

Defining Knowledge

Introduction to the Problem

Epistemology attempts to identify the properties by virtue of which a belief (or set of beliefs) is justified or counts as knowledge. One way of conceiving of this identification of properties is as a conceptual analysis of the concepts of knowledge and justified belief. The aim of a conceptual analysis is to provide a set of conditions, each of which is necessary for the truth of the concept and which collectively are sufficient. 'A is a necessary condition of B' means that B cannot be true unless A is true. 'A is a sufficient condition of B' means that if A is true, then B must be true too. The question 'What is knowledge?' asks for the necessary and sufficient conditions under which a person knows something.

To get a feeling for the kinds of problems connected with conceptual analysis, let's begin with a concept that is considerably simpler than that of knowledge.¹

Give a conceptual analysis of the concept of bachelor.

You might want to say that 'S is a bachelor if and only if (i) S is male, and (ii) S is unmarried'. Yet a little thought reveals that while conditions (i) and (ii) are necessary for being a bachelor, they are not sufficient for being a bachelor. For there are cases where S is an unmarried male but is not a bachelor. Suppose, for example, that S is a dog that is male but unmarried. So we need to add a further condition to our conceptual analysis. Suppose we also require that a bachelor be human: (iii) S is human. Are conditions (i), (ii), and (iii) jointly sufficient to identify all and

¹ The following example is taken from N. Everitt and A. Fisher, *Modern Epistemology* (McGraw-Hill, 1995), pp. 14–15.

only bachelors? One might object that an unmarried human male who is a newly born baby meets these conditions but still would not be considered a bachelor. So maybe we need a fourth condition: (iv) *S* is of marriageable age. But what counts as being of marriageable age will, of course, vary across societies. Is this a weakness of the analysis or does it only show that the concept of bachelor is culturally relative?

The above example illustrates that there are at least two kinds of problems connected with conceptual analysis. First, conceptual analysis depends on linguistic intuitions. Linguistic intuitions may differ between people and among societies. For example, people can be of different opinions on whether to call a monk or a baby a bachelor. Linguistic intuitions may also undergo change across time within the same individual. Second, when one tries to analyse complex theoretical concepts, it is often difficult to find a set of necessary and sufficient conditions which will accommodate all the various intuitions we have about the usage of the concept.

Traditionally, epistemologists have held that the concept of knowledge has three individually necessary and jointly sufficient conditions: justification, truth, and belief. Knowledge is said to be justified true belief. This view, which can be traced back to Plato's *Theaetetus*, claims that what distinguishes knowledge from mere true belief and lucky guessing is that it is based on some form of justification, evidence, or supporting reasons.

The traditional analysis of knowledge applies primarily to propositional knowledge (also referred to as 'knowing that' or 'factual knowledge'). Propositional knowledge takes the form '*S* knows that *p*' where *S* stands for a subject and *p* stands for a declarative sentence expressing some proposition. 'Proposition' is a technical term of philosophers used to denote what it is that is said or asserted by an utterance, a sentence, or a thought. A proposition is something like a meaning. It is important to distinguish the proposition from the event in which that proposition is expressed. Unlike propositions, events happen at a certain time and at a certain place. One must also distinguish a proposition from the words, sounds, symbols, or brain states that express a proposition. When, for example, you say '*S* is a bachelor' and when I say '*S* ist ein Jungeselle' we utter different sentences but express the same proposition. Because propositions are either true or false, only declarative sentences can express propositions. A question or a demand is neither true nor false. And there are even some declarative sentences (e.g., 'Picasso is the greatest painter', 'God is benevolent') where it is debatable whether they have clear truth values. So the class of propositional knowledge is relatively small.

Apart from propositional knowledge, philosophers recognize two main kinds of knowledge, namely practical knowledge (or 'knowing how') and non-propositional knowledge of people, places, and things (or 'knowledge by acquaintance'). Within propositional knowledge one can distinguish between inferential (or demonstrative) and non-inferential (or direct, basic, immediate) knowledge. Inferential knowledge is the product of suitable deductive or inductive inferences from other propositions that serve as evidence or justification. An example of inferential knowledge is when you come to know that today is Tuesday on the basis of observing your garbage being collected and knowing that your garbage is collected every Tuesday.

The truth condition of the justified-true-belief analysis of knowledge states that if you know that p , then p is true. You cannot, for example, know that Bert is a bachelor unless, in fact, Bert is a bachelor. But, of course, p can itself be a claim that something else is false. You can, for example, know that it is false that Bert is married. The belief condition claims that knowing that Bert is a bachelor implies believing that he is a bachelor. A person need not be absolutely certain that something is true in order to know that it is. The belief condition only requires some kind of acceptance in the interest of obtaining truth. Finally, the justification condition requires that a known proposition be evidentially supported. The justification condition is there in order to prevent lucky guesses from counting as knowledge when the guesser is sufficiently confident to believe his own guess. Epistemic justification for a belief is justification for the belief's truth, not its usefulness or its social respectability. According to the tradition, a belief can be justified in a way sufficient for knowledge even if the justification is not conclusive – even if it is the sort of justification one can have for a false proposition (though, of course, one could not know a false proposition to be true by having such a justification for it). Thus the necessity for an independent truth condition.

In 1963 the philosopher Edmund Gettier published a three-page paper challenging the traditional analysis of knowledge as justified true belief.² Gettier does not dispute the claim that belief, justification, and truth are necessary for knowledge, but rather he argues that these alone are not sufficient for knowledge. He presents two examples of justified true belief that are not cases of knowledge. One such example, which is less complex to describe than Gettier's own examples, is the following scenario from Keith Lehrer³: two agents, the ubiquitous S and Mr Nogot, work in the same office. Nogot has given S evidence that justifies S in believing that Nogot owns a Ford car. Imagine that S has seen Nogot driving a Ford, S has been told by persons who have in the past been reliable that Nogot owns a Ford, and so on. From the proposition that Nogot owns a Ford, S then infers the proposition p : 'Someone in the office owns a Ford'. But unsuspected by S , Nogot has been shamming and p is only true because another person in the office, Mr Havit, owns a Ford. Thus S has a true belief, namely the belief p that someone in the office owns a Ford. Moreover, S is justified in believing that p . But we wouldn't want to say in this case that S has knowledge. He has simply been lucky. It just so happened that someone in the office owns a Ford, but not the person S thinks owns a Ford. The kind of reasons S possesses don't seem to be of the right sort.

There is an abundance of Gettier-type examples in the literature. The common feature among all these cases is that the things that justify a subject in believing p are distinct from the things that make p true. What accounts for the justification of the belief is not what accounts for its truth. To see this, reconsider the above example: S clearly has reasons for believing that someone in the office owns a Ford. Yet if we were to explain why this belief is true, we would refer not to Nogot, but rather to

² E. Gettier, 'Is Justified True Belief Knowledge?', *Analysis* 23 (1963), pp. 121–3.

³ K. Lehrer, 'Knowledge, Truth and Evidence', *Analysis* 25 (1965), pp. 168–75.

Havit. So S's justification doesn't direct us to what accounts for the truth of his belief. There is a gap between the satisfaction of the justification condition and the satisfaction of the truth condition. With respect to S's justification, it is merely good fortune or a happy accident that the belief is true. For all he knows, it might have been false that anyone, including Havit, owns a Ford.

Construct your own Gettier-type counterexample to the justified-true-belief analysis of knowledge.

The response to Gettier's short paper, 'Is Justified True Belief Knowledge?', was overwhelming. Some suggested that Gettier-type counterexamples were somehow defective or overlooked some obvious point about justification; others accepted the counterexamples and amended the justified-true-belief analysis by proposing additional conditions on knowledge designed to block Gettier-type cases; still others, instead of adding conditions, suggested changes in how existing conditions (e.g., justification) were to be understood. The section by Pollock discusses the first two kinds of response. Nozick's solution to the Gettier problem belongs to the third group.

Introduction to Pollock

John L. Pollock is professor of philosophy at the University of Arizona at Tucson and works primarily in epistemology, philosophical logic, and artificial intelligence. He not only tries to understand how the mind works but also attempts to build a mind. Pollock directs the OSCAR project, whose purpose is to formulate a general theory of rationality and to implement it in an artificial rational agent. The computer system OSCAR seems already capable of intellectual feats surpassed only by human beings, and Pollock hopes that the time will come when even humans cannot compete with OSCAR on some kinds of intellectual tasks. He has written books about artificial intelligence with such intriguing titles as *How to Build a Person: A Prolegomenon* (1990) and *Cognitive Carpentry: A Blueprint for How to Build a Person* (1995).

'The Gettier Problem' is the introductory section of an appendix to Pollock's *Contemporary Theories of Knowledge* (1986). In this piece he explains the Gettier problem and critically discusses some of the responses to the Gettier problem found in the literature. In the later sections of the appendix, not reprinted in this volume, Pollock set forth his own solution to the Gettier problem. In 1999 Pollock (together with Joseph Cruz) published a second edition of *Contemporary Theories of Knowledge*. The second edition does not contain the appendix on the Gettier problem. Other epistemological books by Pollock are *Knowledge and Justification* (1974), *Subjunctive Reasoning* (1976), and *Nomic Probability and the Foundations of Induction* (1990).

John L. Pollock, ‘The Gettier Problem’

It is rare in philosophy to find a consensus on any substantive issue, but for some time there was almost complete consensus on what is called ‘the justified true belief analysis of knowing’. According to that analysis:

S knows p if and only if:

- (1) p is true;
- (2) S believes p ; and
- (3) S is justified in believing p .

In the period immediately preceding the publication of Gettier’s (1963) landmark article “Is justified true belief knowledge?”, this analysis was affirmed by virtually every writer in epistemology. Then Gettier published his article and single-handedly changed the course of epistemology. He did this by presenting two clear and undeniable counterexamples to the justified true belief analysis.

a → Recounting the example given in chapter one, consider Smith who believes falsely but with good reason that Jones owns a Ford. Smith has no idea where Brown is, but he arbitrarily picks Barcelona and infers from the putative fact that Jones owns a Ford that either Jones owns a Ford or Brown is in Barcelona. It happens by chance that Brown is in Barcelona, so this disjunction is true. Furthermore, as Smith has good reason to believe that Jones owns a Ford, he is justified in believing this disjunction. But as his evidence does not pertain to the true disjunct of the disjunction, we would not regard Smith as *knowing* that either Jones owns a Ford or Brown is in Barcelona.

Gettier’s paper was followed by a spate of articles attempting to meet his counterexamples by adding a fourth condition to the analysis of knowing.

b → The first attempts to solve the Gettier problem turned on the observation that in Gettier’s examples, the epistemic agent arrives at his justified true belief by reasoning from a false belief. That suggested the addition of a fourth condition something like the following:

S’s grounds for believing p do not include any false beliefs.¹

It soon emerged, however, that further counterexamples could be constructed in which knowledge is lacking despite the believer’s not inferring his belief from any false beliefs. Alvin Goldman (1976) constructed the following example. Suppose you are driving through the countryside and see what you take to be a barn. You see it in good light and from not too great a distance, it looks the way barns look, and so on. Furthermore, it is a barn. You then have justified true belief that it is a barn. But in an attempt to

c →

¹ See, for example, Michael Clark (1963).

appear more opulent than they are, the people around here have taken to constructing very realistic barn facades that cannot readily be distinguished from the real thing when viewed from the highway. There are many more barn facades around than real barns. Under these circumstances we would not agree that you know that what you see is a barn, even though you have justified true belief. Furthermore, your belief that you see a barn is not in any way inferred from a belief about the absence of barn facades. Most likely the possibility of barn facades is something that will not even have occurred to you, much less have played a role in your reasoning.

We can construct an even simpler perceptual example. Suppose S sees a ball that looks red to him, and on that basis he correctly judges that it is red. But unbeknownst to S, the ball is illuminated by red lights and would look red to him even if it were not red. Then S does not know that the ball is red despite his having a justified true belief to that effect. Furthermore, his reason for believing that the ball is red does not involve his believing that the ball is not illuminated by red lights. Illumination by red lights is related to his reasoning only as a defeater, not as a step in his reasoning. These examples, of other related examples,² indicate that justified true belief can fail to be knowledge because of the truth values of propositions that do not play a direct role in the reasoning underlying the belief. This observation led to a number of “defeasibility” analyses of knowing.³ The simplest defeasibility analysis would consist of adding a fourth condition requiring that there be no true defeaters. This might be accomplished as follows:

d →

There is no true proposition Q such that if Q were added to S’s beliefs then he would no longer be justified in believing *p*.⁴

But Keith Lehrer and Thomas Paxson (1969) presented the following counterexample to this simple proposal:

e →

Suppose I see a man walk into the library and remove a book from the library by concealing it beneath his coat. Since I am sure the man is Tom Grabit, whom I have often seen before when he attended my classes, I report that I know that Tom Grabit has removed the book. However, suppose further that Mrs. Grabit, the mother of Tom, has averred that on the day in question Tom was not in the library, indeed, was thousands of miles away, and that Tom’s identical twin brother, John Grabit, was in the library. Imagine, moreover, that I am entirely

² See, for example, Brian Skyrms (1967).

³ The first defeasibility analysis was that of Keith Lehrer (1965). That was followed by Lehrer and Thomas Paxson (1969), Peter Klein (1971), (1976), (1979), (1980), Lehrer (1974), (1979), Ernest Sosa (1974), (1980), and Marshall Swain (1981).

⁴ This is basically the analysis proffered by Klein (1971).

ignorant of the fact that Mrs. Grabit has said these things. The statement that she has said these things would defeat any justification I have for believing that Tom Grabit removed the book, according to our present definition of defeasibility. . . .

The preceding might seem acceptable until we finish the story by adding that Mrs. Grabit is a compulsive and pathological liar, that Tom Grabit is a fiction of her demented mind, and that Tom Grabit took the book as I believed. Once this is added, it should be apparent that I did know that Tom Grabit removed the book. (p. 228)

A natural proposal for handling the Grabit example is that in addition to there being a true defeater there is a true defeater defeater, and that restores knowledge. For example, in the Grabit case it is true that Mrs. Grabit reported that Tom was not in the library but his twin brother John was there (a defeater), but it is also true that Mrs. Grabit is a compulsive and pathological liar and John Grabit is a fiction of her demented mind (a defeater defeater). It is difficult, however, to construct a precise principle that handles these examples correctly by appealing to true defeaters and true defeater defeaters. It will not do to amend the above proposal as follows:

If there is a true proposition Q such that if Q were added to S's beliefs then he would no longer be justified in believing *p*, then there is also a true proposition R such that if Q and R were both added to S's beliefs then he would be justified in believing *p*.

The simplest difficulty for this proposal is that adding R may add new reasons for believing *p* rather than restoring the old reasons. It is not trivial to see how to formulate a fourth condition incorporating defeater defeaters. I think that such a fourth condition will ultimately provide the solution to the Gettier problem, but no proposal of this sort has been worked out in the literature.⁵

References

- Clark, Michael (1963) Knowledge and grounds: A comment on Mr. Gettier's paper. *Analysis* 24: 46–48.
- Gettier, Edmund (1963) Is justified true belief knowledge? *Analysis* 23: 121–23.
- Goldman, Alvin (1976) Discrimination and perceptual knowledge. *Journal of Philosophy* 73: 771–91.
- Klein, Peter (1971) A proposed definition of propositional knowledge. *Journal of Philosophy* 68: 471–82.
- Klein, Peter (1976) Knowledge, causality, and defeasibility. *Journal of Philosophy* 73: 792–812.

⁵ A good survey of the literature on the Gettier problem, going into much more detail than space permits here, can be found in Shope (1983).

- Klein, Peter (1979) Misleading “misleading defeaters”. *Journal of Philosophy* 76: 382–86.
- Klein, Peter (1980) Misleading evidence and the restoration of justification. *Philosophical Studies* 37: 81–89.
- Lehrer, Keith (1965) Knowledge, truth, and evidence. *Analysis* 25: 168–75.
- Lehrer, Keith (1974) *Knowledge*. Oxford: Oxford University Press.
- Lehrer, Keith (1979) The Gettier problem and the analysis of knowledge. In *Justification and Knowledge: New Studies in Epistemology*, ed. George Pappas, pp. 65–78. Dordrecht: Reidel.
- Lehrer, Keith, and Thomas Paxson (1969) Knowledge: Undefeated justified true belief. *Journal of Philosophy* 66: 225–37.
- Shope, Robert K. (1983) *The Analysis of Knowing*. Princeton: Princeton University Press.
- Sosa, Ernest, (1974) How do you know? *American Philosophical Quarterly* 11: 113–22.
- Sosa, Ernest, (1980) Epistemic presupposition. In *Justification and Knowledge: New Studies in Epistemology*, ed. George Pappas. Dordrecht: Reidel.
- Swain, Marshall (1981) *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.

Commentary on Pollock

In the paragraph marked [a]→ Pollock recounts one of the two original Gettier examples. The example is a cousin of the one discussed in the ‘Introduction to the Problem’: suppose that Smith is justified in believing that Jones owns a Ford car. On the basis of the belief that

Jones owns a Ford

Smith comes to believe that

either Jones owns a Ford or Brown is in Barcelona.

The inference from the first to the second belief relies on the deductive rule according to which if one is justified in believing a proposition one can disjoin it with any other proposition.¹ Now Smith simply picks the second disjunct (that Brown is in Barcelona) out of thin air. As it turns out, Jones does not own a Ford but Brown happens to be in Barcelona. So it is the first disjunct that Smith has reasons for believing but it is the second disjunct which renders the disjunctive belief true. (A disjunction is true if at least one of its disjuncts is true.) Even though Smith is justified in believing the true disjunction that either Jones owns a Ford or Brown is in Barcelona, we wouldn’t say that he knows it.

One thing to notice with this example is that it assumes that a person can be justified in believing a proposition that is in fact false. So even though a justified belief is more likely to be true than an unjustified belief, justification does not

¹ In ordinary discourse, the operator ‘or’ is sometimes used in an exclusive sense, to say that exactly one of the two disjoint sentences is true. In logic and epistemology, however, ‘or’ is always given an inclusive interpretation according to which at least one and possibly both of the two disjoined sentences are true.

guarantee truth. For this reason Smith can be justified in believing falsely that Jones owns a Ford.

Given that Smith has no evidence for believing that Brown is in Barcelona, why does Pollock (following Gettier) think that Smith is justified in believing that either Jones owns a Ford or Brown is in Barcelona?

Gettier defends this idea by appeal to the following principle: if S is justified in believing p , and p entails q , and S believes q because he deduces q from p , then S is justified in believing q . Some philosophers have challenged this principle.² Can you think of reasons for denying this principle?

As Pollock explains in the passage marked $\boxed{b} \rightarrow$, some philosophers think that Gettier-type counterexamples to the justified-true-belief analysis of knowledge can be ruled out by stipulating that the belief that p must not be based on a false premise. In the preceding example, the true belief that either Jones owns a Ford or Brown is in Barcelona is inferred from the false belief that Jones owns a Ford. So maybe one needs only to add a fourth condition to the justified-true-belief analysis, requiring that the justification does not depend on any false premise. The idea of the 'no-false-premise' approach is that to rule out Gettier examples one needs to stipulate that the justification may not depend on any false premises.

Pollock objects to the no-false-premise condition on the grounds that it is too weak to rule out Gettier-type counterexamples.

Modify the Gettier example discussed by Pollock in such a way that S 's belief to the effect that someone owns a Ford does not rely on a false premise.

Suppose S 's reason for believing that someone owns a Ford is that Nogot told S that he owns one. Let q stand for the proposition

(q) Nogot said that he owns a Ford.

From q Smith infers proposition p

(p) Someone in the office owns a Ford.

Proposition p is true, S believes that p is true and he has reasons for believing p , namely q . Yet S doesn't know that p because it is not Nogot who owns a Ford but Havit. The crucial difference between this example and the previous version is that here S does not rely on a false premise. For even though Nogot did not tell the truth when he said that he owns a Ford, he did say it. Proposition q is indeed true.³

² For example I. Thalberg, 'In Defense of Justified True Belief', *Journal of Philosophy* 66 (1969), pp. 794–803.

³ This example is a variation of an example given by R. Feldman in 'An Alleged Defect in Gettier Counter-Examples', *Australasian Journal of Philosophy* 52 (1974), pp. 68–9.

In the paragraph marked by [c]→ Pollock presents another Gettier-type example in which no false belief is present: *S* is driving in the country and correctly identifies a barn in the distance. Unbeknownst to *S*, the area is full of barn façades which, from a distance, cannot be distinguished from real barns. Pollock concludes that *S* cannot be said to know that he is seeing a barn even though he has a justified true belief. The truth of the belief is just a matter of luck. But since *S*'s failure to know is not attributed to any false premise on which the belief is based, the no-false-premise condition does not succeed in saving the justified-true-belief analysis.

Some epistemologists think that the no-false-premise condition is too strong for knowledge rather than too weak.

Develop an argument to the effect that the no-false-premise condition is too strong for knowledge.

Consider the following variation on the original Gettier example: Smith correctly believes that Jones owns a Ford. Smith's belief is based on his having seen Jones drive a Ford and having seen the registration papers. Suppose that Smith also holds the false belief that Jones bought his Ford in March 2005. However, Jones purchased the car in December 2004. Now Smith infers that someone in the office owns a Ford on the basis of his belief that Jones bought a Ford in March 2005. Since it doesn't matter in which month the car was purchased, we seem to be entitled to say that Smith knows that someone in the office owns a Ford. As long as the premise in question is not essential for the inferred belief, it is possible to base knowledge on a false premise; or so one may argue.

In [d]→ Pollock goes on to discuss the defeasibility analysis of knowledge. The basic idea behind this position is that sometimes justified true belief is not knowledge because the justification is incomplete in certain crucial respects. Not all of the important evidence has been taken into consideration. There are some facts which, if the person were to know them, would seriously weaken his justification for believing. Consider Gettier's example. Smith thinks that either Jones owns a Ford or Brown is in Barcelona (which is true) because he justifiably thinks that Jones owns a Ford (which is false). If he learned that Jones did not own a Ford, though, this would defeat his justification for thinking that either Jones owns a Ford or Brown is in Barcelona. So Smith, though he has a justification for truly thinking what he does, does not have an *undefeated* justification. There are truths which, if Smith learned them, would undermine his justification for thinking that either Jones owns a Ford or Brown is in Barcelona. Requiring undefeated justification for knowledge avoids Gettier's examples.

Explain why the library example discussed in the passage marked [e]→ indicates that the defeasibility analysis of knowledge stated in [d]→ is too strong.

On the defeasibility theory stated in [d]→, what prevents a subject's having knowledge is that a defeating proposition *q* is present, whether or not the subject is aware

of q . If the existence of defeaters alone could rob us of knowledge, then knowledge would become a rare commodity. Therefore, Lehrer and Paxson define knowledge as justified true belief that is undefeated by a proposition which the believer has a justification for believing false. Obviously you can only be justified in believing that a proposition is false if you are aware of this proposition.

And there is another reason for why the defeasibility analysis in $\boxed{d} \rightarrow$ is too strong. As it stands, even *misleading* evidence can undermine justification. To see this, reconsider the library example in $\boxed{e} \rightarrow$. The belief that

(p) Tom stole the book

appears to be defeated by the true proposition that

(q) Tom's mother testified that Tom was somewhere else at the time in question.

But does q really qualify as a defeater of p ? The characterization in $\boxed{d} \rightarrow$ seems to indicate that it does, for if S was made aware of q , he would no longer be justified in believing p . But once we learn the further fact that

(r) Tom's mother is a deranged liar and Tom was indeed at the scene of the crime

we see that we were right in our original inclination to regard S 's belief that p as knowledge. This suggests that a distinction must be drawn between genuine and misleading defeaters. Misleading defeaters such as q are defeaters that can themselves be defeated. Statement r is such a defeating defeater. The article ends with Pollock claiming that a distinction between genuine and misleading defeaters has a good chance of providing a solution to the Gettier problem. In the later sections of the appendix from which the article is taken, Pollock develops his own version of the indefeasibility theory.

Introduction to Nozick

Robert Nozick (1938–2002) was professor of philosophy at Harvard University. He is probably best known for his defence of individualism and free-market libertarianism. His most influential contributions to philosophy outside of political theory have been in epistemology and metaphysics. In *Philosophical Explanations* (1981) – from which the extract is taken – Nozick ranges over a host of fundamental issues, such as the analysis of knowledge, scepticism, free will, the identity of the self, the foundations of ethics and the meaning of life. In dealing with these topics, he rejects the idea of strict philosophical proof and instead adopts a notion of philosophical pluralism. He writes: 'There are various philosophical views, mutually incompatible, which cannot be dismissed or simply rejected. Philosophy's output is the basketful of these admissible views, all together' (p. 21). Nozick thinks that this basketful of views can be ordered

according to criteria of coherence and adequacy and that even low-ranked views can offer important insights.

The most famous piece in *Philosophical Explanations* is the ‘truth-tracking’ account of knowledge: one knows a truth if one believes it, would not believe it were it false, but would believe it were it true. This account is said to solve the Gettier problem and avoid scepticism. *The Possibility of Knowledge: Nozick and his Critics* (1987), edited by S. Luper-Foy, contains 12 essays assessing Nozick’s analysis of knowledge and evidence and his approach to scepticism.

Apart from *Philosophical Explanations*, Nozick wrote *Anarchy, State, and Utopia* (1974), *The Examined Life: Philosophical Meditations* (1989), *The Nature of Rationality* (1993), *Socratic Puzzles* (1997), and *Invariances: The Structure of the Objective World* (2001).

Robert Nozick, ‘Conditions for Knowledge’

Our task is to formulate further conditions to go alongside

- (1) p is true
- (2) S believes that p .

We would like each condition to be necessary for knowledge, so any case that fails to satisfy it will not be an instance of knowledge. Furthermore, we would like the conditions to be jointly sufficient for knowledge, so any case that satisfies all of them will be an instance of knowledge. We first shall formulate conditions that seem to handle ordinary cases correctly, classifying as knowledge cases which are knowledge, and as nonknowledge cases which are not; then we shall check to see how these conditions handle some difficult cases discussed in the literature.*

The causal condition on knowledge, previously mentioned, provides an inhospitable environment for mathematical and ethical knowledge; also there are well-known difficulties in specifying the type of causal connection.

[a] → If someone floating in a tank oblivious to everything around him is given (by direct electrical and chemical stimulation of the brain) the belief that he is floating in a tank with his brain being stimulated, then even though that fact is part of the cause of his belief, still he does not know that it is true.

Let us consider a different third condition:

- (3) If p weren’t true, S wouldn’t believe that p .

Throughout this work, let us write the subjunctive ‘if-then’ by an arrow, and the negation of a sentence by prefacing “not-” to it. The above condition thus is rewritten as:

- (3) $\text{not-}p \rightarrow \text{not-}(S \text{ believes that } p)$.

This subjunctive condition is not unrelated to the causal condition. Often when the fact that p (partially) causes someone to believe that p , the fact also will be causally necessary for his having the belief – without the cause, the effect would not occur. In that case, the subjunctive condition 3 also will be satisfied. Yet this condition is not equivalent to the causal condition. For the causal condition will be satisfied in cases of causal overdetermination, where either two sufficient causes of the effect actually operate, or a back-up cause (of the same effect) would operate if the first one didn't; whereas the subjunctive condition need not hold for these cases.¹ When the two conditions do agree, causality indicates knowledge because it acts in a manner that makes the subjunctive 3 true.

[b]→ The subjunctive condition 3 serves to exclude cases of the sort first described by Edward Gettier, such as the following. Two other people are in my office and I am justified on the basis of much evidence in believing the first owns a Ford car; though he (now) does not, the second person (a stranger to me) owns one. I believe truly and justifiably that someone (or other) in my office owns a Ford car, but I do not know someone does. Concluded Gettier, knowledge is not simply justified true belief.

The following subjunctive, which specifies condition 3 for this Gettier case, is not satisfied: if no one in my office owned a Ford car, I wouldn't believe that someone did. The situation that would obtain if no one in my office owned a Ford is one where the stranger does not (or where he is not in the office); and in that situation I still would believe, as before, that someone in my office does own a Ford, namely, the first person. So the subjunctive condition 3 excludes this Gettier case as a case of knowledge.

[c]→ The subjunctive condition is powerful and intuitive, not so easy to satisfy, yet not so powerful as to rule out everything as an instance of knowledge. A subjunctive conditional 'if p were true, q would be true', $p \rightarrow q$, does not say that p entails q or that it is logically impossible that p yet not- q . It says that in the situation that would obtain if p were true, q also would be true. This point is brought out especially clearly in recent 'possible-worlds' accounts of subjunctives: the subjunctive is true when (roughly) in all those worlds in which p holds true that are closest to the actual world, q also is true. (Examine those worlds in which p holds true closest to the actual world, and see if q holds true in all these.) Whether or not q is true in p worlds that are still farther away from the actual world is irrelevant to the truth of the subjunctive. I do not mean to endorse any particular possible-worlds account of subjunctives, nor am I committed to this type of account.²

¹ ... I should note here that I assume bivalence throughout this chapter, and consider only statements that are true if and only if their negations are false.

² See Robert Stalnaker, "A Theory of Conditionals", in N. Rescher, ed., *Studies in Logical Theory* (Basil Blackwell, Oxford, 1968); David Lewis, *Counterfactuals* (Harvard University Press, Cambridge, 1973); and Jonathan Bennett's critical review of Lewis, "Counterfactuals and Possible Worlds", *Canadian Journal of Philosophy*, Vol. IV, no. 2, Dec. 1974, pp. 381–402. . . .

I sometimes shall use it, though, when it illustrates points in an especially clear way.*

[d]→

The subjunctive condition 3 also handles nicely cases that cause difficulties for the view that you know that p when you can rule out the relevant alternatives to p in the context. For, as Gail Stine writes, “what makes an alternative relevant in one context and not another? . . . if on the basis of visual appearances obtained under optimum conditions while driving through the countryside Henry identifies an object as a barn, normally we say that Henry knows that it is a barn. Let us suppose, however, that unknown to Henry, the region is full of expertly made papier-mâché facsimiles of barns. In that case, we would not say that Henry knows that the object is a barn, unless he has evidence against it being a papier-mâché facsimile, which is now a relevant alternative. So much is clear, but what if no such facsimiles exist in Henry’s surroundings, although they once did? Are either of these circumstances sufficient to make the hypothesis (that it’s a papier-mâché object) relevant? Probably not, but the situation is not so clear.”³ Let p be the statement that the object in the field is a (real) barn, and q the one that the object in the field is a papier-mâché barn. When papier-mâché barns are scattered through the area, if p were false, q would be true or might be. Since in this case (we are supposing) the person still would believe p , the subjunctive

(3) $\text{not-}p \rightarrow \text{not-(S believes that } p)$

is not satisfied, and so he doesn’t know that p . However, when papier-mâché barns are or were scattered around another country, even if p were false q wouldn’t be true, and so (for all we have been told) the person may well know that p . A hypothesis q contrary to p clearly is relevant when if p weren’t true, q would be true; when $\text{not-}p \rightarrow q$. It clearly is irrelevant when if p weren’t true, q also would not be true; when $\text{not-}p \rightarrow \text{not-}q$. The remaining possibility is that neither of these opposed subjunctives holds; q might (or might not) be true if p weren’t true. In this case, q also will be relevant, according to an account of knowledge incorporating condition 3 and treating subjunctives along the lines sketched above. Thus, condition 3 handles cases that befuddle the ‘relevant alternatives’ account; though that account can adopt the above subjunctive criterion for when an alternative is relevant, it then becomes merely an alternate and longer way of stating condition 3.⁴

³ 9. G. C. Stine, “Skepticism, Relevant Alternatives and Deductive Closure”, *Philosophical Studies*, Vol. 29 (1976) p. 252, who attributes the example to Carl Ginet.

⁴ This last remark is a bit too brisk, for that account might use a subjunctive criterion for when an alternative q to p is relevant (namely, when if p were not to hold, q would or might), and utilize some further notion of what it is to rule out relevant alternatives (for example, have evidence against them), so that it did not turn out to be equivalent to the account we offer.

e → Despite the power and intuitive force of the condition that if p weren't true the person would not believe it, this condition does not (in conjunction with the first two conditions) rule out every problem case. There remains, for example, the case of the person in the tank who is brought to believe, by direct electrical and chemical stimulation of his brain, that he is in the tank and is being brought to believe things in this way; he does not know this is true. However, the subjunctive condition is satisfied: if he weren't floating in the tank, he wouldn't believe he was.

The person in the tank does not know he is there, because his belief is not sensitive to the truth. Although it is caused by the fact that is its content, it is not sensitive to that fact. The operators of the tank could have produced any belief, including the false belief that he wasn't in the tank; if they had, he would have believed that. Perfect sensitivity would involve beliefs and facts varying together. We already have one portion of that variation, subjunctively at least: if p were false he wouldn't believe it. This sensitivity as specified by a subjunctive does not have the belief vary with the truth or falsity of p in all possible situations, merely in the ones that would or might obtain if p were false.

The subjunctive condition

$$(3) \quad \text{not-}p \rightarrow \text{not-(S believes that } p)$$

tells us only half the story about how his belief is sensitive to the truth-value of p . It tells us how his belief state is sensitive to p 's falsity, but not how it is sensitive to p 's truth; it tells us what his belief state would be if p were false, but not what it would be if p were true.

To be sure, conditions 1 and 2 tell us that p is true and he does believe it, but it does not follow that his believing p is sensitive to p 's being true. This additional sensitivity is given to us by a further subjunctive: if p were true, he would believe it.

$$(4) \quad p \rightarrow \text{S believes that } p.$$

Not only is p true and S believes it, but if it were true he would believe it. Compare: not only was the photon emitted and did it go to the left, but (it was then true that): if it were emitted it would go to the left. The truth of antecedent and consequent is not alone sufficient for the truth of a subjunctive; 4 says more than 1 and 2.⁵ Thus, we presuppose some (or another) suitable account of subjunctives. According to the suggestion tentatively made above, 4 holds true if not only does he actually truly believe p , but in the 'close' worlds where p is true, he also believes it. He believes

⁵ More accurately, since the truth of antecedent and consequent is not necessary for the truth of the subjunctive either, 4 says something different from 1 and 2.

that p for some distance out in the p neighborhood of the actual world; similarly, condition 3 speaks not of the whole not- p neighborhood of the actual world, but only of the first portion of it. (If, as is likely, these explanations do not help, please use your own intuitive understanding of the subjunctives 3 and 4.)

The person in the tank does not satisfy the subjunctive condition 4. Imagine as actual a world in which he is in the tank and is stimulated to believe he is, and consider what subjunctives are true in that world. It is not true of him there that if he were in the tank he would believe it; for in the close world (or situation) to his own where he is in the tank but they don't give him the belief that he is (much less instill the belief that he isn't) he doesn't believe he is in the tank. Of the person actually in the tank and believing it, it is not true to make the further statement that if he were in the tank he would believe it – so he does not know he is in the tank.*

The subjunctive condition 4 also handles a case presented by Gilbert Harman.⁶ The dictator of a country is killed; in their first edition, newspapers print the story, but later all the country's newspapers and other media deny the story, falsely. Everyone who encounters the denial believes it (or does not know what to believe and so suspends judgment). Only one person in the country fails to hear any denial and he continues to believe the truth. He satisfies conditions 1 through 3 (and the causal condition about belief) yet we are reluctant to say he knows the truth. The reason is that if he had heard the denials, he too would have believed them, just like everyone else. His belief is not sensitively tuned to the truth, he doesn't satisfy the condition that if it were true he would believe it. Condition 4 is not satisfied.⁷

There is a pleasing symmetry about how this account of knowledge relates conditions 3 and 4, and connects them to the first two conditions. The account has the following form.

- (1)
- (2)
- (3) not-1 \rightarrow not-2
- (4) 1 \rightarrow 2

I am not inclined, however, to make too much of this symmetry, for I found also that with other conditions experimented with as a possible fourth condition there was some way to construe the resulting third and fourth

⁶ Gilbert Harman, *Thought* (Princeton University Press, Princeton, 1973), ch. 9, pp. 142–54.

⁷ What if the situation or world where he too hears the later false denials is not so close, so easily occurring? Should we say that everything that prevents his hearing the denial easily could have not happened, and does not in some close world?

conditions as symmetrical answers to some symmetrical looking questions, so that they appeared to arise in parallel fashion from similar questions about the components of true belief.

...

A person knows that p when he not only does truly believe it, but also would truly believe it and wouldn't falsely believe it. He not only actually has a true belief, he subjunctively has one. It is true that p and he believes it; if it weren't true he wouldn't believe it, and if it were true he would believe it. To know that p is to be someone who would believe it if it were true, and who wouldn't believe it if it were false.

It will be useful to have a term for this situation when a person's belief is thus subjunctively connected to the fact. Let us say of a person who believes that p , which is true, that when 3 and 4 hold, his belief *tracks* the truth that p . To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it.

One refinement is needed in condition 4. It may be possible for someone to have contradictory beliefs, to believe p and also believe not- p . We do not mean such a person to easily satisfy 4, and in any case we want his belief-state, sensitive to the truth of p , to focus upon p . So let us rewrite our fourth condition as:

$$(4) \quad p \rightarrow S \text{ believes that } p \text{ and not-(S believes that not-}p\text{).}^*$$

As you might have expected, this account of knowledge as tracking requires some refinements and epicycles.

...

f →

The fourth condition says that if p were true the person would believe it. Suppose the person only happened to see a certain event or simply chanced on a book describing it. He knows it occurred. Yet if he did not happen to glance that way or encounter the book, he would not believe it, even though it occurred. As written, the fourth condition would exclude this case as one where he actually knows the event occurred. It also would exclude the following case. Suppose some person who truly believes that p would or might arrive at a belief about it in some other close situation where it holds true, in a way or by a method different from the one he (actually) used in arriving at his belief that p , and so thereby come to believe that not- p . In that (close) situation, he would believe not- p even though p still holds true. Yet, all this does not show he actually doesn't know that p , for actually he has not used this alternative method in arriving

at his belief. Surely he can know that p , even though condition 4, as written, is not satisfied.

Similarly, suppose he believes that p by one method or way of arriving at belief, yet if p were false he wouldn't use this method but would use another one instead, whose application would lead him mistakenly to believe p (even though it is false). This person does not satisfy condition 3 as written; it is not true of him that if p were false he wouldn't believe it. Still, the fact that he would use another method of arriving at belief if p were false does not show he didn't know that p when he used this method. A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly, we must restate our conditions to take explicit account of the ways and methods of arriving at belief.

[g]→

Let us define a technical locution, S knows, via method (or way of believing) M , that p :

- (1) p is true.
- (2) S believes, via method or way of coming to believe M , that p .
- (3) If p weren't true and S were to use M to arrive at a belief whether (or not) p , then S wouldn't believe, via M , that p .
- (4) If p were true and S were to use M to arrive at a belief whether (or not) p , then S would believe, via M , that p .

Commentary on Nozick

Gettier cases show that for a true belief to qualify as knowledge it must not be a lucky coincidence that the belief is true. The belief must be non-accidentally true. But what does it mean for a belief to not be accidentally true?

The key idea behind Nozick's solution to the Gettier problem is that, for a belief to qualify as knowledge or as justified, it has to stand in some *reliable* relationship to the facts that make the belief true. This raises the question of what constitutes reliability. Consider a petrol gauge in a car. The more reliable it is, the greater the probability that it will correctly indicate the amount of petrol in the tank at a particular time. But, despite its reliability, the petrol gauge may occasionally malfunction. Furthermore we could say of a brand-new petrol gauge that it is reliable even though it has never been used. This shows that the reliability of a mechanism or process is not just a property of its actual performances but also of its possible performances. Analogously, in determining the reliability of the relationship between believing that p and the truth of p , one must take into account, not only whether the relationship is in fact present, but also whether it would be present under other

circumstances. For this reason Nozick expresses the reliable relationship in question by means of two subjunctive conditionals or counterfactuals:

Variation condition: if p were not true, then S would not believe that p [not- $p \square \rightarrow$ not- $(S$ believes that $p)$]

Adherence condition: if p were, contrary to fact, true, then S would believe that p [$p \square \rightarrow S$ believes that p],

where ' $\square \rightarrow$ ' stands for the counterfactual 'if... then' connective.¹ (This is not Nozick's final theory but a first approximation.) In a slightly different formulation, a subject knows that p (i) if he believes that p , (ii) p is true, (iii) he wouldn't believe that p unless p was true, and (iv) he would believe that p if p were true. For example, S knows that the telephone is ringing if he would not believe that it was ringing unless it was ringing; and if it was ringing under slightly different conditions, he would still believe that it is ringing. The two counterfactual conditionals are meant to ensure that to know that p the belief is not accidentally true. A belief that fulfils these conditions is one that, in Nozick's expression, *tracks* the facts that make it true.

There is a subtle but important difference between subjunctive conditionals such as the variation condition and adherence condition and indicative conditionals. To see this difference compare: 'If Lee Harvey Oswald had not killed J. F. Kennedy, then someone else would have' with 'If Lee Harvey Oswald did not kill J. F. Kennedy, then someone else did'. It is not clear that if Oswald had not killed Kennedy, then someone else would have. But it is clear that if Oswald did not kill Kennedy, then someone else did.

Subjunctive conditionals make use of the idea of possible worlds. A possible world is a fictitious situation which differs from the actual world in a certain respect. To say that a belief would be true in a possible world is to say that even in a situation different from this one, the belief would be true. Possible worlds vary in their degree of resemblance to the actual world. Some are more similar to our world than others.

In paragraph [c] \rightarrow Nozick explains that the variation condition and adherence condition do not need to hold in *all* possible worlds. 'If p were true, S would believe that p ' does not mean that in *every* world in which p is the case, S believes that p . To see this, consider the following conditional: 'If kangaroos had no tails, they would topple over.'² The truth of this statement is compatible with a situation in which tailless kangaroos are provided with crutches to prevent them from toppling over. So it is not logically impossible that tailless kangaroos manage to stand on their legs. What the conditional statement really means is that, in worlds which are similar to ours, tailless kangaroos are not given crutches and do topple over. Analogously, the variation and adherence conditions are not supposed to hold in all possible worlds but only in those closest to the actual world. The relation symbolized by ' $\square \rightarrow$ ' is

¹ The subjunctive 'if... then' connective is usually symbolized by ' $\square \rightarrow$ ' and the indicative 'if... then' connective by ' \rightarrow '. However, Nozick uses ' \rightarrow ' for both types of 'if... then' relations.

² This example is borrowed from D. Lewis, *Counterfactuals* (Blackwell, 1973), pp. 8–9.

considerably weaker than the relation of logical implication which is expressed by ' \rightarrow '. The notion of reliability explicated by the variation and adherence conditions does not *guarantee* truth; it does not imply infallibility.

Before setting forth his *reliabilist* analysis of knowledge, Nozick critically discusses the causal definition of knowledge. According to the causal analysis of knowledge, for S to know that p , the belief that p must be caused by the fact that p . To see how the causal condition helps with the Gettier problem, note that in each of Gettier's examples an appropriate causal relation is absent. For instance, what causes Smith to believe that Jones owns a Ford or Brown is in Barcelona is not the fact that makes his belief true (namely, that Brown is in Barcelona) but something else, namely the fact that causes him to falsely believe that Jones owns a Ford. So the causal condition is not satisfied. The facts that make Smith's belief true are not the facts that cause him to believe. That is why, despite having a justified true belief, Smith lacks knowledge.

A problem with the causal definition of knowledge is that it doesn't seem to work for mathematical knowledge or for knowledge of future events. It is not clear how the numbers 3 and 5 cause me to believe that they make 8, or how the future fact that the sun will rise tomorrow causes me to know this fact. Nozick's reliabilist account doesn't run into this kind of problem. I know that 3 plus 5 is 8 if this condition is met: if 3 plus 5 were not 8, then I would not believe that it was.

In the passage marked [a] \rightarrow Nozick provides a further argument against the causal analysis of knowledge. Suppose you are a brain suspended in a vat (or tank) full of liquid in a laboratory, and wired up to a computer which is feeding you your experiences and beliefs under the control of some ingenious scientist. The scientist can induce in you any belief and any experience that he chooses. Now suppose the scientist produces in you the true belief that you are a brain in a vat. In this case the causal condition is met: your belief that you are a brain in a vat is in part caused by the fact that you are a brain in a vat. Does this mean that the brain in a vat knows that it is a brain in a vat? Intuitively the answer is 'no'.

Why does Nozick claim in [b] \rightarrow that cases of causal overdetermination pose a problem for the causal analysis of knowledge but not for the variation condition?

Consider the following example of causal overdetermination: you are looking at a kangaroo and come to believe that there is a kangaroo. A hologram of a kangaroo is arranged to pop up if the kangaroo is not present, but you are looking at the real kangaroo. Since it is in fact the real kangaroo that is causing you to believe that there is a kangaroo, the causal theory of knowledge is committed to granting you knowledge. However, it is questionable whether you do really know, for the situation is such that you would believe that there is a kangaroo even if there were none. This suggests that the causal analysis of knowledge is too weak. In cases of causal overdetermination, we have causation but the variation conditional is false because another causal relation would sustain the belief if the proper cause did not.

In the passage marked [d]→ Nozick discusses the *relevant alternative theory* of knowledge, which bears strong affinities to his own tracking account. According to the relevant alternative theory, a subject knows that p if he can distinguish or discriminate the truth of p from possible alternatives, i.e., states of affairs in which p is false. It is not enough, for example, to truly believe of the woman in front of you that she is Lilly to know it. If you cannot distinguish her from Milly (her twin sister), then despite the truth of your belief, and despite the justification for it (she looks exactly like Lilly), knowledge is absent. If knowledge required the discrimination of *all* possible alternatives, though, there would be little or no knowledge (must you be able to distinguish Lilly from cleverly disguised imposters?). In order to avoid this sceptical outcome, epistemologists restrict the scope of possible alternatives which a knower has to be able to discriminate among to *relevant* ones only. There are several possible views about what might be deemed ‘relevant’.

How does Nozick’s variation condition help to determine which alternatives to p are relevant and thus need to be eliminated before S can know that p ?

The variation condition may be viewed, Nozick claims, as fleshing out a criterion for when an alternative is relevant: the criterion is whether the alternative would have obtained were p false. (The philosopher who is largely responsible for the relevant alternative approach to knowledge is Dretske. In Chapter 6 we will discuss an essay in which Dretske employs the relevant alternative theory to refute scepticism.)

Some claim that the variation condition is problematic because it doesn’t permit us to admit that typically when S knows that p it will also be true that S knows that he doesn’t falsely believe that p .³ Even if S tracks the truth concerning the proposition p , the variation condition is not satisfied concerning the proposition ‘ S doesn’t believe falsely that p ’. S would still believe that he doesn’t falsely believe that p even if it were the case that S was believing p falsely.

Reread [e]→ and the following paragraphs and explain in your own words why Nozick thinks that the variation condition, by itself, is too weak to give us an acceptable definition of knowledge.

There can be cases where the variation condition is met, and yet we would not have a case of knowledge. A case in point is the brain-in-a-vat scenario. Imagine, once again, that a brain in a vat is ‘fed’ the true belief that it is a brain in a vat. When applying the variation condition we get the following conditional: if (contrary to fact) S were not a brain in a vat, he wouldn’t believe that he is a brain in a vat. This condition is indeed met. For if S had not been a brain in a vat, the scientist would not have been able to induce this belief in him via the electrodes to his brain, and hence he would not have

³ E. Sosa, ‘Postscript to “Proper Functionalism and Virtue Epistemology”’, in J. L. Kvanvig (ed.), *Warrant in Contemporary Epistemology* (Rowman & Littlefield, 1996), p. 276.

believed that he was a brain in a vat. But although the variation condition is met, intuitively it seems wrong to say that *S* knows that he is a brain in a vat. The problem is, according to Nozick, that *S*'s belief that *p* is not 'sensitive to the truth' of *p*. It is nothing but a lucky coincidence that the brain in a vat's belief that it is a brain in a vat is true. If the scientist had stimulated a different part of *S*'s brain he would have had a different belief.

Explain why a brain in a vat which believes that it is a brain in a vat meets the variation condition but not the adherence condition.

For the brain in a vat to meet the adherence condition, it must be the case that in close possible worlds the brain in a vat would believe that it is a brain in a vat. Yet it is possible that the scientist chooses to induce a different belief in the brain in a vat. The scientist could make the brain in a vat believe that it is a tailless kangaroo or that it is someone pretending to own a Ford. And because there are close possible worlds in which a brain in a vat does not believe that it is a brain in a vat, the adherence condition is not met.

While the variation condition demands sensitivity to the falsity of *p*, the adherence condition demands sensitivity to the truth of *p*. When combined, these conditions ensure that 'beliefs and facts vary together', that *S* believes that *p* because *p* is true. Nozick describes this situation as one in which the belief *tracks* the truth.

In [f] → Nozick realizes that, as it stands, the adherence condition is too demanding and rules out some obvious cases of knowledge. *S* can know that *p* without it being the case that whenever *p* is true, *S* believes that *p*. The variation condition is also too strict. Consider the following case. Suppose you acquire the true belief that Mr Havit owns a Ford and you do so because Mr Havit has shown you the registration papers. But suppose that if he did not own a Ford, you would still have believed that he did because Tom's mother (who, unbeknownst to you, is a notorious liar) would have told you that he did. Intuitively you do know that Mr Havit owns a Ford, even though the variation condition is not met.

In [g] → Nozick constructs another counterexample to the variation condition. A grandmother might know that her grandchild is well on the basis of good perceptual evidence, but she would still believe that he is well, even if he were sick, because other family members would tell her that he is well. So the adherence condition is not fulfilled even though intuitively the grandmother does know that her grandson is well, if he is well.

Nozick deals with these cases by relativizing the variation and adherence condition to a particular method. The method by which *S* acquires a belief must be held constant from the actual to the possible world. Concerning the first case, in the actual world you rely on perception to come to believe that Mr Havit owns a Ford, whereas in the counterfactual situation you rely on testimony of Tom's mother. Similarly in the second case, if the grandmother would rely only on her perception rather than on testimony of relatives, then she would not believe that her grandchild is well when in fact he is sick. Since in either of the counterexamples there is a shift in belief-formation

method between the actual world and the counterfactual situation, relativizing the subjunctive conditionals to a particular method allows Nozick to rule out the counter-examples.

When evaluating Nozick's reliabilist analysis of knowledge, we have to examine two kinds of question. First, is the analysis demanding enough so as to ensure that a knower's belief be non-accidentally connected with the fact that makes it true? Second, is the analysis liberal enough so as to include all cases of knowledge?

Some critics think that the adherence condition is too strong. To understand their worry, consider a doorbell with a short circuit.⁴ Whenever the doorbell rings inside the house, someone is outside pressing the bell-push. But sometimes pressing the bell-push does not result in the bell ringing. Given this scenario, whenever the bell rings, you know that someone is at the door. But when the bell does not ring you cannot be sure whether someone is at the door. Now the question is whether, by hearing the bell ring, you can come to know that there is someone at the door. Since the adherence condition is not fulfilled, Nozick seems to be committed to answer in the negative. Intuitively, however, you do know that someone is at the door when the bell rings.

Where does all this leave us? We seem to be caught in a dilemma: if we get rid of the adherence condition, we run into trouble with respect to the brain-in-a-vat scenario. And if we embrace the adherence condition we get counterintuitive results concerning the doorbell example. Some philosophers have taken this and similar problems to indicate that the search for a definition of the concept of knowledge is misguided. Maybe the concept of knowledge has vague boundaries and the only kind of conceptual analysis possible is one where the analysans is as vague as the analysandum. (The analysandum consists of the expression to be analysed and the analysans is the phrase by means of which it is analysed.)

⁴ This example was suggested to me by Fred Dretske.