

EPISTEMIC VALUE AND FORTUITOUS TRUTH

COLIN CHEYNE

University of Otago

ABSTRACT

Why are the conditions for propositional knowledge so difficult to discover or devise in this post Gettier age? Why do not most epistemologists agree on roughly the same analysis as they appear to have done in the pre Gettier paradise? I argue that the problem lies in that fact that the epistemologists' intuitive concept of knowledge appeals to desiderata that probably cannot be satisfied. Unfortunately, if we abandon some of these desiderata, it is difficult to settle on a concept of knowledge which is not too remote from our initial intuitions. I suggest that epistemology may be better off without the concept of knowledge. There are plenty of other interesting notions for us to be going on with.

1. Introduction

Epistemologists live in a post-Gettier age. This age is characterised by a voluminous literature spawned by a search for the necessary and sufficient conditions for propositional knowledge. Why are such conditions so difficult to discover or devise? Why don't most epistemologists agree on roughly the same analysis as they appear to have done in the pre-Gettier paradise? It is not that they have widely differing concepts of knowledge as some have suggested. At the intuitive level, most epistemologists agree (roughly) as

to whether or not a given item is an item of knowledge (Whether those not engaged in the search agree with those intuitions is another matter) I shall argue that the problem lies in that fact that the epistemologists' intuitive concept of knowledge appeals to desiderata that probably cannot be satisfied Unfortunately, if we abandon some of these desiderata, it is difficult to settle on a concept of knowledge which is not too remote from our initial intuitions I suggest that epistemology may be better off without the concept of knowledge There are plenty of other interesting notions for us to be going on with

2. 'Interesting' Kinds of Belief

Items of (propositional) knowledge constitute a special kind of belief¹ It has been claimed that belief is not necessary for knowledge² I disagree This dispute has been well aired elsewhere, and I have nothing further to say on the matter here³ I take it as a basic assumption that belief is necessary for knowledge

Not just any belief will do as an item of knowledge Knowledge is a particularly 'worthy' or 'interesting' or 'valuable' kind of belief Or so it seems to epistemologists Our shared intuitions distinguish items of knowledge from items of non-knowledge, and so we assume that there is a difference between knowledge and non-knowledge And that difference is important to us There are beliefs which constitute knowledge and there are all our other beliefs We think it best to maximise the former and minimise the latter

What is it about knowledge which makes the difference? A good analysis of knowledge should explain this to us We hope that by setting out the necessary and sufficient conditions for knowledge we will make clear what it is

that makes knowledge such an interesting or valuable kind of belief

Let us adopt the term 'interesting' as a technical term to apply to any set of beliefs which seems naturally to form an epistemically interesting, coherent, and worthy kind. I shall examine various of these interesting kinds and assess them as putative candidates for knowledge.

3. True Beliefs

One of the ways in which we sort beliefs into different kinds is to distinguish between true beliefs and false beliefs. We are rational beings who prefer to have true beliefs rather than false beliefs. I believe that true beliefs have intrinsic value and false beliefs have intrinsic disvalue. This claim that true beliefs are intrinsically valuable is sometimes disputed.⁴ 'Why should I have true beliefs?' or 'What's so great about truth?' are questions which seem to make perfectly good sense. But whether or not true beliefs have intrinsic value, it is highly plausible that they have instrumental value. On the whole, true beliefs are necessary for our survival and assist in making our lives more comfortable. Acquiring true beliefs may not be the only game in town, but it is certainly major league.⁵

So true beliefs form an interesting kind. But true belief is not enough for knowledge. I may come to believe that it is raining in Eketahuna as a consequence of looking at the tea-leaves in the bottom of my cup after breakfast in Invercargill, but I would not then know that it is raining in Eketahuna. True beliefs can be acquired as a matter of chance or luck, but knowledge cannot. Actually, this is not strictly true. Suppose I glance out the window and see Moana walking past. As a result, I come to believe that Moana is now in Invercargill. If that was the only time I

glanced out the window all morning, and it coincided with the only time Moana passed by, then my acquiring that knowledge was a very chancy matter. But although I may have been lucky to have acquired a true belief in those circumstances, it was not a matter of luck that the belief I did acquire was true. We must distinguish between cases where luck plays a role as to whether or not the belief is acquired and those cases where luck plays a role as to whether or not the belief is true. I shall refer to the latter as cases where the belief is *fortuitously* true.

4 Non-fortuitously True Beliefs

Being a true belief is not sufficient for being an item of knowledge. Perhaps what we are looking for is the set of true beliefs which are not fortuitously true. Of course, without an analysis of what it is for a belief to be non-fortuitously true, we still do not have an explanatory analysis of knowledge. But if the suggestion is correct, then such an analysis of knowledge would just be the set of necessary and sufficient conditions for a belief to be non-fortuitously true. At the very least, this suggestion may be pointing us in the right direction.

But there are counterexamples to the suggestion. Suppose someone who knows about my eccentric belief-producing procedures also knows that it is raining in Eketahuna and surreptitiously rearranges the leaves in my cup so that I form the belief that it is raining in Eketahuna. My belief is not fortuitously true but I still do not know that it is raining in Eketahuna. Or so say the intuitions of most epistemologists.

There is some controversy here. Some externalists claim that if the belief acquired is non-fortuitously true (perhaps because it was caused by a reliable process or

caused by the fact believed) then according to their intuitions this is sufficient for knowledge. Thus, I can be tricked into knowing that it is raining in Eketahuna even in circumstances in which I am not justified in so believing and in which the process by which I come to believe is quite bizarre. These are minority intuitions. But I return to this issue later.

5. Epistemically Valuable Beliefs

Epistemic justification covers a whole raft of different notions. It is controversial whether only one of those notions is the only true notion of justification or whether 'justification' is a generic term which legitimately covers a number of different notions or whether it is an incoherent notion which appeals to conflicting desiderata.⁶ I wish to avoid this controversy. It is not relevant to the point I wish to make. When a belief is described as being justified (for the believer) then some epistemic value is being ascribed to that belief. Either the believer has good grounds available for the belief, or the believer can give good grounds for the belief, or the believer has not violated any epistemic duties in acquiring it, or the believer has acquired it via a reliable method, or the belief coheres with the believer's other beliefs, and so on. The point is that an ascription of epistemic justification (whatever one's notion of justification) is an ascription of epistemic value. The notion of epistemic value is wider than that of justification. We may accept that some property is epistemically valuable even though we would deny that it provides justification. For example, many (but by no means all) epistemologists would agree that being caused by a reliable process is epistemically valuable but deny that this provides justification (unless the believer is aware that it has been so caused). There are epis-

temic values to which no-one would be likely to attach the label of justification. Even if we restrict our attention to epistemic values which are truth-conducive, this may still be the case.

By truth-conducive, I mean any value which makes a belief more likely (in some sense) to be true. There are other properties which a belief may have which we may value. Certain kinds of belief may bring comfort or pleasure to the believer but such properties need not correlate highly with truth. But when considering values which are relevant to the notion of knowledge it is only the truth-conducive values which are of interest. Therefore, I restrict my use of the term 'epistemic value' to truth-conducive properties which a belief may have.

Another restriction is that epistemic value applies to beliefs *qua* mental states, not beliefs *qua* belief-contents. Generally we think of beliefs as a certain kind of mental state. They belong to the class of propositional attitudes along with desires, fears, doubts, suppositions, etc. As such, they have propositional content, that is, the content of each belief-state is a proposition (whatever that may amount to).⁷ Often we use the term 'belief' to refer simply to the content of a belief-state. When we say that two people have the same belief, we usually mean that they have beliefs with the same belief-content (i.e. they believe the same thing). As mental states, their beliefs may have widely differing properties. One may be rational, the other irrational, one certain, the other tentative, one obtained by inference, the other by authority, and so on. On the other hand, their contents will have identical properties, and these will be semantic properties. They will have the same truth value, meaning, and reference.

Epistemic values are not semantic properties. It follows that for any epistemic value it is possible to have false

beliefs with that value, no matter how high the degree of value. It is sometimes claimed that it is impossible for a self-evident belief to be false. But this claim arises from an equivocation on whether self-evidence is a psychological property (one that applies to a belief-state) or a semantic property (one that applies to a belief-content). If it is a psychological property then its applying will not guarantee truth and if it is a semantic property then its applying will not guarantee any particular psychological attitude to the belief.

For any given epistemic value, the set of beliefs possessing that value to a sufficiently high degree forms an interesting kind of belief. The more worthy and interesting the epistemic value, the more interesting the kind of belief will be.

Which values are most worthy is a matter of dispute. The major dispute is between internalists and externalists. Internalists put great importance on the ability of the believer to have access to the epistemic value in question. We can assess whether or not our beliefs have good grounds (so long as what we mean by 'grounds' is restricted to those things to which we have cognitive access), or whether they cohere with our other beliefs, or whether they have been validly inferred, and so on. Externalists, on the other hand, do not restrict themselves to what the believer can be cognitively acquainted with. We cannot always be aware of how reliable the method employed to acquire a belief may be. Nor can we always be aware of what the proper function of our cognitive faculties may be and thus whether a belief was acquired by properly functioning cognitive faculties. And so on.

But the problems I wish to draw attention to obtain irrespective of what one takes to be the important epistemic value or values. For 'epistemic value' you may plug in

any epistemic value (internal or external) which you consider to be necessary for knowledge

For any particular epistemic value, the set of epistemically valuable beliefs form an interesting kind. But some of the beliefs in the set may be false. For this reason, we reject any such kind as constituting knowledge. A false belief cannot be knowledge, no matter how epistemically worthy that belief. A natural suggestion is that it is the epistemically valuable beliefs which are also true which constitute the items of knowledge. In other words, the suggestion is that the set of items of knowledge is the set of epistemically valuable, true beliefs. At this point, we must consider Gettier.

6. The Gettier Problem

In his famous 1963 paper, Edmund Gettier presents two counterexamples to what he took to be the traditional analysis of knowledge. According to this analysis

S knows that p iff (a) p is true,
 (b) S believes that p , and
 (c) S is justified in believing
 that p

Gettier's counterexamples show that although each of these conditions may be necessary for knowledge, they are not jointly sufficient. In other words, it is possible for someone to have a justified true belief and yet not have an item of knowledge. An important proviso is that, although it matters little what we assume is required for a belief to be justified, it must be possible for a justified belief to be false.

Gettier's examples were a little more complicated than was required to make the point. Here is a simple but

typical example which does the job. Suppose I often know what the weather is like in Eketahuna. I acquire this knowledge because I have an aunt who lives in Eketahuna and she tells me what the weather is like when I ring her up. Suppose one day she tells me that it is raining, but she tells me this, not because she has seen the rain, but because she has heard what sounds like rain on her window. In fact, what she hears is water from her neighbour's garden sprinkler splashing against the window. However, it is, indeed, raining in Eketahuna. So, it is true that it is raining in Eketahuna, I believe it, and I am justified in believing it. All three conditions of the traditional analysis of knowledge have been met, and yet, our intuitions (or at least, the intuitions of most epistemologists) tell us that I do not know that it is raining in Eketahuna. They tell us that because in this case my belief is fortuitously true.

7. The Gettier Lesson

What do the Gettier counterexamples tell us? What lessons can we learn from them? They teach us that it is not just epistemically valueless beliefs (such as beliefs based on reading tea-leaves) which can be fortuitously true. It is possible for an epistemically valuable belief to be false. With this possibility comes another: the possibility that in circumstances similar to those which would result in a false belief, that belief could be fortuitously true. That is what happens in the Gettier counterexamples. And this directs our attention to what I believe is the important lesson to be drawn from those counterexamples. Any condition on our beliefs which does not guarantee their truth will not provide a condition for knowledge which is sufficient to accommodate our intuition that knowledge cannot be a fortuitously true belief.

External epistemic values, such as reliability of the process of belief acquisition, are just as subject to the Gettier problem as internal values. A reliable process cannot guarantee truth, unless it is 100% reliable. It is unlikely that we have any processes available to us which are 100% reliable. If knowledge could only be acquired via such processes then we would have far less knowledge than we have according to our intuitions. Telephoning my aunt is a reliable way of acquiring true beliefs about the weather in Eketahuna, but not all true beliefs acquired by that method are items of knowledge. It might be objected that a process which involves my aunt listening to sounds on the window, rather than looking out the window, is an unreliable process. But I do not see why such a process should not be highly reliable, and yet occasionally lead to a false belief and even less often a fortuitously true belief.

One way of identifying analyses of knowledge which can be 'gettierized' in ways similar to that described is that the truth condition must be included as a separate condition. This gives the set of beliefs which are both epistemically valuable and true a certain lack of coherence. It is as though the truth condition is tacked on to ensure that no false beliefs can count as knowledge. The set of true beliefs form an interesting kind for one reason, the set of epistemically valuable beliefs form an interesting kind for another, albeit related, reason, but there are members of their intersection which lack a connection between their truth and their truth-conducive property. We must turn to examining sets of beliefs which are distinguished by a property which guarantees their truth.

8. Internalist Guarantees of Truth

Properties of beliefs which purportedly guarantee truth might be internalist or externalist. The classic internalist account is Cartesian foundationalism. On this account the beliefs which constitute knowledge are those which are self-evident or those that have been deduced from self-evident beliefs. I have already argued that claims that self-evidence ensures truth trade on an equivocation between the property's being psychological or semantic. Even if I am wrong, there are other well-known problems with this account that I shall not detail here. It is doubtful that any of our beliefs are self-evident in the sense required. Even if we can have self-evident beliefs, those that are usually proffered as examples are not the sort of contingent facts about the external world of which we suppose that we can have knowledge nor do they logically imply such facts. I reject all such internalist accounts and consider externalist ones.

9. Causal Connections

A classic attempt to specify a connection between the truth of a proposition and the belief in that proposition is by means of a causal condition.⁸ In simple terms, the beliefs that p that are caused by the fact that p form an interesting kind. All beliefs of this kind will be true and none of them will be fortuitously true. But this interesting set of beliefs cannot be the set of all items of knowledge. For example, when I am sitting in front of the fire in my sitting room I know that smoke is rising from the chimney, but the fact that the smoke is rising is not a cause of my belief that the smoke is rising. However, that fact and my belief do have a common cause, namely the fire. Although the fact does not cause the belief, they are *causally connected*. They are

causally connected via a common cause. We, therefore, consider the interesting kind which is the set of all beliefs that are causally connected to their facts.

This interesting kind of belief yields a causal account of knowledge according to which *S* knows that *p* if and only if *S*'s belief that *p* is causally connected to the fact that *p*. As a core notion this causal theory has strong intuitive appeal. It is particularly attractive to those who seek a naturalized epistemology. The causal relation is a genuine relation which connects much of what occurs in the world. But the theory has been widely criticised. Some of the objections have been readily dealt with, others can be accommodated by adjustments to the theory. Space precludes dealing with all objections. I can only mention those I believe to be most troublesome. If I have omitted your favourite objection, it may be because I do not consider it insurmountable. On the other hand, it may be because I have overlooked it or it has never come to my attention, but I hope not.

Conceptual explications can be criticised because the conditions suggested are insufficient or because at least one of them is not necessary. First, we consider the charge that a causal connection is not sufficient for knowledge. Stories can be told in which, although there is a causal connection between fact and belief, our intuitions tell us that it is not a case of knowledge. There are familiar examples of knowledge being denied by the existence of relevant alternatives in the vicinity. For example, Jim sees Judy and forms the true belief that he has seen Judy. But it could just as easily have been her twin sister Trudy. If it had been, he would still have believed it was Judy, so Jim does not know.⁹ There are also examples where knowledge is denied by the fact that the causal connection is bizarre or deviant in some way. For example, there is the brain

tumour which has the effect of causing the sufferer to believe (in the absence of any other symptoms) that she has a brain tumour¹⁰ Intuitions concerning some of these cases are by no means unanimous, but if we wish to go along with the intuitions of the majority we shall have to find a way of ruling out such cases The causal connection will have to be an 'appropriate' one, and an account will have to be given as to what constitutes an appropriate connection

A promising approach is to use the notion of proper function Roughly, the proviso is that a causal connection can only provide knowledge if it is part of a process one of whose proper functions is to provide true belief, and that process must be functioning properly and reliably in an appropriate environment¹¹ A process can acquire a proper function either by intentional design or by natural selection

Insufficiency objections are a serious problem for a causal connection analysis of knowledge Even if we can give an account of 'appropriateness' for causal connections which conforms with intuition, it is likely to be a rather messy account But we can put this problem to one side since objections that a causal connection is not necessary for knowledge turn out to be even more serious

Once again, many counterexamples have been devised Many of them can be overcome by relaxing our notion of what constitutes a causal connection We may need to regard inference as a causal process whereby beliefs give rise to further beliefs, and to allow causal overdeterminants to count as causes (And let us not worry about what is going on causally in quantum situations until there is less controversy as to how the results and theories should be interpreted)

It is knowledge of universal generalisations which gives rise to the most obdurate problem for a causal account and which is the problem most often raised. The argument is that we can know that all emeralds are green and yet we are not causally connected to every emerald, so it cannot be the case that our belief is causally connected to the fact that all emeralds are green. But it has been suggested that we *can* be causally connected to the fact that all emeralds are green.¹² The suggestion goes as follows. It is reasonable to suppose that there is some fact that is the cause of all emeralds' being green. This fact (whatever it is) will also be the cause of the limited number of emeralds with which we have causal contact's being green and this is the cause of our belief that all emeralds are green. So our belief that all emeralds are green and the fact that all emeralds are green are causally connected via a common cause, the fact (whatever it is) that causes all emeralds to be green.

This account may not sound very convincing but suppose we can make it plausible for cases like the greenness of all emeralds, the blackness of all crows, etc. Such accounts depend on there being some fact that is the cause of the universal fact. Unfortunately, it is reasonable to suppose that there are universal facts which do not have causes. Why should there not be universal facts which are fundamental 'brute' facts, true because that is the way the universe is and always has been? Consider the fact that all protons have a rest mass of 1.672614×10^{-24} grams. If protons have always existed and their rest mass is not explicable in terms of more fundamental particles, then this is an example of a universal fact that does not have a cause. Hence, it is a fact that cannot meet the causal-connection condition.

It is possible that there are no such facts. If there is a god that created the world and we can have appropriate causal contact with the singular facts about this god which were causally responsible for the world being the way it is, then there will be no universal facts with which we cannot have causal contact. The same might be true if the world was caused by an uncaused big bang or if the 'Russian doll' model of the world is true.¹³ But these are contingent matters and a theory of knowledge is beginning to look very dodgy if it relies on the truth of such scenarios.

One response to this objection could be to point out that if the causal-connection condition denies knowledge of fundamental universal facts, then, although this yields scepticism about such facts, this is an acceptable scepticism that we can live with. So we cannot know the brute universal truths about the world? Why should we think that we ever could uncover and come to know such truths? As long as they are the only sort of facts we cannot know, then the problem is not serious. Unfortunately, it is unclear how far this scepticism spreads. Many of our inferences seem to depend on universal facts. If those facts are facts that cannot be known, then nor can the conclusions we draw from such inferences. Or so it would seem.

Enough has been said to suggest that, although the set of all beliefs that are causally connected to their facts constitutes an interesting kind of belief, it does not constitute the set of all items of knowledge. There are other kinds of connection between fact and belief to be considered.

10. The Truth-Tracking Connection

The fact that p and S 's belief that p may be connected according to the following subjunctive conditionals

- (1) If p were not true, then S would not believe that p
- (2) If p were true, then S would believe that p ¹⁴

It is claimed that any belief which meets those conditions 'tracks the truth'. By this, it is meant that (on a possible-worlds analysis of subjunctive conditionals¹⁵) in the possible worlds close to the actual world, p and S believes that p have the same truth value. Or, putting it another way, in all situations sufficiently similar to the actual situation, if p is true then S believes that p and if p is false then S does not believe that p . Since the actual world is itself a close possible world and S actually believes that p , it follows that p is actually true.

Thus, the truth-tracking conditions appear to pick out an interesting kind of belief which cannot be gettierized ¹⁶ Do they then pick out the items of knowledge? Robert Nozick claims that they do. Unfortunately, they fare little better than the causal connection. First, they are not sufficient to rule out non-knowledge when the belief arises from a deviant causal process. Consider the brain tumour which causes the belief that one has a brain tumour, and suppose it does this in a law-like way. If I have the tumour, then I believe that I do. If I do not, then I do not so believe. But as with the causal connection, perhaps some further condition can be added to disallow such deviant processes.

More serious is the fact that the conditions are too strong. Suppose I turn on the heater and leave the room. I know that the heater is on in the next room. But if the heater were not on, I would still believe that it was. In terms of the possible worlds analysis of subjunctive conditionals, I assume that the closest world in which the heater is not on is one in which an overload switch has been trig-

gered or a wire has become disconnected or my wife has switched it off. An adherent of the truth-tracking condition could respond by arguing as follows. Either the world I describe is very close, in which case I do not have knowledge because I cannot have knowledge about a heater which is on the verge of being disconnected, or the world I describe is very distant, in which case there are closer worlds in which I am not oblivious to the fact that the heater is likely to be disconnected, and in those worlds I would not believe that the heater is on.

What are the relevant close worlds? I claim that my reading of the conditional is a perfectly natural one. It is easy to suspect that our respondent is choosing worlds to fit the theory rather than the other way around. This points to a more general problem for the truth-tracking account of knowledge. The truth values of subjunctive conditionals are simply too indeterminate and too context-dependent for such an account to bear the weight of the task demanded of it, that of distinguishing knowledge from non-knowledge. Our intuitions as to what counts as knowledge may be rough and ready but our intuitions as to the truth values of subjunctive conditionals are even rougher and much less ready to make fine distinctions.

However, the context-dependence of subjunctive conditionals may be seen as a virtue of the truth-tracking account because knowledge ascriptions are themselves context-dependent. That knowledge ascriptions are to some extent context-dependent is a reasonable claim. In some contexts (for example, criminal trials or drug-testing), we set a more stringent standard for knowledge than in other contexts. But if the truth-tracking analysis is correct, then variation in ascription will not just vary along a stringency scale. Much will depend on what is in the mind of the ascriber and this can vary in multiple different ways. That

being the case, knowledge will not be an 'interesting' kind of belief, in my sense of interesting, and we should not expect to produce a neat set of necessary and sufficient conditions for it. On the other hand, given a fair degree of agreement in actual cases, we may doubt that knowledge is context-dependent to such an extent.

Another familiar objection, which I shall not detail here, is that on the truth-tracking account, knowledge is not closed under known implication. Although Nozick takes this to be a virtue, the resultant absurdities effectively constitute a *reductio* of his position.¹⁷ I conclude that the truth-tracking conditions are either too strong or they are not clear enough to make the required distinction.

11. The Explanatory Connection

Alan Goldman suggests that the important condition for knowledge is that the fact that p must explain, and render (significantly) more probable than it otherwise would be, the belief that p .¹⁸ He further suggests that a necessary condition for this connection to hold is

$$\text{pr}(S \text{ believes that } p / p) \gg \text{pr}(S \text{ believes that } p)$$

The idea is that in the usual case of my phoning my aunt in Eketahuna, the fact that it is raining will play a significant role in an explanation of my belief that it is raining. Furthermore, in the circumstances, I am much more likely to believe that it is raining if it is raining than I am to believe it regardless of the weather. However, in the 'gettierized' version, the fact that it is raining does not play a role in explaining my belief, and the fact that it is raining does not, in the circumstances, increase the likelihood of my having that belief.

One kind of possible counterexample to this account can be readily dealt with. When I am looking at the fire, the fact that the smoke is rising from the chimney does not explain, or make more probable, my belief that it is and yet I know that the smoke is rising. But we can patch up the theory along the same lines as for the causal theory, that is, we can weaken the requirement so that the fact and belief may be connected via some other fact which explains both the fact and my belief.

The main problem with this condition is the difficulty in determining if it has been met. Does the fact that all protons have a rest mass of 1.672614×10^{-24} grams explain my belief that they do, any better (or to any greater extent) than the fact that all observed protons have such a mass and the fact that I assume an appropriate inductive principle? Does the universal fact make my belief any more probable than it would otherwise be? It certainly does not make it significantly more probable than it would be if a few distant (or even a few close but unmeasured) protons had a different mass.

Goldman suggests that the relevant probabilities can be calculated by ranging over close possible worlds. Accordingly, the $\text{pr}(A)$ is the proportion of close possible worlds in which A is true, while the $\text{pr}(A/B)$ is the proportion of close possible worlds in which both A and B are true compared to all those in which B is true. That only close worlds are relevant he sees as an advantage, particularly in comparison to the truth-tracking condition. He gives the following example. A father is playing tennis with his young son when he hears from a radio broadcast that an assassination has occurred at a distant location. Thus he knows that his son is not the assassin. But suppose that if his son had been guilty, the father would still not have believed that he was. His love for his son would, in the cir-

cumstances, drive him to distraction. The truth-tracking condition is not met but, according to Goldman, the explanatory condition is because facts that explain the son's innocence also explain the father's belief in his innocence. What is puzzling is that the probability condition does not appear to have been met. In all close possible worlds the son will be innocent. After all, he is only ten years old, a great distance from the scene of the crime, and lacks the motivation, temperament, and skills of an assassin. It follows that on Goldman's method of calculating the probabilities, the probability that the father believes that his son is innocent given that he is will be identical to the antecedent probability that he so believes.

A similar problem arises if we suppose that the belief concerns a law of nature. If we make the reasonable assumption that close worlds are those with the same laws of nature as the actual world, then $\text{pr}(S \text{ believes that } p / p)$ will equal $\text{pr}(S \text{ believes that } p)$. Defenders of the proposal will no doubt object that we must obviously range over possible worlds in which p is false, but then we are in danger of choosing our possible worlds in order that our calculations meet our intuitions.

Perhaps Goldman's analysis of probability or the probability condition itself are a mistake. It does not follow that his explanatory condition is also mistaken. But what counts as a good explanation? I incline to the view that the only genuine explanations of singular facts are causal explanations. That S believes that p is a singular fact, so it follows that all explanations of beliefs are causal explanations. If I am right, then the only way a fact can play a significant role in an explanation of a belief in that fact is if the fact is causally connected to the belief.¹⁹ I suggest that it is only when such a causal connection exists that our intuitions that the fact explains the belief or renders it more

probable are firm enough to match our intuitions as to the belief's being an item of knowledge. In other words, a clearly explicated explanatory condition will reduce to a causal condition.

As with the truth-tracking analysis, we might suppose that these problems arise because knowledge is highly context-dependent. If so, my previous remarks in this regard will apply. My conclusion is the same as that for the truth-tracking account. Either the explanatory condition is too strong or it is too unclear.

12. The Nondefeasibility Condition

Nondefeasibility accounts of knowledge provide another kind of condition on beliefs that entails their being true. A simple version requires that

S knows that p only if there is no truth q such that if S believed that q , then S would no longer be justified in believing that p .²⁰

Notice that a subjunctive conditional has already entered the picture. This condition is much too strong. It can be readily shown that there is always a truth that will do the job of rendering the belief unjustified.²¹ Nondefeasibility accounts rapidly become complex. This alone is a good reason to doubt that such an account captures our intuitions as to what counts as knowledge or that it will explicate our concept. The complexities usually involve counterfactuals in which we suppose that S has acquired all (relevant) true beliefs and then we must decide if she still has her original justification. Once again, our intuitions as to the truth of such subjunctive conditionals will not bear the weight put on them.

I conclude that the truth-tracking, explanatory, and nondefeasibility accounts are only coherent enough to provide an interesting kind of true belief when they are too strong to constitute a necessary condition for our intuitive concept of knowledge. One result of tightening those accounts is that they will tend to be virtually indistinguishable from a causal account.

13. Conclusion

I have argued that putative accounts of knowledge can be divided into two broad types. First, there are the epistemically valuable, true beliefs, where epistemic value is conducive to a belief's being true but does not guarantee its truth. This notion is too weak to be that of knowledge because it cannot rule out the possibility of a belief's being fortuitously true. Secondly, there are the beliefs that are appropriately connected to their facts. This notion is too strong to be that of knowledge. The only plausible connection is a causal connection and there are facts which do not seem to have any causal connection with our beliefs in those facts.

It might be suggested that we could live with a concept of knowledge which allowed for the possibility of fortuitous truth as long as that possibility were extremely remote. There are two problems with this suggestion. One is that in order to lower the probability of fortuitous truth, we must lower the probability of having an epistemically valuable, false belief. In other words, we must raise our standard for a belief to have sufficient epistemic value for knowledge. If that standard is too high, then too few of our beliefs will qualify as knowledge. This will conflict with our intuition that we do, in fact, know lots of things. The other problem is that if we allow some fortuitously true be-

iefs to count as knowledge, on the grounds that they are extremely unlikely, then it is not clear why we should not allow others. We are venturing on a slippery slope at the bottom of which is the view that any true belief counts as knowledge. And that conflicts with our intuition that there is more to knowledge than true belief.

Our intuitions as to what constitutes knowledge are too demanding. Philosophers can of course try legislating in such cases. 'Now that I have clarified the issue I suggest that analysis X is what we *ought* to mean by knowledge.' But what analysis are we to suggest? If we include the strong connection requirement then we have what amounts to scepticism. On the other hand, epistemically valuable, true beliefs just do not form an interesting kind as long as they include the fortuitously true. So they are not interesting enough for such an important label. True beliefs do form an interesting kind, as do epistemically valuable beliefs, but these kinds are too remote from our intuitions to be imposed by legislation.²²

Many before me have concluded that we should abandon the search for a neat set of necessary and sufficient conditions for knowledge. The increasingly baroque nature of the accounts proposed in the post-Gettier flood provided good inductive grounds for their conclusion. I have demonstrated why the search is unlikely to be successful. My argument allows us to abandon the search with a philosophically clear conscience.

I have identified three kinds of beliefs which should be of interest to epistemologists. There are the true beliefs, the justified or rationally-held beliefs, and the beliefs which are appropriately connected to their facts. The beliefs which are thought to be items of knowledge may be safely ignored. In other words, I advocate epistemology without knowledge.

Bibliography

- Alston, W P 1989 "Concepts of Epistemic Justification" In his *Epistemic Justification* Ithaca Cornell University Press
- Armstrong, D M 1969-70 "Does Knowledge Entail Belief?" *Proceedings of the Aristotelian Society* 70 21-36
- Brown, J R 1990 " π in the Sky" In A D Irvine (ed) *Physicalism in Mathematics* Dordrecht Kluwer Academic Publishers, pp 95-120
- Feyerabend, P 1978 *Science in a Free Society* London NLB
- Gettier, E L 1963 "Is Justified True Belief Knowledge?" *Analysis* 23 121-3
- Goldman, A H 1988 *Empirical Knowledge* Berkeley University of California Press
- Goldman, A I 1967 "A Causal Theory of Knowing" *Journal of Philosophy* 64 357-72
- Goldman, A I 1976 "Discrimination and Perceptual Knowledge" *Journal of Philosophy* 73 77-91
- Harman, G 1973 *Thought* Princeton Princeton University Press
- Lehrer, K 1990 *Theory of Knowledge* London Routledge
- Lehrer, K & Paxson, T D 1978 "Knowledge Undefeated Justified True Belief" In Pappas & Swain (1978) 146-54
- Nozick, R 1981 *Philosophical Explanations* Cambridge, Mass Harvard University Press
- Pappas, G S & Swain, M (eds) 1978 *Essays on Knowledge and Justification* Ithaca Cornell University Press
- Plantinga, A 1990 "Justification in the 20th Century" *Philosophy and Phenomenological Research* 50 45-71
- Plantinga, A 1993 *Warrant and Proper Function* Oxford Oxford University Press
- Radford, C 1966 "Knowledge - By Examples" *Analysis* 27 1-11
- Ruben, D -H 1990 *Explaining Explanation* London Routledge

Sartwell, C 1992 "Why Knowledge is Merely True Belief" *Journal of Philosophy* 89 167–80

Sosa, E 1969 "Propositional Knowledge" *Philosophical Studies* 20 33–43

Keywords

Gettier, knowledge, belief

Colin Cheyne
Department of Philosophy
University of Otago
P O Box 56
Dunedin
New Zealand

Colin Cheyne@stonebow.otago.ac.nz

Notes

¹ For those uncomfortable with the reification of beliefs, we can say that knowing that *p* is a special category of believing that *p*

² For example, Radford (1966)

³ See Armstrong (1969–70) pp 21–36 and Lehrer (1990) pp 27–38

⁴ For example, Feyerabend (1978) p 131 fn

⁵ The 'truth game' is, of course, more complicated than this. The piling up of true beliefs by disjunction introduction has little value, as does the acquisition of true beliefs if it is at the expense of acquiring many more false beliefs. Roughly, the aim is to increase one's significant true beliefs, both proportionally and in total. I note in passing that those who raise questions about the value of true beliefs do appear to be seeking true answers to those questions, rather than merely useful ones.

⁶ See Alston (1989) and Plantinga (1990)

⁷ I shall not enter into a discussion on the ontological status of propositions

⁸ Goldman (1967)

⁹ Goldman (1976) p 778

¹⁰ Sosa (1969) p 39

¹¹ See Plantinga (1993)

¹² Brown (1990) pp 113–14

¹³ Every fact is caused by another fact and so on down *ad infinitum*

¹⁴ Nozick (1981) pp 172–78

¹⁵ Nozick discusses all examples in terms of this analysis. His account does not appear to fare any better under alternative analyses

¹⁶ In fact, the conditions need to be amended so as to fix the method of belief acquisition across the relevant counterfactual situations (*ibid* p 179), but my discussion does not depend on this point

¹⁷ Nozick (1981) pp 204–17

¹⁸ Goldman, A H (1988) pp 19–37

¹⁹ David-Hillel Ruben argues that non-causal relations (in particular identity and part-whole relations) may be cited in an explanation. See his (1990) pp 209–33. But in this context, causal theorists will no doubt accept that causal connections may be extended by such relations

²⁰ Adapted from Pappas & Swain (1978) pp 27–28. See also Lehrer & Paxson (1978)

²¹ See Harman (1973) p 152

²² Crispin Sartwell (1992) argues, for quite different reasons than mine, that the concept of knowledge is either merely true belief or it is incoherent. He assumes that the concept is coherent and therefore concludes that it must be merely true belief. My response is that knowledge is not merely true belief, so his conclusion should be that the concept of knowledge is incoherent.