

Knowledge, Belief and Justification

Table of contents

Introduction and Overview

1. **Belief and justification**
 - 1.1 **The dual-component view of belief**
 - 1.2 **Realist and instrumentalist notions of belief, behaviour and action**
 - 1.3 **Can beliefs be voluntary?**
 - 1.4 **Internalist and externalist theories of justification**
- 2 **Knowledge**
 - 2.1 **Kinds of knowledge**
 - 2.2 **The tripartite definition of knowledge**
 - 2.3 **Gettier objections to JTB**
 - 2.4 **Responses to Gettier**
 - 2.4.1 **Infallibilism**
 - 2.4.2 **Indefeasibility**
 - 2.4.3 **Appropriately caused beliefs**
 - 2.4.4 **No False Lemmas**
 - 2.4.5 **Epistemic virtue**
- 3 **Suggested strategy and techniques for teaching JTB**
 - 3.1 **Technique**
 - 3.2 **Sequence**
- 4 **Some helpful resources and further reading**

Introduction and overview

This document is intended as a teaching resource for nonspecialists teaching philosophy at A-level, and in particular for those teaching the AQA syllabus. By 'nonspecialists' I mean those who have little or no previous background in the subject. Other documents in this series will eventually cover all aspects of the syllabus, but for ease of use it makes sense to have specific resources for specific topics, units or sections. The syllabus deals with knowledge in different ways in different units. Here I have focussed exclusively on the section 'Knowledge, Belief and Justification' that appears in *Unit 3: Key themes in philosophy*. However, all this was based on the old specification. The new specification, for first examination in 2015/16, makes substantial changes, and I have added extra sections so that these changes are covered. The result is that if you are doing the new specification you will have more information than you need, but the structure of this document is based on the old specification.

I stay fairly close to the syllabus all the way: every heading in the section of the AQA specification that I am dealing with has its own numbered section here. There are also one or two sections that do not appear in the specification, but which it made sense to have for the sake of completeness and clarity.

This document is designed to be used in different ways for different purposes. Each section has a title and is numbered as part of the overall sequence. If you simply want to look up a certain topic, then you can consult the table of contents and go straight to a section you might want for a particular lesson, for example. On the other hand, it has been written as a piece of continuous prose, so if you read it straight through you will get a sequential journey through this section of the specification.

What I have tried to provide are short and concise but very clear explanations of what your students need to know about each subtopic (each section is a subtopic). I have pitched it at nonspecialist adults, so it should be accessible to those unfamiliar with the subject. Technical terms are explained when they are introduced, and I have tried not to take too much for granted.

The explanations given in the main text are designed to set out what you need to know about each sub-topic in order to teach it. However, in order to keep this document to a manageable length, those explanations are necessarily short, and you may wish to explore some of these topics further. One way to do this is to consult the further reading I have indicated, in combination with other resources, in order to get a more detailed picture. One thing I would warn you against, however, is randomly searching the internet. For every decent blog, article and encyclopedia entry there are a hundred really bad ones, and if you follow a bad source your students will suffer.

There is a separate section, after the main text, that contains a brief guide, including lesson sequence and techniques, on how to teach the traditional account of knowledge – as justified true belief.

It is important to be clear about exactly what the AQA syllabus requires of your students. On the one hand, the syllabus contains a significant volume of material, and some of that material is quite complex. It can seem quite daunting at first sight. On the other hand, your students are not expected to have in-depth knowledge about any particular component. That is good news, because it makes the teaching manageable according to a kind of formula. For each subtopic, your students should be able to do the following:

- (1) Define each position, in a few sentences at most, in a way that clearly distinguishes it from similar positions.
- (2) Explain the position using illustrations and examples.
- (3) Evaluate the position by identifying the standard arguments for and against it.

If you look at the marking scheme, you will see that 1-3 above are all necessary if your students are to attain the highest grades. But, crucially, your students do not need any more depth or detail than is required to meet 1-3 above. Bear in mind that the standard arguments for and against each position are well-known, and your students are merely expected to be able to rehearse these accurately. Teaching the syllabus successfully therefore depends on communicating these positions to your students in such a way that they *understand* them. That will stop the process being a mere test of memory, which in turn will free up your students' valuable cognitive resources for other tasks. If you can succeed in that, exam preparation will be much easier and results will be much better.

Since the syllabus requires only a relatively superficial understanding of a broad range of topics, these notes will provide only what is necessary for you to teach that far. Inevitably, this means leaving out a great deal of detail, argument and counter-argument that is found in the professional literature, with which neither you nor your students are expected to be acquainted.

1. Belief and justification

From the AQA specification:

- Belief: the dual-component view of belief (as advanced by, for example, Hume); realist and instrumentalist notions of belief, behaviour and action; whether beliefs can be voluntary.

1.1 The dual-component view of belief

Belief is one of a number of mental states that philosophers refer to as *propositional attitudes*. Beliefs are *attitudes* in the sense of *attitudes of mind*: psychological stances. If you currently believe that the UK will never experience a decent summer again, that is a fact about what is in your mind.

Propositions articulate thoughts. Propositions are not the same as sentences, since two different sentences can express the same proposition. Beliefs are *propositional attitudes* in the sense that their contents – the thoughts they contain – can be explicitly stated in sentences containing a *that-clause*. I did this above with the belief *that the UK will never experience a decent summer again*. The *that-clause* states the *content* of the belief, which is *what is believed* by the person who holds the belief.

We can take differing attitudes towards the same proposition. We might **hope** *that the UK will experience good weather*, but we might nevertheless not **believe** it.

Technically, philosophers think of belief as a relation between a person and a proposition. To believe that the UK will never experience a decent summer again is to stand in the relation of believing to the proposition 'The UK will never experience a decent summer again'.

Propositions, as expressed in declarative sentences (sentences which make a claim or say how things are) can be true or false. Consequently, beliefs too can be true or false, depending on whether the proposition in question is true or false. The fact that we can evaluate beliefs in terms of truth and falsity distinguishes them from other propositional attitudes such as hopes, fears, desires, and so on, which are not *truth-evaluable*.

If your belief is true, what is in your mind reflects how the world actually is. It is widely believed by philosophers that the point of holding beliefs is that one's mind accurately reflects how things are in the world; therefore we should aim to hold only *true* beliefs. Belief is often contrasted with desire in terms of 'directions of fit'. When we believe, we aim to

change our minds so that they fit how the world is; when we desire, we aim to change the world so that it fits what our mind wants.

This distinction between belief and desire in terms of direction of fit is often attributed to David Hume. Hume also tried to distinguish between beliefs and other propositional attitudes by claiming that beliefs are more vivid than any imaginative state like fear. This attempt to explain what it is to have a belief is widely regarded as a failure, mainly due to its implausibility (many of us experience our fears as being at least as vivid as our beliefs).

1.2 Realist and instrumentalist notions of belief, behaviour and action

I said above that if you hold a belief then that is a fact about your mind. This assumes that beliefs are part of reality – real things that exist in the mind. So if you acquire a new belief, for example, there is an item in your psychology that was not there previously. So the contents of your mind have changed.

Some philosophers think that the mind is identical with the brain, a very strong claim. Most philosophers accept a weaker claim: that the mind is very closely connected with the brain. If we believe either claim we might assume that a change in the mind will result in a change in the brain (e.g. altered neuronal pathways), and vice versa.

Additionally, it is widely believed that beliefs play an important role in intentional action. We cannot explain why you turned on the tap unless we attribute to you the belief that water would come out when you did so. Many modern philosophers have accepted the view that actions are caused by pairs of beliefs and desires: I want a drink, and believe that there is a drink in the fridge, so I take a drink from the fridge. And anything that can be a cause of something else must be real; so beliefs must be real.

All this points to the conclusion that beliefs are real things that play a role in the real world. This view is called *realism about beliefs*. It vindicates our intuitive, common-sense view of beliefs.

But, on a competing view called *instrumentalism*, attributing beliefs is merely a convenient way to explain the behaviour of something. For example, if Martians were discovered and they behaved much as we do, we would find it profitable to attribute beliefs to them even though we knew nothing about the structure of their brains and minds. If we find it useful to explain the behaviour of computers in the same way, then perhaps we should ascribe beliefs to them too.

For the purposes of explaining and understanding something, it might not matter if there is nothing real that corresponds to the idea of belief. The equator is just an idea – there is no line painted round the middle of the Earth – but it is still a very useful idea.

The idea behind instrumentalism is that there is not much more to having a belief than to be disposed to respond in certain ways to certain circumstances. The problem with this view is that inanimate objects like thermostats also have dispositions: if a thermostat activates a central heating system, should we say that it believes that the house is too cold? If not, it

may be that something's having dispositions is insufficient evidence for us to ascribe beliefs to it.

1.3 Can beliefs be voluntary?

It is widely believed by philosophers that we do not, at least ordinarily, choose which beliefs we have. We are compelled to believe only what we think is true. You can see why if you try to believe that you are now reading this while on the moon – you should find it impossible. Notice the difference between *imagining* that you are now reading this while on the moon – which is easy – and *believing* it, which should be very hard. So we cannot just will our beliefs.

Having said that, there are psychological phenomena where people believe things as a result of some psychological – as opposed to rational – influence. If I am in love with someone, I may continue to believe that they are innocent of a crime when the evidence says otherwise. If so, my belief in their innocence is the outcome of merely psychological (including emotional) processes rather than rational processes like reasoning – though perhaps it is more realistic to suppose that in such cases emotion distorts rather than replaces rationality.

There are a wide range of merely *causal* reasons why people hold certain beliefs, including mental illness, hypnosis, indoctrination and psychological phenomena such as confirmation bias (where we only acknowledge evidence that confirms what we are already predisposed to believe). Hypnosis is interesting because, in contrast to my other examples, it seems to offer the possibility of making a *voluntary*, and possibly *rational*, decision to acquire a belief through *causal* (psychological) means. In the other cases, we might not even know why we hold the beliefs we do.

However, philosophers tend to idealise persons as rational agents who think, do and say what it makes sense to think, do and say. Obviously, real people very often fail to live up to this ideal.

1.4 Internalist and externalist theories of justification

There are rival accounts of what justification is. In particular, there is a distinction between internalist and externalist accounts of justification. Confusingly, this distinction can be made in at least three different ways, though your students only need to know the first two, which are related to each other.

On the first way, internalists think that if you hold a justified belief, you will always have access to the grounds for that belief; externalists deny this. For example, if I believe that I am sitting at my desk right now, then I have access to the justification for this belief, which is the fact that I perceive that I am sitting at my desk. Alternatively, if I know that David Cameron is Prime Minister, then my justification is that I have got this information from a reliable media outlet, for instance. An externalist would agree that I can *sometimes* have this access, but deny that I *must always* have it – whereas the internalist insists on this.

The second way of making the distinction is related to the first, but importantly distinct.

On the internalist view, the items which justify a belief are immediately available to the subject – they are internal in the sense that they are internal to the subject’s consciousness, i.e. they are mental states. On the externalist view, those items are things in the external world. The two ways of making the distinction are related in the sense that it is more plausible that I will always have access to the justifications for my belief if those justifications are my own mental states.

Note that this distinction, applied above to justification, can also be applied to knowledge. In that context, an internalist will support what is called *the KK thesis*: if I know something, then I must know that I know it. An externalist will deny this. I mention this only to avoid possible confusion: the syllabus refers to “internalist” and “externalist” theories of *Justification*’ (my emphasis) – not of *knowledge*. But it would be easy for students to get confused about this under examination conditions, and if they go on to write an essay about knowledge instead of justification then they may not get any marks at all, because the examiner may take the view that they have simply not answered the question.

2 Knowledge

From the AQA specification:

- Knowledge: the tripartite definition of knowledge; ‘internalist’ and ‘externalist’ theories of Justification [I have dealt with these in s1.4 above]; Gettier-type objections to the tripartite definition and responses to Gettier, for example indefeasibility, whether beliefs are appropriately caused, whether they track the truth. Whether such approaches and responses are successful.

2.1 Kinds of knowledge

In epistemology, philosophers usually distinguish between three distinct kinds of knowledge. The first is *knowledge by acquaintance*, which we acquire by having some experience of whatever it is we have knowledge of. This is the kind of knowledge in question when we say things like ‘I know Paris well’ or ‘I’ve known Joe for many years’. The second is knowledge of how to do something, usually referred to as *practical knowledge* or *knowledge-how* (in philosophy) or know-how (in ordinary English). We acquire knowledge-how when we acquire skills or abilities, like riding a bicycle or speaking a language. The third is *knowledge-that*, or *propositional knowledge*, which we acquire by learning facts. For instance, you probably know that David Cameron was Prime Minister of the UK in 2013, that water is a liquid, and that you need a cup of coffee right now.

The relationship between propositional and practical knowledge is highly controversial in contemporary philosophy. At one extreme it is claimed that all practical knowledge can be reduced to propositional knowledge: every practical skill can in principle be codified as a series of facts. At the other extreme, it is claimed that all propositional knowledge can be reduced to practical knowledge: knowing facts is at bottom just an ability that rational animals have. There is also a wide range of less extreme positions. The AQA specification concentrates exclusively on propositional knowledge, and does not require your students to

be too familiar with this debate; but they *must* be able to distinguish *propositional* knowledge – defined as knowledge which is introduced into a sentence by a *that*-clause – from the other two kinds.

The specification focuses on the traditional account of propositional knowledge and modern criticism of that account. There is a distinctive line of reasoning that leads to the traditional account, and it is important that your students understand it. For this reason, I strongly recommend that you stick to a certain sequence when presenting this topic to them. Towards the end of this document is a section that provides detailed guidance on how to teach this central part of the syllabus, but for now I shall provide the same sort of summary as I have for the other sections.

2.2 The tripartite definition of knowledge

Let us put ourselves in the position of philosophers trying to formulate necessary and sufficient conditions for knowledge. What is knowledge?

It seems that there is something psychological or mental about knowledge: I can know that China has an official ideology, and my dog can know that it is time for its dinner, but my shoe cannot have knowledge at all. Can my computer? Or should we say that only things with minds can have knowledge? Certainly if computers could have minds then it seems likely that they could have knowledge. The mental aspect of knowledge is usually taken to be belief: I can't *know* that my eyes are green and not *believe* that they are. But belief is not enough for knowledge, since we can have *false* beliefs, and we can only have knowledge of *truths* (you can't *know* that you are a teacher unless you *are* one).

But true beliefs are not enough for knowledge either, since we can have true beliefs by chance, and these beliefs are not cases of knowledge. For example, if I guessed the winning numbers on the lottery last week, we would normally deny that I *knew* what the winning numbers would be since they were determined by chance – and only after I made my selection. It was just good luck, and a belief that is true due to sheer good fortune is not knowledge.

If you doubt this, consider the circumstances under which, in everyday life, you would and would not be prepared to say that somebody knows something, and ask yourself what difference luck makes. Consider as many cases as you can.

Having accepted that beliefs which are the result of guesses, coincidences, accidents, or luck exclude knowledge, we now need to establish what distinguishes knowledge from mere true belief. It seems that if my true belief is to count as knowledge then I must have a good reason for holding it. If so, what is required to turn true belief into knowledge is *justification*. If I have rigged the lottery then my *justified* true belief about the results *does* count as knowledge.

The traditional conception of knowledge is called the JTB (Justified True Belief) analysis, also known as the tripartite definition of knowledge. The idea is that justified true belief is both necessary and sufficient for knowledge. The JTB account is an *analysis* in the sense that it breaks knowledge down into three component parts – justification, truth, and belief. These

components give us individually necessary and jointly sufficient conditions for the correct application of the concept of knowledge. The JTB analysis tells us what knowledge *is*: nothing can be knowledge unless it is a justified true belief, and if something is a justified true belief then it must be knowledge; so knowledge *just is* justified true belief.

This is known as the tripartite analysis of knowledge because it lists three conditions which must be met if something is to count as knowledge. We use schematic letters to represent this: for any subject of experience (S), S knows some fact (p) if and only if conditions i-iii are satisfied.

S knows that *p* iff

- i. *p* is true;
- ii. S believes that *p*;
- iii. S is justified in believing that *p*.

2.3 Objections to JTB

Since the JTB analysis of knowledge purports to provide necessary and sufficient conditions for knowledge, there are two ways in which we can take issue with it. Firstly, we can deny that one or more of the conditions are necessary for knowledge. Secondly, we can deny that the conditions are jointly sufficient for knowledge. We shall take each of these in turn.

2.3.1 The individual conditions are not necessary

Is truth necessary for knowledge? Almost everybody thinks it is. But it has been denied. One way to do this is to look for examples of cases where something is known but not true. 'We knew that smoking was safe until we discovered that smoking causes cancer.' Is this convincing? Or should we say 'believed' instead of 'knew'? Someone can take themselves to know something, but if it later turns out that they are wrong then surely they never knew it – they only believed it. Most philosophers find it overwhelmingly plausible that you cannot know something unless it is true. If something is not true, then it cannot be known.

Is belief necessary for knowledge? Again, most philosophers think it is. If I come home and my house is on fire, I might say 'I don't believe it'. Since I know that my house is on fire, and since I have said that I don't believe it, it looks like we have knowledge without belief. But most philosophers would deny that your claim not to believe it should be taken literally. You do believe it, and your denial is, for example, an expression of shock. So, as with truth, it is not really controversial that belief is required for knowledge.

Is justification necessary for knowledge? Again, almost everyone thinks it is. But some have argued for "lightweight" knowledge – where we have truth and belief but not justification. If you ask your students what the capital of France is and they say 'Paris', then you could argue that they know what the capital of France is – even if they have no justification for what they think. The objections to this are to say either that there will be some justification for their beliefs somewhere, or that they don't really know – not in the sense we are interested in. Even philosophers who argue for lightweight knowledge usually admit that there is also "heavyweight" knowledge, which does require justification, and is what the JTB analysis is about.

2.3.2 The conditions are jointly insufficient - Gettier objections

In 1963, Edmund Gettier provided some counterexamples to the JTB analysis, which appear to show that JTB is not sufficient for knowledge. Here are two examples (neither of which are Gettier's own, but which are clear to most students):

(1) Suppose I look at the clock on the wall and thereby form the belief that it is 15:00, and suppose also that it really is 15:00. I have a justified true belief, so according to the JTB analysis I know that it is 15:00. But now suppose that the clock stopped at 15:00 yesterday, and only tells the right time for one minute each day. It is only by chance that my justified belief is true, and so it is not knowledge.

(2) Suppose I am watching Wimbledon on TV, and I watch Federer beat Nadal in the men's final. I therefore have a justified true belief that Federer beat Nadal to take the title. As it happens, my belief is true; but only through chance, because I have unwittingly been watching a replay of last year's final, where the same thing happened. Again, JTB is satisfied, but I do not have knowledge.

2.4 Responses to Gettier

Gettier counterexamples depict cases where someone has a true justified belief that does not count as knowledge because their belief is only true by chance. If these examples succeed – and most philosophers agree that they do – then they show that the JTB analysis is insufficient for knowledge. If so, then we do not know exactly what knowledge is in the sense of having a recipe for it. Since knowledge is a foundational concept for philosophy, philosophers have been unwilling to accept this conclusion, and the debate over how to respond to Gettier's argument continues in contemporary discussion. Below I take a brief look at the most prominent suggestions, since these are part of the syllabus.

2.4.1 Infallibilism

Infallibilism is the view that beliefs are only justified when they are certain (in the sense that things could not be otherwise), and consequently that we can only have knowledge when we can be certain that our belief is true. The argument for this view goes as follows. Since we can only have knowledge of truths, if I know that p then p must be true. I cannot be mistaken about p being true, because if p was false then I could not know that p . So if I know that p , then I must be justified in believing that p . Consequently, if I do not know that p , then I was not justified in believing that p in the first place.

The last sentence above contains the crucial claim. Ordinarily, we do not have to be certain of our beliefs in order for them to be justified: it just has to be rational to hold those beliefs, as demonstrated by our having a good reason for holding them. But infallibilism raises the threshold for a belief to count as justified: a good reason is no longer good enough – now we must be certain. The effect of this is that when I have a justified true belief I must have knowledge, and if I do not have knowledge then I do not have a justified true belief (since to be justified is to be certain). If we accept this line of argument then the JTB analysis is saved from Gettier, because his counterexamples now show only that in such cases our beliefs were not justified in the first place.

But should we accept this argument? There are at least two good reasons not to. Firstly, it is implausible to suppose that justification requires certainty, partly because this claim is inconsistent with our normal practice, which accepts a much lower standard for the justification of beliefs: having good reasons for them. If a theory diverges from normal practice, that is often a strike against it. Secondly, a belief can be justified and yet be false: we might have very good reasons to believe something and still be mistaken about what we believe. So justified beliefs need not even be true, let alone certain. This approach seems to fail.

2.4.2 Indefeasibility

A defeasible condition is one that can be defeated in the right circumstances; an indefeasible condition cannot be defeated. Recall the clock case above (one of the Gettier counterexamples). Here I have a perfectly good justification for my belief that it is 15:00 – I have looked at a clock. But in this case my justification is defeated by the circumstances, because it is only through luck that my belief is true. Supporters of the ‘fully-justified true belief’ theory of knowledge propose adding an extra condition to the three that make up JTB:

- 4 My justification for believing that *p* is not defeated by the circumstances.

If we add this condition in, knowledge is ruled out in this and all similar cases, removing the possibility of Gettier-style counterexamples since my justification cannot be defeated. So far so good. But this extra condition imports a difficulty of its own. Whenever knowledge is in question, I am always vulnerable to things I am unaware of. I might be mistaken but be unaware of this. Either I am right and I have knowledge or I am wrong and I do not. But in any given case I might be wrong about which position I am in. It is precisely this fact that is exploited by the Gettier counterexamples to JTB – they all depend on me not realising that I am watching a replay, or that the clock has stopped, or something similar. Condition 4 rules this out by appealing to the objective circumstances. But in so doing, it makes whether I have knowledge dependent on a factor that I cannot reliably assess: since I do not know what position I am in, I cannot tell whether the circumstances defeat my justification. The result is that I can *never* know whether I have knowledge, and that seems implausible, for there are surely at least some occasions on which I can know that I know something (the KK thesis – see S1.4). So this approach too seems to fail.

2.4.3 Appropriately caused beliefs

If a belief is caused in the right kind of way – for instance, as in perception, by the thing that the belief is about – then that might offer some guarantee that chance has been excluded. There are different versions of this approach, and what ‘appropriate’ means varies accordingly. Here I shall look at *reliabilism*, the view that knowledge is true belief that is caused by a reliable process. (This is enough to conform with the AQA specification.)

Note the difference between the following two claims:

1. Knowledge is JTB, and justification requires that the relevant belief is caused by a reliable process.
2. Knowledge is true belief that is caused by a reliable process.

(1) is reliabilism about *justification*; (2) is reliabilism about *knowledge*. Be sure not to get these two mixed up. Reliabilism about justification follows the same pattern as most other responses to Gettier: it tries to rescue the JTB analysis by adding an extra condition that is supposed to exclude accidental true belief. Reliabilism about knowledge, by contrast, abandons JTB, and in place of justification it adds a condition which says that the belief I hold must be appropriately caused.

A 'reliable process' is one that 'produces true beliefs in actual situations, but would produce true beliefs, or at least inhibit false beliefs, in relevant counterfactual situations' (Goldman 1976 p.771). A counterfactual situation is one that does not in fact obtain, as articulated in a type of sentence called a counterfactual conditional: What *would have* happened *if* Hitler had won the Second World War? Reliabilists think of visual perception as a reliable process, because although there are visual illusions, vision is generally a reliable way to acquire knowledge about the world. If I saw an apple, it is the apple itself that caused my perception of it and hence my belief about the apple. If there was no apple near me, I would not have perceived an apple and thus would not have formed the belief that there was one. So far so good.

But this view, in the form of either (1) or (2), still seems vulnerable to Gettier-style counterexamples. Henry is driving through the countryside with his son, and points out tractors, barns, silos and so on. He points out one barn in particular, which he is sure is a genuine barn, as indeed it is. The viewing conditions are good etc. So far, it seems clear that Henry knows that this object is a barn. But, unknown to Henry, this area is full of papier-mache facsimiles of barns. They look exactly like barns, but are actually only facades, without back walls or interiors. The object Henry has in mind is actually a barn, but if he were to encounter one of the facades he would mistake it for a barn. So it is only accidental that Henry is right about its being a barn. It therefore seems that introducing the notion of a reliable process, and especially if that process is causation, cannot eliminate the element of luck.

2.4.4 No False Lemmas

Another remedy that has been offered to cope with Gettier-type cases is the No False Lemmas approach. The idea is that we add a condition which specifies that the belief in question has not been inferred from any falsehood. That takes care of cases where such an inference occurs; but in many cases, such as the Barn County case above, no such inference is required. So this approach also fails.

2.4.5 Epistemic virtue

Finally, we could add a condition of epistemic virtue – that is, we could stipulate that, to be knowledge, the belief must result from skill on the part of the believer. But in the Barn County case above Henry's belief is the result of skill; yet still he does not have knowledge.

4 Suggested strategy and techniques for teaching JTB

4.1 Technique

The way to approach this is by asking the students questions in a carefully-structured sequence.

In order that your students understand what the point of the JTB analysis is, it helps to situate it within the methodological tradition from which it emerged. At least since Plato, philosophers have tried to define philosophically interesting concepts: causation, justice, beauty, meaning, art, knowledge and so on. It is very hard to do this: after 2500 years of Western philosophy, there are no uncontroversial successes. Nevertheless, we have learnt a great deal about these concepts from trying to define them. To get this across to your students, you might consider getting them to define concepts in class. DO NOT begin with philosophically interesting concepts – it is just too hard. Begin with everyday concepts, and move from easy ones to harder ones. Begin with *definition*, so that they are reminded what one is. Then move to *triangle*, since geometrical concepts can be defined relatively easily. Next an everyday object like *chair*, then *game*, and only then something like *art* or *meaning*. Do not spend more than two or three minutes on each concept. By the time you have spent ten or fifteen minutes on this, your students will have a vivid appreciation of how hard it is.

Next you need to introduce the concept of necessary and sufficient conditions. Start by making it explicit that you are moving from definitions, which govern the correct uses of words, to necessary and sufficient conditions, which govern the correct application of concepts. When we use a word, we call up the relevant concept (concept is just a modern word for idea, with the difference that only entities with minds can possess concepts, whereas ideas are a slightly more abstract notion). Concepts can be applied correctly or incorrectly – i.e. if I say that grass is purple I have applied the concept *purple* incorrectly.

By telling us when a concept is correctly applied, necessary and sufficient conditions tell us what it is to *be* something. For example, if we want to know what it is to be a bachelor, we can try to find necessary and sufficient conditions for the correct application of that concept. As it happens, *bachelor* is one of the rare cases where it is easy to supply these conditions: a bachelor is an unmarried man (definition); so there are two conditions for being a bachelor – that something is unmarried, and that it is a man. These two conditions are individually necessary (because without either one of them we do not have a bachelor) and jointly sufficient (because both of them are enough for something to be a bachelor) for the correct application of the concept *bachelor*. We can now say that x is a bachelor if and only if x is an unmarried man. And this is what philosophers are aiming for with concepts like *causation*, *art*, *meaning*, and *knowledge*. Again, get your students to try this themselves. I suggest the following sequence of concepts: *car*, *neighbour*, *horse*, *friend*, *justice*.

Now we have all that in place, it will be easier for your students to see the point of trying to find necessary and sufficient conditions for knowledge. We want to know what knowledge is, and this is how we find out. To engage them further, we want them to go on this journey themselves, insofar as that is possible.

Part of what philosophers do – one aspect of philosophical method – is conceptual analysis. By inspecting the way we ordinarily use words, we can arrive at conclusions about the content of the concepts we all share. And, since our concepts relate to reality, this method can reveal truths about reality. Otherwise the search for necessary and sufficient conditions would be pointless.

For instance, in everyday English we would not normally speak of house bricks or tyres having knowledge. From this fact about how we use words we can draw a conclusion about our concept of knowledge: that we restrict correct application of the concept 'knows' to certain types of object or entity, and tyres and house bricks are not among them. And, since our concept of knowledge relates to reality, we can in turn draw a conclusion about knowledge itself: that house bricks and tyres cannot have knowledge.

So the way we ordinarily use words is a kind of evidence that we can exploit to find out how things really are. We have to be quite careful with this, because our evidence will only be as reliable as our grasp of how words are correctly used, and anybody can get this sort of thing wrong, let alone A-level students. But in the case of knowledge we can use this method quite reliably to let the students develop the correct line of reasoning themselves with minimal prompting from the teacher. This approach also helps the students to understand that they are already familiar with the evidence (the way we use words) and that they already have the skills (their grasp of English) that we need to do philosophy in this way.

We will combine this method with another, called *Socratic questioning*, or the *Socratic method* after the ancient Greek philosopher Socrates, who (according to Plato) relied on it for teaching. The core of the method is that we ask students questions, and get them to give us the answers. The point of this is that it forces the students to think about the issues for themselves, which gives them a much more thorough understanding, and one they will remember, in contrast to spoonfeeding them facts and then asking them to regurgitate those same facts in an exam. Asking questions in this way is a skill that takes time to develop, but if you took a teaching course within the last ten years (and statistically most of you will have done) you probably took a module on this, and perhaps you even use the technique already. In any case, I have laid out below the precise sequence that works best for this topic. If you stick to the sequence, you should find that the students miraculously arrive at the same conclusion as the philosophers, i.e. the JTB analysis. The discussion is rigged, of course, since the sequence of questions is designed to elicit certain answers. But that does not matter here: you have to communicate the content to them somehow, and better like this than getting them to memorise something they don't really understand.

4.2 Sequence

Begin by asking the students about the different ways in which we ordinarily use the verb 'to know' and the noun 'knowledge'.

You can use the results to make some important distinctions:

- 1 Between the three main types of knowledge: knowing that, knowing how, and knowing by acquaintance.
- 2 Between different literal uses – for instance, we ordinarily say that animals can have knowledge, but is it the same thing for a human to have knowledge?
- 3 Between literal and metaphorical uses. For instance, we do not ordinarily think that tyres or house bricks can know things.

You can then move on to building the JTB analysis. I recommend that you start with belief, then move to truth, then to justification. So, following on from where you left off (tyres and house bricks) ask this question:

Can computers literally know things?

Hopefully, by considering the evidence, your students will arrive at the conclusion that only creatures with minds can literally know things. If they don't then you can short-cut this part of the process by asking them whether they think it is possible to know something but not believe it. Once they see that this is impossible you have the first condition of JTB in place, and it is all downhill from there. It is also worth making the connection between the ability to hold beliefs and the literal attribution of knowledge: bricks can't know things because they can't hold beliefs, and they can't hold beliefs because they don't have minds. Then continue with the sequence as below. If they are having trouble taking one of the steps just use examples to prod them along, as illustrated below.

Q: Is belief enough for knowledge? A: No.

Q: Why not? A: Because you can have false beliefs.

[Here you could ask, for example, 'If I **believe** that I'm a pink elephant then do I **know** it?]

Q: Are true beliefs enough for knowledge? A: No.

Q: Why not? A: Because you can have a true belief through good luck, and that is not knowledge.

[Here you should use an example of a lucky true belief, such as correctly predicting the winning lottery numbers. If they don't think that you knew the numbers in advance, then they should reject the idea that you can have knowledge through luck.]

Q: Why isn't it knowledge? What is missing? A: Justification.

4 Some helpful resources and further reading

The following sources may help to provide more depth and detail than I could provide in this document. However, while searching out additional materials may help you to understand and teach the topics better, please bear in mind that your students only need a certain level of depth and detail in order to pass the exam – if you provide more than they need then you run the risk of confusing or overloading them.

The links below all contain many further links and suggestions for further reading.

General resources for epistemology

<http://www.ucl.ac.uk/~kka7409/EpistemologicalResearch.htm>

<http://plato.stanford.edu/entries/epistemology/>

<http://www.rep.routledge.com/article-links/P059>

Instrumentalism

http://www.philosophybasics.com/branch_instrumentalism.html

<https://sites.google.com/site/minddict/intentional-stance-the>

Internalist and externalist theories of justification and knowledge

<http://plato.stanford.edu/entries/justep-intext/>

<http://www.iep.utm.edu/int-ext/>

The tripartite analysis of knowledge (justified true belief)

<http://plato.stanford.edu/entries/knowledge-analysis/>

http://www.reading.ac.uk/AcaDepts/Id/Philos/jmp/Theory%20of%20Knowledge/Analysis_of_Knowledge.htm

Gettier counterexamples to JTB

<http://philosophyfaculty.ucsd.edu/faculty/rarneson/courses/gettierphilreading.pdf>

(This is Gettier's original article.)

Infallibilism

<http://philpapers.org/archive/HOWIAG.pdf>

<http://documents.routledge->

interactive.s3.amazonaws.com/9781138793934/A2/Descartes/DescartesInfallibilismCartesiancircle.pdf

<http://www.iep.utm.edu/fallibil/>

Reliabilism

<http://plato.stanford.edu/entries/reliabilism/>

<http://www.iep.utm.edu/reliabil/>

<http://www.trinity.edu/departments/philosophy/sluper/The%20Reliabilist%20Theory%20of%20Rational%20Belief.pdf>