausible to count a belief as *knowledge* only if it's not an *accident* – not a “mere” matter of chance (of whatever numerical degree) – that it's true.
- But how do we pin this down?
 - Is it mere “chance” that my corroding speedometer is still sufficiently reliable to provide an accurate reading (when perhaps in a month's time it won't be)?
 - Suppose I very occasionally hallucinate that P , is it “chance” that my current perceptual belief that P is not an hallucination?

Contextualism

- Yet another problem, especially pressing for an “internalist” account of knowledge, is that sometimes our criteria can vary.
 - “I know that the train leaves at 17:36” (because I always take that train).
 - “But do you really *know* that it does? It really is essential that I make that appointment.”
 - “OK, I’ll check on the Web to make sure. Then I’ll know.”
- This suggests that the “hurdle” for what counts as adequate justification can vary.

The Role(s) of the Concept of Knowledge

- Consider the contrast between:

“Does she know that her husband is cheating on her?”

which could just mean “Does she believe that he’s cheating on her, as we all do?”

“Do you know that her husband is cheating on her?”

which is more likely to mean “Is it genuinely the case?”, rather than an epistemological enquiry.

Is “Knowledge” a Genuine Category?

- It is very unusual, in ordinary life, to ask “Does S know that P ” in a situation where:
 - We are totally confident that S believes that P ;
and
 - We are totally confident that P is true.
- This might suggest that it’s a mistake to search for some single consistent account of what “knowledge” is, which can deal with all the contexts in which it is applied.
- But we can still ask whether P is true ...

Back to G.E. Moore's Hands

- If we agree with Moore, then we may see externalism about knowledge and justification as a way of reconciling his claim that we know this is a hand, with the sceptical arguments that seem to show that we can't know that we know.
- An externalist can say to the sceptic:
“I can't prove to you that I know this is a hand, or that my belief is justified, but nevertheless I claim that I do know it, and it is justified.”

Externalism and Scepticism

- Suppose we accept an externalist account of justification. So *if*, say, my perceptual beliefs are, *in fact*, caused by a reliable causal process, *then* I do in fact know that this table is in front of me.
- But of course the sceptic can still ask: “How do I *know* – or if you prefer, *what right do I have to be at all confident* – that my beliefs are in fact so caused?” Externalism does not exclude sceptical doubt “from the inside”.

Putnam's Semantic Externalism

- The sceptic claims “I might be a brain in a vat (BIV), so this hand might be just part of the image created artificially.”
- But what do I mean by “hand”? According to Putnam, meanings aren't purely mental.
- If I am a BIV, then my word “hand” actually means a “hand-in-the-image” ...
- ... in which case this is genuinely a “hand”, because it is a hand-in-the-image.

Sceptical Responses (1)

- Is the meaning of “hand” just determined by what we’re actually referring to when we think we’re pointing to a real hand?
- Or do we have some further idea of the kind of thing that a hand really is?
- Can we thus make sense of the possibility of a “God’s eye view” (unavailable to us), from which it would be clear that it is all a clever simulation, rather than involving a real entity something like what we take a hand to be?

(2) Post-Linguistic Envatting

- Suppose that I am “envatted” after I have become linguistically competent.
- So then my word “hand” has already established its “outside vat” meaning.
- It seems to follow that when I later say “this is a hand” from within the vat, I can manage to mean a real hand rather than a mere “hand-in-the-image”. If so, I can raise the question as to whether this really is a hand.

Back to Induction

- With vertical scepticism (evil demon, BIV, *The Matrix* etc.), it's tempting to ask in a semantic externalist spirit: "Why should I care if it's all an illusion? I'm quite happy to continue with 'life as I experience it' either way."
- But Hume's "problem of induction", as a form of horizontal scepticism, evades this response: whether the world I experience is real or not, I still have the problem of inferring from past to future, from "observed" to "not yet observed".

The Ethics of Belief

- Hume avoids indiscriminate scepticism by rejecting Descartes' "ethics of belief" – the view that we should withhold assent to anything that's not known with total certainty.
- Hume sees belief as typically involuntary, so withholding assent isn't even an option.
- Note that epistemological externalism also involves a similar rejection.
- We seem to be forced to accept this, if we are to hold out against the sceptic.