

## Moving the Bar

This short talk will not be a traditional 'paper', but a personal account of where I have got to with epistemology. Since it covers a lot of ground, I am hoping that most people will find at least some aspects of it interesting. I started out in epistemology, like most people, with Descartes' *First Meditation*, and the agenda that knowledge should aspire to a state of certainty. Like most people, I then became a bit philosophically dispirited when I decided that Descartes had failed to deliver the goods. If you want certainty, then the possibility that I am producing reality out of my own mind (as a dream), or that some other force is creating a bogus reality (as an Evil Demon, or as a virtual reality machine), just seems undeniable. Having total confidence that my own thinking exists (as in the Cogito argument) is no great help, since the thinking itself may still be tainted. You can embrace this result, and assert that knowledge is unattainable, as certain ancient Greeks, and Peter Unger in modern times, have done, or you can drift (as I did) towards a position known as 'fallibilism'. This derives from the pragmatist Charles Sanders Peirce, and says that you can have knowledge even in the face of what is called a possible 'defeater' – that is, some evidence which could undermine your knowledge. Thus we might say that I 'know' that I am in Rewley House, while conceding that it is faintly possible that Rewley House doesn't exist, and that some mysterious visitor might suddenly reveal this alarming truth to us. At this point, though, it is tempting to think that the meaning of the word 'know' has shifted a bit from its usage in discussions of Descartes. The metaphor which comes to mind is that of raising or lowering the bar of justification. In effect, Descartes may have placed the bar so high that it seems impossible to clear it.

Once we begin to talk of moving the bar up, it is worth noting that the bar could be pushed considerably higher than the place where Descartes set it. For example, Descartes rests his arguments on claims that certain truths are 'clear and distinct' to him, which he understands as the view that his own reason, when applied with suitable sceptical caution, is fully reliable. But clearly we can doubt this, since the world seems to be well-stocked with people who believe all sorts of nonsense which they take to be clear and distinct to themselves. Furthermore, Descartes failed to recognise sceptical questions about the Ego, which is supposed to be a thing clear and distinct to the introspective reasoner. There seem to be no real grounds, beyond intuition, for saying that the existence of thoughts requires some further entity called a 'thinker', and if the thoughts must have a thinker, then each separate thought will have a separate thinker, unless something unites them. Descartes spotted the latter problem, but divine intervention was the only proposal he could think of to solve it. Yet another scepticism Descartes might have embraced concerns language. Only in the twentieth century have those sorts of doubts been explored, with proposals that there is no such thing as meaning, and that translation is impossible, or that we can never accurately interpret one another. Thus Descartes, being already in the grips of scepticism about his sense experiences, could have added further scepticisms, about reason, about personal identity, and about meaning. I think we can safely say that he would then know nothing whatsoever.

So we embrace Fallibilism, and lower the bar. But how far are we allowed to lower it? It would clearly be outrageous if each of us were to set the bar at the level which suited us, and which turned our pet prejudices into knowledge. Hence bar-setting cannot be a private activity. Modern commentators have increasingly noticed that Descartes' approach to the problem of knowledge is highly individualistic, founded in his famous solitary and almost solipsistic winter meditation inside a stove. It is, I think, in keeping with the birth of modern science that he wanted to push his own private bar up to a higher level than was to be found in scholastic philosophy or alchemy. But he seems to have no objections within his scheme of thought to some epistemological wrecker who says we should lower the bar rather than raise it.

We now find ourselves faced with a bar height which arises from a consensus within a community. But which community? If we ask which group of people sets the standards for a good justification within our own community, we realise that there is no settled answer. Anyone who was accepted to be in possession of very high-grade justifications in some area of study would be labelled an 'expert'. So consider who might be ranked as an expert, on (say) nuclear physics, or rose growing, or epistemology. When it comes to philosophy I suspect that many people in this room are familiar with the experience of being considered a major expert in one social group, and finding themselves to be a complete non-entity in another. The status of expert is a relative one. Hence the question of whether a person should be considered justified about a proposition will also be relative. If justification is the main hallmark of knowledge (given an initial true belief), then we are suddenly faced with the proposal that knowledge is thoroughly relative.

## Context or Invariance

If we accept this line of reasoning, have we just consigned knowledge to an intellectual dustbin? The proposal that justification and knowledge rest on facts about the circumstances in which they arise has been labelled Contextualism, and has given rise to a lively debate in the past decade. The thesis of

contextualism is expressed as the claim that 'S knows that p' has truth-conditions which vary according to context. Thus 'George knows that Hume was a sceptical empiricist' might be true at the local pub, but not true at a conference on David Hume, the reason being that the word 'know' has changed its meaning between the two locations. The phenomenon of contextualism also seems to occur in much ordinary language, as well as when discussing the word 'know', the popular example being the word 'flat'. The truth-conditions for the sentence 'this surface is flat' seem to vary according to circumstances; it might be flat if you want to put a vase on it, but not flat if you want it to support your writing on thin paper.

In this light, contextualism about knowledge looks quite plausible, and certainly picks out a genuine phenomenon when it labels a person as an expert in modest company, but a novice in a more high-powered company. The standard reply to the contextualist interpretation, though, is known as Invariantism. This approach says there is no variation in the truth-conditions for 'S knows that p', and there is said to be a genuine debate about the precise and objective nature of knowledge, or about the precise height which is correct for the bar of justification if it is to trigger a state of knowledge. The relativity of being acknowledged as an 'expert' is not denied, but is interpreted differently. The other day on the train I heard a child ask its mother where Borneo was, and the child was well satisfied when mum told him confidently that it was in South America. It is probably best not to enquire into mum's justification for this assertion, but even if it was quite good (for example, the fact was told to her by a drunken geography teacher) it would be a bit silly to say that the truth conditions for 'mum knows where Borneo is' are just different in that family. Wise epistemologists will say that mum was wrong, and that the truth conditions for 'S knows that Borneo is in South America' are invariant, and were misunderstood in this case. The sentence 'Hume was a sceptical empiricist' is a much trickier example, but Invariantists will say, not that there is necessarily a determinate truth-value for that sentence (since it may involve vagueness and insufficiency of evidence), but that the sentence 'S knows that Hume was a sceptical empiricist' has a determinate truth-value, provided that there is a truth available, and that a good justification for asserting it is within reach of a serious Hume scholar. If something is knowable, then either you know it or you don't.

If we go along with Invariantism, and concede that there is a correct standard for knowledge, then what is that standard? Both the Invariantists and the Contextualists are keen on the observation that almost anyone who claims to 'know' something will waver in confidence, either if they are vigorously challenged, or if the truth of what they claim to know suddenly became very important. Almost anything that I think I know could be undermined if everyone in this room conspired to tell me I was wrong, and almost anything you claim to know would be repudiated by you if you were told that your life depended on really knowing it. Contextualists will want to say that such cases prove their point that the bar height can easily be shifted, but the Invariantists observe that in these circumstances we retreat to something normally labelled 'really and truly' knowing, and that this phrase brings us close to the objective standard for adequate justification.

It is immediately clear that 'really and truly knowing' sets the bar pretty high. But if we return to this question of the correct bar height, it seems to me once again that contextualism is the more inviting account. It would seem unlikely that anyone would propose to set the bar so high (as in my case of adding multiple new scepticisms to Descartes' doubts about the senses) that no one could possibly know anything. Equally we would reject the very low bar which permits any old riff-raff to join the epistemological elite. It strikes me as obvious that the bar of really and truly knowing is set at just the right height to produce experts. The only explanation I can think of for this, apart from conspiracy theories, is that we all want to act according to the best information available, and so we want to identify the people who possess that best information. We set standards of justification to find the right intellectual guides, be it in researching rose-growing, nuclear physics, or the correct way to categorise David Hume. If this doubt (about the relativity of experts) is not sufficient to make us doubt the Invariantist dream of a fixed standard, then the fact that any group could, I think, be persuaded to raise its bar even further seems to me conclusive. If a group is happy that it 'really and truly' knows the answer to some question (a group of research scientists, say), they could easily be persuaded to search for more conclusive evidence and justification, by a clever but more sceptical visitor. A further difficulty for Invariantists, I think, is that if they are asked to define their standards for really and truly knowing, there seems to be nothing they can offer except intuitions and speculations. They can talk of evidence, coherence, foundations, self-evidence, consensus, consistency, conservatism and probability, but there seems to be no meta-standard for arbitrating among these proposals. The fact is that we set the bar at the height which suits our society, just as witch doctors and religious fundamentalist do.

Two further points are worth making about the contextualist-invariantist debate. There is a difference of opinion among commentators over whether the paradigm case should be my own attempts to judge whether I know, or rather our judgements of other people's knowing. On the whole the contextualists are struck by the subjective problem, and invariantists say their case looks better when we judge other people. We sense our own standards shifting in different contexts, but when looking objectively at others we are more inclined to apply fairly rigid standards about whether someone really knows something.

Also it is, in my view, important to see that the debate is not really about a shift in the meaning of the word 'know', which invariantists claim, not unreasonably, does not change its meaning, even in philosophical discussion about scepticism; the claim is rather that it is the word 'justified' which shifts in meaning. When someone says 'I know this' we understand pretty well what they are claiming, but when they say 'I am fully justified in believing this' the word 'justified' can mean all sorts of things.

## A Different Approach

I am now sounding like a thoroughgoing relativist, but relativism is a view of which, in general, I disapprove. I am with Plato, who thought relativism was a slippery slope that ended in us according massive respect to the point of view of a tadpole. And yet the line of thinking I have sketched has, to me, a horrible inevitability. Please note that nothing in my sketch has challenged the notion of truth, since thinking that truth itself is entirely relative seems to be the proposal that truth is nothing more than what is believed, either by an individual or a social group. Such a proposal strikes me as intellectual suicide. Whether a clear and unambiguous proposition is or is not true seems to me to be a fairly determinate fact. But in my survey I have been addressing the standards by which we should judge that someone is in *possession* of the truth, rather than judging the truth itself of their belief. It is also part of my rejection of relativist that I am a thoroughgoing realist about the external world (even, as Tim Williamson put it at one of these weekends, a 'rothweiler realist' on the subject). The issue here I take to be whether, as Plato put it, nature has objective 'joints' which we are trying to pinpoint in our descriptions and theories. Anti-realists like to tell us that the slicing of reality is entirely a matter of human mental structure and human social convention, rather than capturing the facts. I take realists like myself to reject such views, and claim that many of our descriptions are successful in capturing reality. I assert my commitment to truth and reality at this point, in order to give a context to my horrified discovery that knowledge is riddled with relativism.

I do not wish to end this talk by putting the skids under knowledge, and then grandly leaving the building. If we take a broader perspective on the line of thought I have followed, it is notable that the quest began with Descartes' desire for certainty, and homed in on justification as the problem. This way of approaching the matter goes back to the Hellenistic Greeks, who became gripped by a scepticism they had initially found in the questions of Socrates, and which led to a desire to find what they called the 'criterion' for knowledge. It didn't take long for some of them to decide that no criterion for knowing could be found, and to reach a spectacular state of intellectual paralysis, most beautifully caught in the reports of the philosophising of Cratylus, who in philosophical debates would eventually do nothing more than waggle a finger, to express his total uncertainty. And yet Aristotle, who was contemporary with some of this vigorous scepticism, regarded it with abhorrence, laughing (for example) at thinkers who claimed that there seemed to be two realities, because you could see two realities if you went cross-eyed. So why was Aristotle not interested in this problem? He didn't just reject scepticism, but he also didn't spend time trying to find the criterion (and this despite the fact that his teacher Plato had explored the problem quite thoroughly in his dialogue *Theaetetus*). As we might put it, Aristotle wasn't interested in justification, and in searching for the rules for setting the bar at some specified height.

Aristotle's most direct discussion of epistemological issues is in *Posterior Analytics*. The agenda of that text (and this is where I wish to conclude my talk) is not with justification and knowing, but with explanation and understanding. The famous first sentence of Aristotle's *Metaphysics* is normally read as 'All men desire to know', but the scholar Julia Annas says that this translation is incorrect, and that he actually said that 'All men desire to understand'. As soon as you place understanding on the epistemological throne, thus deposing knowledge, the issue I have been discussing becomes, in my view, dramatically clearer. If we seek knowledge, then we seem to be intimidated by more advanced knowers than ourselves, or even future possible knowers, and these great experts are in turn intimidated by the even higher knowledge which might possibly be in the possession of the Archangel Gabriel, or of God himself. If we set our sights on understanding, however, the picture is quite different. On the one hand, understanding is a highly human activity, and understanding is an attainment which is available to people with limited experience and little education. Small children come to understand the causal connections that tie the world together, and can give highly cogent answers to the question 'Why?' in all sorts of circumstances. On the other hand we all recognise that understanding comes in degrees, and that the question of whether someone understands something is likely to be highly contextual. It only makes sense to ask whether I, or some person I am observing, understands something, even in the sense of 'really and truly' understands something, if we have some end in view, usually a practical one. If someone really and truly understands how some model of car works, we would probably expect that person to be able to analyse its faults, and to give it a good service. This is what we actually usually mean by an 'expert', rather than someone who can rattle off a set of well-justified beliefs. You don't need a PhD in Physics to understand a car, and a thorough understanding of a car has fairly well defined limits and obvious standards. If you are studying for an A Level, your notion of understanding aspires to the level of an A\*, not to world-class expertise.

If we ask whether there is such a thing as 'absolute knowledge', then Invariantists, and especially those who think that some justifications can have rock-solid foundations, such as sense experiences, or obvious a priori truths, are inclined to think that there could be such a thing. But if you ask whether there is 'absolute understanding' most people would laugh at such an aspiration. If we move from the quest for understanding to the connected quest for explanation, the implicit relativism in the whole game becomes even clearer. Explanations are what we ask for when we seek understanding for ourselves, and also when we want to test the understanding of other people. But the notion of an absolute explanation, or unique specification of what an explanation should consist of, is an absurdity. If you ask me for an explanation, this is going to be relative to our mutual needs. An explanation could be a superb one for an audience of children, but inadequate elsewhere. Even a high-powered explanation at a PhD viva interrogation might be cut off as going into more detail or depth than is required for the occasion.

Since I shifted my epistemological attention away from the quest for certainty, knowledge and justification, and towards the nature of understanding and the characterisation of a good explanation, I have cheered up considerably. There are still lots of juicy philosophical problems to worry about, and I will mention eight such problems. **First**, are all good explanations causal in character? Most explanations seem to involve an account of the causes of an event, and Aristotle liked this approach, but some explanations seem better labelled as 'structural' or 'constitutional' rather than 'causal'. **Second**, must a good explanation necessarily invoke a law of nature? This has been a popular view, but it strikes me as a rather superficial approach to explanation. **Third**, are there objective explanations – facts which will stand as explanatory for any enquirer – or is all explanation dependent on subjective interests (thus pushing us once again towards strong relativism)? **Fourth**, is a good explanation the thing which increases the possibility of an event? **Fifth**, could something be a good explanation, and yet be completely false? **Sixth**, is all good explanation reductive in character, giving an account of the reality which underlies some event, or can good explanation exist on the same level as the thing to be explained? **Seventh**, do some things demand an explanation, and other things have an openly self-explanatory character? And **finally**, and this is the area which has caught my main attention, in order to explain and understand the behaviour of things in reality, should we aim to define what is essential to them? And if we could manage such an account, would it be the bottom level of all possible explanation?

Despite these difficult questions, if we seek understanding there doesn't appear to be much danger of sceptical paralysis, and no danger of becoming entangled in an intense and doomed attempt to define the thing we are seeking. If I ask myself whether I understand my car, or the classification of Hume's philosophy, I can chart the degree of my understanding with confidence and reasonable accuracy, and the same applies if I ask whether other people understand such things. But if I ask whether I 'know' these things my head starts to swim a bit, especially if knowing is of the 'yes or no' kind, as in quiz shows. It takes a bit longer than thirty minutes to decide whether someone understands something, but we could probably arrive at a clearer picture of what *University Challenge* students understand than what they are labelled as 'knowing'.

The epistemological dichotomy I have presented here, asking you to choose between knowledge and justification on the one hand, and understanding and explanation on the other, is clearly somewhat simplistic, and I think there may be a meeting ground when we compare the notion of 'understanding' with justifications which are labelled as 'coherent'. Nevertheless, I think the second view is the more promising. When you track understanding and explanation back to Aristotle, for example, you discover the interesting fact that the famous Four Causes which he cites, and which have rather perplexed modern students of causation, are actually much better understood as the Four Modes of Explanation (since the Greek word *aitia* will translate both ways). Thus we can explain something by its matter, by its form or essence, by what gives rise to it, or by its ultimate purpose – known as the Material, Formal, Efficient and Final Explanations. Another interesting point about the second view of epistemology is that it seems more in keeping with the interests of modern scientists and other researchers, who care very much about finding good explanations, but care very little about what would trigger a status for their theories that got them labelled as something called 'knowledge'. To me the most striking implication of the dichotomy is that if you pursue knowledge and justification, you are led into questions of context and social need, and the ultimate epistemological battleground becomes politics and the competition to establish rival candidates as experts, so that the key concept is power, as explored in the writings of Michel Foucault. If, on the other hand, you follow the path of understanding and explanation, you find yourself in the thickets of metaphysics, trying to understand laws, causation, identity and essence. If for no other reason, I will continue to pursue the latter path, simply because it looks more interesting.

September 2010. [www.philosophyideas.com](http://www.philosophyideas.com)