

WHEN DEDUCTION LEADS TO BELIEF

by Tobies Grimaltos and Christopher Hookway

Abstract

The paper questions the common assumption that rational individuals believe all propositions which they know to be logical consequences of their other beliefs: although we must acknowledge the truth of a proposition which is a deductive consequence of our beliefs, we may not genuinely *believe* it. This conclusion is defended by arguing that some familiar counter-examples to the claim that knowledge is justified true belief fail because they involve propositions which are not really believed. Beliefs guide conduct or issue in assertion by answering questions which arise in the course of deliberation and conversation, but the troublesome cases present propositions which do not present the agent's answer to any question. The paper concludes by sketching the conditions under which the deductive consequences of our beliefs can be believed.¹

1. Introduction

It is generally assumed that deductive logic systematizes norms governing the formation of rational belief through inference: if you are aware that some proposition is a logical consequence of your beliefs, and you have no independent reason to doubt it, then you ought, rationally, to add it to your stock of beliefs; and if you do not believe the conclusion of a deductively sound argument, and you have no doubt of its soundness, you cannot, rationally, remain confident of the truth of all of the premises. Our aim in this paper is to question this assumption, in the light of some claims about belief and inference and about the role of *questions* in structuring our beliefs and inferences.

Our target is a closure principle for belief. As a first approximation, we may express it thus:

If P is a deductive consequence of propositions that X believes, and X recognizes that this is the case, then X believes that P

Scepticism about this principle is not uncommon: if X is antecedently doubtful of P , then recognizing that it is a consequence

¹ The work contained in this paper was carried out during the academic year 1991–92, while Tobies Grimaltos was in receipt of a grant from the Dirección General de Investigación Científica y Técnica of the Spanish Ministry of Education. An earlier version was given at the University of Oviedo in 1993: we are grateful for comments received during a stimulating discussion on that occasion.

of his other beliefs may lead him to view his other beliefs with suspicion rather than accept *P*. Accounting for this would require us to be careful in arriving at an accurate specification of the principle, adding a rider to exclude cases where *X* has independent reasons for doubting *P*. And puzzles such as the lottery paradox may suggest that we can be agnostic about a proposition which is a consequence of a large number of beliefs if these beliefs are not absolutely certain. But the challenge we want to present is rather different: it does not depend upon features of unusual cases such as those just described.

Our claim is defended through close examination of some cases that have emerged from discussion of the 'Gettier problem'. Since Edmund Gettier published his original challenge to the thesis that knowledge was justified true belief in 1963,² the problem to which he gave his name has exercised an extraordinary attraction for analytical philosophers. As will be well known, the 'problem' is to provide a philosophical analysis of knowledge: a set of necessary and sufficient conditions for the truth of propositions of the form *x* knows that *p*. Its appeal might be explained by pointing out that some philosophers like trying to solve puzzles where the rules for a solution being successful are clearly stated: the analysis must match our intuitions about the truth values of knowledge sentences. But since most philosophers would now be sceptical of the claim that all (or even most) philosophically interesting notions can receive reductive definitions of the sort sought by the Gettier industry, it is curious that the need for such a definition of *knowledge* is so widely taken for granted. Why do we need such a definition? And why do so many philosophers whose views elsewhere in the philosophy of language show much sophistication, continue to look for a reductive definition in epistemology?

We do not propose to answer these questions here. However two observations which are relevant to answering them are worth noting. First, the counterexamples that are produced in the course of testing and refining definitions of *knowledge* are frequently of independent interest: cases are introduced where someone's epistemic position is plainly flawed and it is very hard to see how or why; and other cases are introduced where we cannot see why someone's beliefs possess epistemic merit. Hence the examples are interesting, as pointing to complexities in our practice of epistemic

² See Edmund Gettier 'Is Justified True Belief Knowledge', *Analysis*, 23.6, 1963, pp. 121–123. The examples to be discussed are all found in this paper.

evaluation, even if they do not direct us to a correct analysis of *knows* – or even if we do not attach much importance to the search for such an analysis. And second, the concepts introduced in the course of searching for an analysis have often been of independent interest even if they have not yielded an acceptable solution to the Gettier problem. For example, reliabilist and externalist themes in epistemology may originally have emerged in the course of attempts to define *knows* but their importance is now acknowledged independently of that connection. And few could deny the importance of Robert Nozick's work on tracking or Alvin Goldman's ideas about discrimination and perceptual knowledge: debates about causal theories of perception and about what it is to see something have been informed by the distinctions and conceptions uncovered by the Gettier debate, and these benefits are real whatever one thinks of the interest of the debate itself. If analyses of knowledge in terms of causation, reliability and defeasibility now have few defenders, these concepts have not lost their epistemological significance (a significance which was rarely remarked before they were used in furthering the Gettier debate). Thinking about the examples developed in trying to solve the Gettier problem can be epistemologically valuable even for someone with little real investment in 'solving' the original puzzle.

Gettier's original counter-examples to the claim that knowledge was justified true belief involved belief resulting from inference: true beliefs obtained through deductive reasoning from false but justified beliefs were plainly justified but were not knowledge. The believer who relies upon such an inference, we might suppose, is blameless: if you cannot trust the deductive consequences of your justified beliefs, what can you trust? On the other hand, in spite of the truth of the conclusion arrived at, the believer's epistemic position is flawed or unsatisfactory. Blamelessly acquired true beliefs can, for all that, be epistemically flawed.

We shall use these examples to raise some questions about the role of inference (particularly deductive inference) in producing new beliefs. We shall argue that Gettier's own counter-examples to the standard account of knowledge rely upon the closure principle for belief which we want to reject. And we shall find that reflection upon these examples will help us to see why the principle is false. If we are correct, it will follow that Gettier's putative counter-examples are nothing of the kind: they are not justified true belief because the belief condition is not satisfied. But we must emphasize that we are not offering a general solution to the Gettier problem.

We claim only that reflection upon a limited range of the examples thrown up by the debate can help us to gain a better perspective upon some questions about belief and inference.

2. Gettier's examples: deductive inference and belief

Our strategy, then, is to begin by discussing the examples presented by Gettier in his original 1963 paper and to see what can be done by way of rejecting them on the grounds that the propositions that are alleged to be cases of justified true belief are not genuinely *believed* by the agents in questions. We begin with Gettier's second example: we are to suppose that Smith forms the belief that Jones owns a Ford on the basis of strong evidence. Smith then selects three place names at random and constructs the following three propositions which, since they are entailed by the claim that Jones owns a Ford, he accepts:

Either Jones owns a Ford or Brown is in Boston.

Either Jones owns a Ford or Brown is in Barcelona.

Either Jones owns a Ford or Brown is in Brest-Litovsk.

Despite the evidence, Jones no longer owns a Ford: perhaps he sold his Ford last week and Smith saw him driving a rented car. However since 'by the sheerest coincidence and entirely unknown to Smith', Brown has just landed in Barcelona, Smith's belief in the second of these propositions is both true and justified.

The underlying structure of the example can be represented thus:

1. X believes that P and X has reason to believe that P
2. X is aware that P entails $P \vee Q$.
3. X thus has reason to believe $P \vee Q$
4. X infers and thus believes that $P \vee Q$.
5. Although P is false, Q (and hence $P \vee Q$) is true.
6. So X has a justified true belief that $P \vee Q$.

The assumption we have identified is involved in step 4. If the assumption is abandoned, one might question whether, in this case, X can infer the disjunctive proposition. Or, more paradoxically, we could allow that he can infer it, but deny that his attitude towards his conclusion is one of belief. But why should that be a plausible claim? At first sight, it seems absurd.

One way to defend it would be to point out that beliefs have a role in the explanation of action and inference: if we know what

someone believes, we can understand and sometimes predict their behaviour. Hence we should ascribe to people no more beliefs than is warranted by this explanatory need. To ascribe to someone a belief that was explanatorily inert, in that case, makes no sense. Perhaps we can argue that the disjunctive 'belief' in Gettier's case is unreal because it can have no explanatory role. In general, of course, disjunctive beliefs can have a role in explanations of actions and inferences. If I believe that my copy of Kant's *Critique of Pure Reason* is either at home or in my office, and my search of my office is fruitless, I shall go home to find it. In effect, the disjunctive belief is used as a premise for a step of disjunctive syllogism: once I learn that one disjunct is false, I plan my activities on the basis of the other disjunct. In Gettier's example, this does not occur: it would be irrational for Smith to conclude (on hearing that Jones has sold his Ford) that Brown is in Barcelona, Boston and Brest-Litovsk. Indeed, any piece of Smith's behaviour which could be explained as appropriate given the truth of one of the disjunctions could also be explained by reference to his non-disjunctive belief that Jones owns a Ford. This even explains his assent to the disjunctive proposition when asked its truth value: since he believes that Jones owns a Ford, he can see that *Jones owns a Ford or Brown is in Barcelona* must be true. It also enables us to understand why he is not tempted to draw any conclusion about Brown's location on hearing that Jones is no longer a Ford owner. We might argue that our readiness to make such a step of disjunctive syllogism is a mark of a disjunctive belief being 'psychologically real'. So if, when challenged, Smith acknowledges the truth of 'Jones owns a Ford or Brown is in Barcelona', we could explain this by saying that Smith believes that Jones owns a Ford and he knows that this entails that Jones owns a Ford or Brown is in Barcelona. There is no need to ascribe to him a psychologically real belief in the disjunctive claim.

The 'inference' involved in the example we have discussed involved a step of 'or-introduction'. We shall not be surprised to find strictly analogous examples which involve the parallel inferential process of existential generalization. Thus Keith Lehrer presents an example of a teacher who believes on good evidence that one of his students, a Mr Nogot, owns a Ferrari, and infers that at least one person in his class owns a Ferrari. This belief, it is alleged, would be justified and true if Nogot was not a Ferrari owner while another student does, in fact, own one.³ In line with

³ See Keith Lehrer, *Knowledge* (Oxford: Oxford University Press, 1974), p. 19.

the argument of the previous paragraph, we could reasonably argue that for the existentially quantified belief to be psychologically real, the teacher must be prepared to conclude 'So it must be someone else' on hearing that Nogot does not own a Ferrari. Unless that condition is satisfied, we can explain all of the teacher's actions and inferences by reference to the belief about Nogot without appeal to the existentially quantified belief: and we explain his avowal that someone in the class owns a Ferrari by reference to his belief that Nogot does.

Gettier's own first example is slightly more complex – but only slightly so. Smith and Jones have applied for a job. Smith believes, with reason:

1. Jones will get the job and Jones has ten coins in his pocket

He concludes:

2. The man who will get the job has ten coins in his pocket.

In fact, of course, Smith himself will be appointed to the post, and he too has ten coins in his pocket. The principal difference between this example and Lehrer's is that the conclusion involves a definite description rather than an existential quantifier. It is easy to show that this does not affect the substantive point being made. To see this, we must take into account the pragmatic ambiguities to be found in uses of definite descriptions: are we to take the use of 'The man who will get the job' in (2) referentially or attributively? If the former, then the assertion is simply an alternative expression of the belief that Jones has ten coins in his pocket and does not present a counterexample to the traditional analysis of knowledge. If the latter, the assertion has a distinct content but, as with the other two examples, it is reasonable to deny that it is psychologically real. On hearing that he rather than Jones has the job, he will not exploit (2) and conclude 'So I must have ten coins in my pocket'. Nothing is explained by belief in the conclusion which is not, as easily, explained by belief in the premises. Psychological reality of a belief calls for a sort of autonomy or robustness which is lacking in this case.

We are not claiming that any general morals about how to solve the Gettier problem can be derived from this. Thinking these examples through is a useful way of focusing on some issues about the role of inference in the enlargement of our knowledge, but it probably does not promise a general recipe for defending a traditional conception of knowledge. But we would expect the

suggestions made so far to convince no one. Given the initial plausibility of the claim that, for example, the teacher in Lehrer's example believes that at least one student owns a Ferrari, then we would expect the reader to question our initial suggestion about when a belief is psychologically real or to question its relevance to the epistemic issues under discussion.

It would be natural to complain that if Smith would assert the disjunctive proposition that *Jones owns a Ford or Brown is in Barcelona*, or if he would assent to it, then we are correct to ascribe to him a belief in that proposition. Even if our claims about when beliefs are psychologically real are correct, the critic could suggest that there is a kind of ambiguity in the concept of belief and conclude that there is a second conception which is at issue when we are doing epistemology. Failing that, a counter-example to justified true 'assent' (or 'preparedness to assert') might be as interesting as one to justified true belief. We shall address this below (sections 3 to 6). To forestall a second line of attack, we should note that we are not arguing that deduction never produces belief: our argument depends upon special features of the deductive inferences under discussion, for example their great simplicity. We shall discuss when deduction does generate belief in section 7; and we shall consider what a true Gettier style counter example relying on deduction would have to be like.

3. *Two conceptions of belief?*

The argument presented in the last section rests upon the assumption that the concept of belief employed in the definition of knowledge as justified true belief is the same as that invoked in familiar styles of belief/desire explanation of behaviour. We argued, in effect: since the inferential 'beliefs' that Gettier and Lehrer invoke do not have a fundamental role in the explanation of inference or action, they are not real beliefs. However it is natural to object that we have relied upon too restrictive a conception of belief; it does not do violence to our everyday concept (which is presumably what is appealed to in the traditional analysis of knowledge) to ascribe belief in a proposition to someone who is ready to assert it, to commit him or herself to its truth. If the teacher would bet on the proposition that someone in his class owns a Ferrari, or if he would volunteer that information when asked about the car owning habits of his class, then, plainly, he believes the proposition.

We can clarify this point by distinguishing two conceptions of belief:

*Belief*₁: somebody believes₁ something when their acceptance of it has a fundamental role in accounting for their behaviour;

*Belief*₂: somebody believes₂ a proposition when they are disposed to assert it.

In that case, the objection runs, we have argued only that someone may believe₂ a proposition which he or she does not believe₁. Consider the first example we discussed. Smith's 'belief' that Jones owns a Ford or Brown is in Barcelona, we urged, is not psychologically real: it is not a case of what we have now called belief₁. But, surely, Smith does stand ready to assert that proposition: he stands in the belief₂ relation to it. We have shown, at best, that the psychological basis of belief₂ in a proposition can in fact be the state of believing₁ some different proposition. This is relevant to the evaluation of Gettier's counter-examples only if the traditional analysis of knowledge identifies it with justified true belief₁. But that is very implausible; the traditional analysis of knowledge would surely hold that a justified readiness to assert a true proposition is sufficient for knowledge: knowledge is justified true belief₂.

Since assertion is itself a species of action, this distinction of senses of 'believe' may seem questionable. We shall now argue that the intuitive distinction underlying the response is best accounted for within the perspective that we defend. First, it is doubtful that the propositions in question are believed even in this putative second sense. Does Smith stand ready to assert this peculiar disjunctive proposition? It is difficult to imagine circumstances in which he would be able to do so, in which it would be intelligible that he would freely make the assertion that Jones owns a Ford or Brown is in Barcelona. Any such assertion would naturally be taken as expressing a psychologically real disjunctive belief – or as trying to induce such a psychologically real belief in an audience. Plainly if someone were to ask him whether the proposition is true, he would acknowledge that it was: this is the sort of thing that happens in a logic class or logic examination. This gives him an opportunity to *assent* to the proposition, but that falls short of actively asserting it. So we shall allow that you can assent to propositions which you do not believe₁, but that you can only assert propositions which you either believe₁ yourself or those which you believe that your audience should rationally believe₁. So far, the

distinction between two kinds of belief is of little help unless it is used to identify those propositions to which you would assent. What the objection rests upon is the evident truth that Smith will assent to these disjunctive propositions which he may not believe₁. In such a case the subject believes a certain disjunctive proposition to be true – this is a belief₂ and may even be a belief₁. Our claim will be that believing a proposition to be true is all that is required for ‘assent’ and that this is distinct from believing the proposition: to believe₂ that *p*, and to believe₂ that *p* is a true proposition, are not always the same.

This distinction may also seem unpromising. Indeed, one of the ‘obvious’ intuitions used to motivate deflationary theories of truth is that to believe that *p* and to believe that it is true that *p* are the same. The belief about the truth of a proposition may be described as second order. The reader may be reminded of a more familiar distinction: someone who knows no Spanish may believe that the sentence ‘En España, los buzones de correos son amarillos’ is true without believing that in Spain, post boxes are yellow. Other examples may be introduced without invoking other languages: if someone is ignorant of Physics but has overheard remarks on a radio documentary, they may believe that ‘Electrons have negative charge’ is true without believing that electrons have negative charge. Our examples differ from these: the ‘second order beliefs’ concern ‘propositions rather than sentences; and there has been no suggestion that these propositions are not understood. Unpromising as it may seem, we believe that something can be done with this distinction: we shall argue that if we understand the notion of the content of a belief or assertion correctly, these odd ‘propositions’ do not identify such contents. In a defensible sense, the second order claims do specify the contents of possible assertions while the corresponding first order propositions do not. In that case, since our elucidation of belief₂ linked it to a readiness to *assert* the proposition in question, it is reasonable to claim that they are not believed₂.

4. *Knowledge, assertions, questions*

As we noted in the first section, most epistemologists have focused on claims to propositional knowledge. For example, contributors to the Gettier debate have tried to offer analyses of sentences of the form *X knows that p*, where the complement to ‘knows’ specifies the content of a proposition. This has the advantage of preserving a

logical and grammatical parallel between knowledge and belief. On the assumption that knowledge is a variety of belief (belief which is true and justified, belief which tracks the truth etc), then it is not surprising that these parallels are emphasised. Knowledge takes a propositional complement because it involves belief which is a propositional attitude.

Moreover if these expressions are used to refer to propositional attitudes, then another source of the implausibility of some of our claims emerges. It is natural to assume that such attitudes can be taken to any proposition: if, say, I have a pro-attitude towards the truth of a proposition, then I believe it. And since connectives such as 'or' can be used to form complex propositions out of simpler ones, it is unsurprising that Gettier's subjects form beliefs in these disjunctive propositions. However this strategy can prevent our taking full account of distinctive features of the logic of 'knows': these grammatical parallels between 'knows' and 'believes' are superficial.

We can see this by noting that many knowledge claims (possibly most of those that are made in ordinary conversation) do *not* concern propositional knowledge. Consider the following examples where 'knows' takes an indirect question complement:

Jones knows where Brown is.

The teacher knows who owns a Ferrari.

The Physics professor knows why water expands on freezing.

The weather forecaster knows whether the sun will shine tomorrow.

We shall refer to these knowledge claims as *Q-claims* and the more familiar propositional knowledge sentences as *P-claims*.⁴ It is instructive that there is no indirect question form of belief-sentence: 'John believes where Brown is', 'The teacher believes who owns a Ferrari' are ungrammatical although we can use awkward expressions such as 'The teacher has a belief concerning who owns a Ferrari'. The significance of such grammatical disanalogies is always controversial, but we believe that focusing upon *Q-claims* is a source of considerable philosophical insight.

We are often interested in whether someone possesses knowledge because we wish to escape from ignorance concerning the answer to some question. If I am agnostic about why water expands on

⁴ The approach to the understanding of knowledge defended here was developed in C. J. Hookway *Scepticism* (London: Routledge, 1990), chapter X.

freezing, then when I learn that the physics professor knows this (knows why water expands on freezing), I shall be ready to accept whatever answer he offers to the question: he or she is identified as a reliable or trustworthy source of information. I then rely upon a distinctive form of inference which we refer to as the *knowledge inference*.⁵

X knows *Q*
 X's answer to *Q* is *P*.

So, *P*.

So we can think of the possessor of knowledge as having the ability to answer a question correctly. As Edward Craig has emphasized, we use the concept of knowledge in the course of identifying sources of good or reliable information.⁶ This need not only be exploited in the context of learning from someone's testimony. Suppose that I am aware that a friend is undertaking a journey which requires her to catch the morning train to Manchester. In the course of making her journey, she has to put to herself the question *When is the morning train to Manchester?* Unless she 'answers herself' correctly, her overall plans will not succeed. So I can be interested in whether she has the correct answer to a question because she needs this answer in order to succeed in some project. Indeed, one way to think of the role of beliefs in guiding conduct is that they provide answers to the different questions that arise in the course of planning a course of action. So if someone believes that *P*, then we can expect that whenever a question arises to which *P* is an answer, then, so long as then can see that *P* is such an answer, they will act on the basis that *P*, *ceteris paribus*.⁷

So questions have a role in ascribing knowledge to people, in planning courses of action, and in seeking testimony. If you meet my request for information by asserting that the next train out of town is tomorrow morning, then your assertion is explicitly the answer to a question. It is natural to hold that even those assertions

⁵ This formulation differs slightly from that found in Hookway, *Scepticism*. The reasons for the change can be found in Christopher Hookway, 'Knowledge, Questions and Context: a response to Fogelin', *Analysis*, 53.3, (1993), pp. 132–136.

⁶ See Edward Craig, *Knowledge and the State of Nature*, (Oxford: Oxford University Press, 1990), *passim*.

⁷ We need this *ceteris paribus* qualification because the agent may have other answers to the question available, or because recognizing that *P* is an answer to *this* question can contribute their belief being weakened.

which are not responses to explicit questions, can be understood only when we recognize that the speaker saw that a question arose in the context and was responding to it. If we can see no intelligible question that the speaker could be answering, we can make no sense of the linguistic action that is performed. So if we think of knowledge as, primarily, the ability to answer a question, we can understand the connections between knowledge, action and the planning of conduct, and assertion. And if we think of beliefs as stored answers to questions that might arise in planning actions or in conversation, then there are prospects for explaining the connection between knowledge and belief.

If these connections can be clarified and made precise, then the following principles hold:

- I. X can believe₁ that *P* only if there exists a question to which *P* is the answer.
- II. X can assert (i.e. believe₂) that *P* only if there exists a question to which *P* is the answer.
- III. X can know that *P* only if there exists a question to which *P* is the answer.

In that case, we must defend our view of Gettier's original counter-examples by showing that there are no questions to which these propositions are answers. And we must defend our claim about knowing propositions to be true by showing that there can be a question to which '*P* is true' is the answer even if there is no question to which '*P*' is the answer.

Thus there is a question 'Who will get the job' to which Smith, in Gettier's example, possesses the wrong answer; similarly, the teacher possesses the wrong answer to the question 'Who owns a Ferrari?'. And Smith has an incorrect answer to the question 'What kind of car does Jones drive?'. This much is straightforward. If Gettier's examples are genuinely to count against the claim that knowledge is justified true belief, there must be genuine questions to which his subjects possess true, justified but epistemically flawed answers. It is this that we must now consider.

5. Questions and Gettier's Examples

Let us now return to the case of Smith and the proposition that Jones owns a Ford or Brown is in Barcelona. If this proposition is to be asserted, known or believed, then, in the light of the three principles stated above, there must be a question to which it is an

answer. Before asking what this question might be, we should look at less problematic examples of disjunctive knowledge or belief. Suppose that somebody believes that the next train to London leaves at either 16.10 or at 16.40. Such a subject does not know when the train leaves, nor do they have a definite opinion about which answer to that question is the correct one. But their epistemic position with respect to that question is not wholly inadequate; they have some information with a bearing on the matter. If we associate with a question its set of possible answers, then our subject can rule out many of the answers that the question can receive: indeed, they can narrow the range of choice down to two. They know whether the next train leaves at 17.00 and they know whether it leaves at 15.50 and so on. Their 'disjunctive' belief is an acknowledgement that they can narrow the range of answers that need to be taken into account in discovering the answer but that they cannot lead us directly to the correct response. This is why the disjunctive knowledge licenses the use of disjunctive syllogism: once we or they can eliminate one of the disjuncts, we can recognize the other disjunct as the correct answer to the question. This suggests that we can make sense of disjunctive belief, assertion and knowledge when there is a question to which each of the disjuncts is a possible answer. We assert disjunctive propositions when we wish to provide a partial but incomplete solution to an implicit or explicit question. A disjunctive proposition can never be the complete answer to a question.

For it to be any kind of answer to the question at all, then, we suggest, it has to rule out some of the possible answers. If a disjunctive proposition were simply the disjunction of all possible answers, then it is not even a partial answer to the question and cannot intelligibly be asserted, known or believed. Consider the proposition that *Libya is in Europe or Asia or Africa or North America or South America or Australasia or Antarctica*. It rules out no possible answer to the question *In what continent is Libya?* So long as that question provides the context of the assertion, that claim is a kind of pragmatic tautology: it cannot be sensibly asserted. At best it records one's understanding of the question but it is not even a partial answer to it.

Let us now turn back to the second Gettier example. Our question is: can a single step of 'or-introduction' ever give rise to a new belief or to something that the speaker can assert? There are two different cases to take into account. First, suppose, varying Gettier's example, that Smith infers *Jones owns a Ford or Jones owns a*

Peugeot. If we have in mind the implicit question *What kind of car does Jones own?* then, in this case, Smith disjoins two possible answers to that question. But in that case, the disjunction is not Smith's answer to the question: his answer is *Jones owns a Ford*. This is evident from the fact that evidence that Jones has sold his Ford will not lead Smith to infer the other answer. Whenever the question arises, in conversation or in planning his actions, Smith should go by *his* answer and not by this disjunction. Hence Smith infers a conclusion that can have no role in sincere assertion or in planning conduct: the proposition is neither believed₁ or believed₂. There is no question to which the disjunction is Smith's answer.

The second case, Gettier's own example, offers a further complexity. It is far from clear that there is a question to which the propositions *Jones owns a Ford* and *Brown is in Barcelona* are both competing answers. We are strongly tempted to deny that there is such a question; but since we fear that someone with more ingenuity than we have could come up with a complicated counter-example to that claim, we shall simply state that we do not know if there is one. If there is such a question, the case parallels the one just considered. If there is not, then the proposition is not one that can be asserted, known, or made the object of a psychologically real belief. This is because there is no scenario in which it could be appealed to in order to answer a question that confronts us. In general, if I confidently believe one disjunct of a disjunctive proposition, there can be no question to which the whole disjunction is my answer, in which case no sense can be made of my sincerely asserting the proposition, believing it, or knowing it.

The other Gettier examples, together with Lehrer's case of the teacher and the Ferrari do not differ in kind from this one. Since our response to these examples can be guessed from the previous paragraphs, we shall not consider them further here. However, we should comment on an interesting variant of Lehrer's 'Nogot' example (section 2 above). The teacher recalls once believing of a particular student that he or she possessed a Ferrari but now forgets which student that was. Consequently he now believes only that *some* student has a Ferrari, inferring this from the recollection that he once believed of a particular student that he or she did so.⁸

⁸ There is a further variant where the quantified belief is caused by the original forgotten belief but the teacher has no recollection of where it originated. Such a belief would be psychologically real, and so might provide the basis of a genuine counter-example; but one might also reasonably doubt whether it is really justified.

If the quantified belief is made true by the car ownership of a student other than the one that was the object of the lost belief, do we then have a Gettier style counter-example? This conclusion can be avoided by claiming that even in this case 'some student in the class' functions referentially rather than as a straightforward existential quantifier. Confronting the question of who owns a Ferrari, the teacher commits himself to the correctness of a particular answer which he can specify only indirectly: he is committed to the correctness of the answer he once explicitly endorsed. In that case, the ownership of a Ferrari by some other student would not make the belief true.

6. *Knowing whether a proposition is true*

We claimed above that although Gettier's subject does not believe that *Jones owns a Ford or Brown is in Barcelona*, he does believe that this disjunctive proposition is true. We must now explain how this is possible. Consider a simple question such as *When does the next train leave?* If we were to list its possible answers, it would contain:

The next train will leave at 16.00

The next train will leave at 16.10

The next train will leave at 16.20

and so on. In this case, the different answers have the same subject or topic (the next train) and ascribe a range of different properties to that subject. Moreover the different properties belong to a set – e.g. the set of times after 16.00 this afternoon: they are all related times. So in choosing an answer, we decide which of these properties belongs to this thing. We might schematize this:

a is F_1

a is F_2

a is F_3

Not all questions have this simple form. And even among those that do, not all exploit lists of properties that can be worked out in the a priori way that is possible with this list of times.

If someone is to have genuine knowledge of the disjunctive proposition involved in the Gettier example, there would have to be a question whose list of possible answers includes:

Jones owns a Ford

Brown is in Barcelona

This does not fit the pattern we gave above: the answers have different subjects, and the predicates they use, ('owns a Ford', 'is in Barcelona') lack the systematic unity found in our simple case. If we are to find a question to which this disjunction was an intelligible response it would probably be a 'why-question': we are considering two explanations of (say) a murder, and the first can be sustained only if Jones owns a Ford whereas the second requires Brown to be in Barcelona. If we are convinced that one of these explanations is correct but cannot work out which, we might intelligibly carry out an investigation into whether Jones owns a Ford or Brown is in Barcelona. But this is an unusual case and is not involved in Gettier's example.

On the other hand, when we affirm that *Jones owns a Ford or Brown is in Barcelona* is true, then we answer a question with the possible answers:

Jones owns a Ford or Brown is in Barcelona is true
Jones owns a Ford or Brown is in Barcelona is false.

This fits the 'a is F_n ' pattern perfectly. And we can easily imagine situations – in a logic class, or in the course of writing a textbook – in which this question can arise. Even defenders of a deflationary theory of truth can agree that the concept of truth has value in discussing the semantics of propositional languages.

7. *When does deduction yield belief?*

In this concluding section, we address two questions. We have shown that certain deductive 'inferences' do not produce belief (or knowledge). Does this result apply to all putative beliefs resting upon deduction? Or are there cases where deduction leads to genuine (psychologically real) belief? And (this is the second question) are there genuine 'Gettier cases' based upon deductive reasoning? It would be interesting to go on to consider whether the strategy we have used in addressing Gettier's counter-examples could be applied to other cases from the literature too – including cases where no deductive reasoning is involved. Unfortunately, we lack the space to do that on this occasion.

So, first, when does a chain of deductive reasoning give rise to a genuine (psychologically real) belief? Two different sorts of answers could be given. First, we could argue that an inferred belief is psychologically real when the inference is sufficiently complicated that it is impracticable to rely upon the premises and repeat the

inference whenever it becomes necessary to answer a question to which the inferred belief is relevant. Or second, we could argue – relying upon the views about questions that have been defended here – that an inferred belief is ‘real’ whenever there is a real question to which it is an answer. In order to assess these answers, let us consider an example of a belief which would be real by the second criterion but not by the first. Do you believe that Madonna is mortal? If ever the question arose, you could infer this from:

Madonna is a human being
All human beings are mortal.

There are interesting questions in Cognitive Science concerning how such information is stored. It is implausible that (at least before today) you stored your view of Madonna’s mortality as a separate item of information: the information is readily available to you, whenever required, without being separately stored. This might suggest that it is not ‘psychologically real’. On the other hand, if anyone were to proclaim her immortality, you would (we hope) be affronted that one of your better embedded convictions was contradicted. This seems to be a case where the conclusion of the inference is thought of as an object of belief, however it is stored, and however simple the inference may be which is required to elicit it.

This suggests we should adopt the second approach. On this occasion deductive reasoning yields belief (even if you have not yet carried out the deduction) because there is a question to which the deduced proposition is an answer (partial or complete). This seems correct although further qualification is needed. Consider the following example.

John believes that the murder was committed by someone who was (say) a body builder, and he is aware that Peter and Paul were the only body builders who had the opportunity to commit the crime. He also happens to believe that the murderer made his getaway by car, and he has the information that Paul is a non-driver.

From the first two pieces of information he can deduce that Peter or Paul was the murderer; and this is a partial answer to his question. But, we assume, he does not believe this because he also has available information that rules Paul out of consideration. We would expect him to believe that Peter was the criminal. Perhaps he can believe a proposition when there is a question to which it is

the strongest answer that he can infer from the rest of his beliefs. Or must it be the strongest answer which he does infer? If he fails to notice the relevance of the information about Paul's being a non-driver, then, it seems, he believes the disjunctive proposition. Thus our settled answer is:

Deduction gives rise to belief when there is a question to which the conclusion of the inference is the strongest answer possessed by the believer.

The answer to the second question – can deductive reasoning ever give rise to Gettier examples? – is 'Of course'. Amend the example just considered as follows: in spite of strong evidence to the contrary, Paul can drive; but, in spite of that, Peter is the guilty party. Our subject may still (through deduction) arrive at a true, justified but epistemically flawed identification of the murderer. The case differs from Gettier's own examples because the inference provides an answer to a question which is psychologically real and is answered by none of the premises alone.

The moral to be drawn from this discussion – that the study of questions should be central to the epistemological study of belief and inference – provides another example of the fertility of reflection on the Gettier examples. Inference is an attempt to arrive at an answer to a question; believing a proposition involves using it to answer questions; and knowledge is a relation between an agent and a question. Treating propositional attitudes in abstraction from this context distorts our understanding of cognition.

*Universitat de València
Facultat de Filosofia i CC. de l'Educació
Blasco Ibáñez, 21
46010 València
Spain*

*Department of Philosophy
University of Birmingham
Birmingham B15 2TT
United Kingdom*