Philosophy 5340 – Epistemology

Topic 2: The Problem of Analyzing the Concept of Knowledge

Analyses of the Concept of Knowledge

Part 2: More Detailed Discussion

1. The Traditional Analysis of the Concept of Knowledge

A traditional account of the concept of knowledge is as follows: knowledge = justified, true belief. Or, to put it in standard, expanded form:

Person A knows that *p*

means the same as

(1) A believes that *p*,
(2) It is true that *p*, and
(3) A is justified in believing that *p*.

2. A Brief Digression: How Does One Arrive at Analyses?

How might one arrive at this analysis? One standard way of attempting to construct analyses involves searching for all necessary conditions, and then seeing whether the combination of all of the necessary conditions will provide one with a sufficient condition. If it does, then one may very well have arrived at an analysis of the concept in question.

To understand this technique, one needs to understand the distinction between *necessary* conditions and *sufficient* conditions. This can be explained as follows:

Suppose that if *p* is true, then *q* **must also** be true. Then the truth of *p* suffices to ensure that *q* is also true, and philosophers say that *p* is a **sufficient condition** for *q*.

Suppose, on the other hand, that if *p* is **not** true, then *q* **cannot** be true either. Then the truth of *p* is necessary if *q* is to be true, and philosophers say that *p* is a **necessary condition** for *q*.

The definitions of "sufficient condition" and "necessary condition" can, equivalently, be put as follows:

(1) "A is a sufficient condition for B" means the same as "A logically entails B".

(2) "A is a necessary condition for B" means the same as "B logically entails A".

How, then, might the above, traditional analysis of the concept of knowledge have been arrived at? One natural way is by searching for *necessary* conditions of person A's knowing that *p* – that is, for conditions that must be satisfied if A is to know that *p* – and then conjoining all of the necessary conditions so discovered in the hope that the conjunction of those necessary conditions, taken together, will be a *sufficient* condition for A's knowing that *p* – that is, will be a condition that logically suffices to ensure that A knows that *p*.

Illustrations: (1) If Bruce is the brother of John, what more specific things must be the case? (2) If Mary knows that there is a cat on the mat, what more specific things must be the case?

3. Objections to the Justified True Belief Analysis of Knowledge

3.1 Suggested Counterexamples to the Analysis

3.1.1 Gettier's Counterexamples to the Traditional Analysis

The traditional analysis of knowledge as justified true belief is exposed, as Edmund Gettier shows in his famous article, to counterexamples, since there are cases where a person has a justified true belief that something is the case, but where that does not count as a case of knowledge. (Alvin Goldman, in footnote 1 of his essay "A Causal Theory of Knowing" – page 463 of Michael Huemer's anthology, *Epistemology - Contemporary Readings*, describes Gettier as simply reminding us of a point that had been noticed much earlier – in 1912 – by Bertrand Russell. Gettier's point, however, does not seem to be **explicitly** present in Russell's discussion, since what Russell says is not that something can be a case of a justified true belief, and yet not be a case of knowledge, but rather: "Thus it is clear that a true belief is not knowledge when it is deduced from a false belief." But this leads rather quickly to Gettier's point, by means of the additional claims that false beliefs can be justified, and that what follows from a justified claim is itself justified.)

Gettier offers two counterexamples to the traditional analysis of the concept of knowledge:

(1) Smith, Jones, and the person with ten coins in his pocket who will get the job.

Basic structure: Smith has strong evidence for the conjunctive proposition that Jones will get the job and Jones has ten coins in his pocket. Smith concludes that the person who will get the job has ten coins in his pocket. The latter belief turns out to be true, but not because Jones gets the job: for the job goes to Smith, who happens also to have ten coins in his pocket. So Smith's belief that the person who will get the job has ten coins in his pocket appears to be both true and justified, but the claim is that one would not regard it as a case of knowing.

Why is the belief a justified one? The idea here is that if you're justified in believing p, and p logically entails q, then you're **potentially** justified in believing q. And if, in addition, you do believe q, and believe q because you know that p logically entails q, then you are actually justified in believing q.

This transmission of justification via the relation of logical entailment seems intuitively plausible. But one can also offer an argument for it, if one can connect up justification with probability in a certain sort of way – namely, if it is true that if the proposition that B is at least as likely to be true as the proposition that A, relative to your evidence, then you are at least as justified in believing B as in believing A.

The inference here is via **existential generalization**, where existential generalization is an inference of the following logical form:

a is F

Therefore: There is something that is F.

Or, to use standard logical notation:

F*a*
Therefore: (∃x)F*x*.

(2) Either Jones owns a Ford, or Brown is in Barcelona.

Basic structure: Smith has strong evidence that Jones owns a Ford, and, as a result, comes to believe that either Jones owns a Ford, or Brown is in Barcelona. (Smith has no evidence concerning the whereabouts of Brown.) It turns out, contrary to Smith's strong evidence, that Jones no longer owns a Ford. Brown, however, just happens to be in Barcelona. So Smith's belief that either Jones owns a Ford, or Brown is in Barcelona, is in fact true. Moreover, since that disjunctive belief follows deductively from something that Smith does have strong evidence for – namely, the proposition that Jones owns a Ford – it would seem that Smith also has strong evidence for the belief that either Jones owns a Ford or else Brown is in Barcelona. But, it is claimed, though this belief is true and justified, it is not a case of knowledge.

In this case, the inference is via **disjunctive addition**, where disjunctive addition is any inference of the following logical form:

p
Therefore: *p* or *q*

Or, to use standard logical notation:

p
Therefore: *p* v *q*

3.1.2 The Absent-Causal-Connection Type of Counterexample to the Traditional Analysis

Suppose that one sees a table, and has a visual experience of there being an apple on the table, and forms the belief that there is an apple on the table. Suppose further that there really is an apple on the table, of just the sort that there appears to be. Then presumably one has a justified true belief that there is an apple on the table. But suppose, finally, that there is no causal connection at all between the apple on the table and your visual experience of the apple variety, since there is an opaque screen in front of the apple – a screen that one cannot see – and in front to the screen there is a holographic image of an apple that produces, by accident, just the sort of visual experience that the apple would produce if the screen and the holographic image were absent.

In that situation, would one know that there was an apple on the table? Most people certainly think that one would not know this. If so, then we have a different sort of counterexample to the analysis of knowledge as justified true belief.

3.1.3 The Abnormal-Causal-Connection Type of Counterexample to the Traditional Analysis

A variant on the previous counterexample is this. Suppose that a computer produces a certain holographic image of an apple only if a scan of the table shows that there is an apple on the table that happens to match exactly the image that the computer will produce. Once again, most people say that one does not know, in that situation, that there is an apple on the table. If so, then one has a counterexample not only to the analysis of knowledge as justified true belief, but also to **certain** ways of supplementing that analysis by adding on a condition involving causal relations between one's perceptual state and the object of which one has knowledge.

3.1.4 Indiscriminability Counterexamples to the Traditional Analysis

Suppose that one is driving through a part of Nebraska that traditionally has had many barns that were visible from the road. It is a sunny day, and one points at a nearby barn that one can see clearly, and says, to one's young child, "That's a barn." Given that it was in fact a barn, did one know that it was a barn?

It is very natural to say that one did. But suppose that this part of Nebraska is being used by Hollywood to make farming movies, and that the particular piece of road that one was driving along has only one barn – the one that one pointed to – and that all of the other things that look like barns are just barn facades – Hollywood sets, with barn fronts, but no sides or backs, etc. If that were the case, so that it was just an accident that what one pointed at was a barn, would one have *known* that it was a barn? (People go different ways on the answer to this question.)

Here's a slightly different case. Suppose that one points at a person, Mary, whom one knows very well, and says, "That's Mary." Suppose, further, that Mary has an identical twin who looks just like her, in appearance, style of dress, etc., but that one is unaware of this. Since relatively few people have identical twins, let alone ones from whom they are indistinguishable later in life, one is surely **justified** in believing that the person one had pointed to was Mary. But is it true that one would **know**, in that situation, that it was Mary?

3.1.5 Undermining Evidence Counterexamples to the Traditional Analysis

Suppose that you go into a room with some friends, and see an apple on a table, and form the justified, true belief that there is an apple on the table. In normal circumstances, most people would say that you would know that there is an apple on the table. But suppose that all of your friends who are now in the room with you were in the room a few minutes earlier, at which time they saw a holographic image that exactly matches the real apple that you, and they, now see, and that they were shown at that time that it was just a holographic image. So at this moment, they think that they are once again seeing only a holographic image, and they are wondering if you realize that it is a holographic image.

If you knew what they remember, and what they are now thinking, you would, presumably, no longer be justified in believing that you were seeing an apple. So there is **potential evidence** that would undermine the justification for your belief, and some

philosophers maintain that the existence of that potential evidence makes it the case that you do not **know** that there is an apple on the table.

This case is a variant on a 'Tom Grabit' case, originally introduced by Gilbert Harman. In that case, you see a person that you know very well – Tom Grabit – steal a book from a library. Later, however, and unbeknownst to you, Tom's mother lies, and says that Tom was out of town on that day, and says that Tom has a twin brother – Buck Grabit – who stole the book. Harman maintained that the existence of that undermining evidence – evidence of which you have no knowledge – would make it the case that you do not know, at any time after Tom's mother tells the lie, that Tom Grabit stole the book. If Harman is right, we have another type of counterexample to the analysis of knowledge as justified true belief. Other philosophers, however – such as Keith Lehrer and Thomas Paxson – reject Harman's claim that one would not have knowledge in this sort of case.

3.2 A Different Type of Objection

Most objections take the form of counterexamples. But one can also argue that the justified-true-belief analysis fails to sustain a certain entailment that seems necessary to the concept of knowledge:

The justified true belief analysis doesn't entail what intuitively appears to be the right relation between (Kp and Kq) and Potential-K(p & q), since it doesn't entail the following:

$$(Kp \text{ and } Kq) \Rightarrow \text{Potential-}K(p \,\&\, q).$$

Explanation

(1) A belief can surely be justified even if the epistemic probability of its being true is less than 1. Suppose, then, that a belief is justified if and only if its epistemic probability is greater than some threshold k. (A natural idea – and in my view the correct one – is that $k = 0.5$, but all that matters for the present argument is that there is some threshold that is greater than 0 and less than 1.)

(2) Suppose, then, that there are two propositions p and q, such that the epistemic probability of p for person S is greater than k, and similarly for q. Then S is justified in believing that p and also justified in believing that q. But the epistemic probability of the conjunction p and q for S could perfectly well be less than k, in which case S would not be justified in believing that p and q. In short, given the following notation,

"Prob(p) = k" means that the epistemic probability that p has for person S is equal to k

"JBp" means that S is justified in believing that p

the following entailments do **not** hold where k is some number greater than 0 and less than 1:

$$[\text{Prob}(p) > k \text{ and } \text{Prob}(p) > k] \Rightarrow \text{Prob}(p \,\&\, q) > k$$

$$\text{JB}p \,\&\, \text{JB}q \Rightarrow \text{JB}(p \,\&\, q)$$

(3) As a consequence one can, given the justified true belief analysis of knowledge, know that p and know that q, without its being the case that if one infers the conjunction of p and q from one's belief that p and one's belief that q, one is justified in believing that p and q, and so without its being the case that one knows that p and q. One can, in

short, know that *p* and know that *q* without that entailing that one thereby **potentially knows** that *p* and *q*.

4. 'Skeptical Doubts' Concerning the Soundness of the Concept of Knowledge

Before going on to consider how the traditional analysis of knowledge as justified true belief might be revised to avoid these objections, I think one should pause briefly to consider the somewhat radical idea that the whole idea of knowledge may be flawed, in that it may be at bottom incoherent.

4.1 The Problem of the Relation between Knowledge and Justification

How is knowledge related to justification? It seems very plausible that knowledge entails justified belief. But if that's right, what is the relevant level of justification? If one knows that *p*, is one justified in being completely certain that *p*, in the sense of believing that there is no chance at all that *p* is false? But then how much knowledge will one have?

An alternative is that if one knows that *p*, then one is justified in believing that it is more likely that *p* is true than that it is false. But is a probability that *p* is true of just more than 50% really enough for knowledge?

If neither a 100% probability nor a probability of more than 50% seems satisfactory, then should one opt for some probability greater than 50% and less than 100%? But now the problem is that any such choice would seem to be arbitrary.

4.2 The Question of the Closure Condition

The best formulation of the 'closure condition' for knowledge is a tricky matter. For our purposes here, however, the following formulation will do:

The Closure Condition for Knowledge

Suppose:

(1) *S knows* that *p*;
(2) *p* entails – that is, logically necessitates – *q*;
(3) *S* knows that *p* entails *q*;
(4) *S* comes to believe that *q* because she believes both that *p*, and that *p* entails *q*.

Then:

(5) *S* knows that *q*.

Most philosophers hold, I think, that the closure condition is true. But some philosophers, such as Robert Nozick, have argued that the closure condition is false. Could this disagreement be due to the presence, in the concept of knowledge, of ideas that are inconsistent?

4.3 The Closure Condition and Justified Belief

Here is one way of attempting to argue that there is an inconsistency.

(1) If closure does hold, then one special case of it is that if one knows that *p* and one knows that *q*, then one potentially knows that (*p* and *q*). (Let's abbreviate that as (*Kp* and *Kq*) ⇒ Potential-*K*(*p* & *q*).)

(2) Knowledge entails justified belief.

(3) The level of justification cannot in general be equal to one, since, at the very least, there are very few propositions where there is no chance at all that the proposition is false, and if one had knowledge only if one was justified in believing that there was no chance that a proposition was false, one would have very little knowledge, if any.

(4) But if there are a number of propositions, each of whose probability is less than one, but greater than some threshold *k* (such as $k = 0.5$, or $k = 0.9$, or $k = 0.99$), it may very well **not** be the case that the probability of the **conjunction** of those propositions has a probability that is greater than the threshold *k*. So if the level of justification can be less than one, then it cannot be a necessary truth that (*Kp* and *Kq*) ⇒ Potential-*K*(*p* & *q*), that knowledge is closed even under conjunction, let alone under entailment in general.

5. Possible Responses to Counterexamples That Appear Sound

In general, there are three main ways in which one might go, given the Gettier counterexamples, which are surely sound, along with any other counterexamples or other objections to the traditional tripartite analysis of the concept of knowledge that appear sound:

(1) One possibility is that what is required is not supplementation by a fourth clause, but a **stronger version** of clause (3) – the justification clause – and probably also of clause (1);

(2) A second possible reaction is to conclude that clauses (1) though (3) in the above analysis of the concept of knowledge need to be <u>supplemented</u> by a fourth clause;

(3) A third possibility is to replace clause (3) by one or more other clauses. (Doing so typically leads to what are known as strongly '**externalist**' accounts of knowledge, since, in the case where the proposition that *p* is about some state of affairs external to the mind of the believer, such accounts involve not only the condition that *p* is true, but that certain other external states of affairs exist as well, such as that the causal process that generated the belief in question is a reliable one.)

6. The First Type of Response: The Strengthening Strategy

6.1 A. J. Ayer's Strengthening Strategy: Knowledge and Certainty

A. J. Ayer offered the following analysis of the concept of knowledge:

S knows that *p* = def.

(1) It is true that *p*,
(2) *S* is sure that *p*, and
(3) *S* has a right to be sure that *p*.

Objections/Possible Problems?

What is it to be sure that something is the case? One natural answer is that to be sure that *p* is the case is to hold that there is **no chance at all, however small, that *p* is false.** To have the right to be sure will then be to be justified in believing that there is no chance at all that *p* is false, that the probability that *p* is false is equal to zero.

On this interpretation, it does seem that the Gettier counterexamples will be blocked: those cases will not be cases of knowledge, since in those cases one is not justified in being certain that the proposition in question is true. But this will also be so in almost all cases where one claims to have knowledge, since there will almost always be **some** chance that the proposition in question is false. (This issue is discussed at more length in section 6.2.)

The alternative will be to go with some less demanding interpretation of what it is to be sure, and to have the right to be sure. But once there is some possibility that the proposition that one believes is false, then it would seem that Gettier cases once more are possible.

Summing Up

It seems that either Ayer's analysis entails that one has virtually no knowledge, or else it, too, is exposed to Gettier counterexamples.

Plus Features of Ayer's Analysis?

This analysis, under the strong interpretation of what it is to be sure, and to have the right to be sure, entails what intuitively appears to be the right relation between (*Kp* and *Kq*) and Potential-*K*(*p* & *q*). The reason is that while, as was noted earlier, if there are a number of propositions, each of whose probability is less than one, but greater than some threshold *k* (such as $k = 0.5$, or $k = 0.9$, or $k = 0.99$), it may very well not be the case that the probability of the **conjunction** of those propositions has a probability that is greater than the threshold *k*, if instead there are a number of propositions, each of whose probability is precisely equal to one (and not just infinitesimally close to one), then the probability of the conjunction of those propositions must also have a probability that is equal to one.

6.2 William W. Rozeboom and the Strengthening Strategy

Another advocate of the strengthening strategy is William W. Rozeboom, in his article "Why I Know So Much More Than You Do" (*American Philosophical Quarterly* 4 (1967): 281-90). Rozeboom's discussion is superior to Ayer's in certain respects, in that Rozeboom offers more support for his analysis, and addresses the crucial objection.

Rozeboom's starting point is clause (1), where he argues that it is absurd to say, "He knows that *p* but isn't entirely sure of it." So clause (1) needs to be replaced by:

(1*) A is completely certain, subjectively, that *p* is the case.

Rozeboom then wants to maintain that that subjectively certain belief will then have to be justified if one is to have a case of knowledge. So (3) needs to be replaced by something like:

(3*) A is justified in being completely certain that *p* is the case.

Next, Rozeboom suggests that a person, A, is not justified in being completely certain about the truth of *p* on the basis of some body of evidence, *E*, unless, given *E*, *p* couldn't possibly be false. But if this is right, then, in the first of Gettier's examples, Smith would not be justified in being completely certain that Jones would get the job, for the evidence that Smith had was perfectly compatible with Jones's not getting the job – as in fact turned out to be the case. As a consequence, neither would Smith have been justified in being completely certain that someone with ten coins in his pocket would get the job. So given Rozeboom's amended analysis of the concept of knowledge, Smith didn't know that someone with ten coins in his pocket would get the job.

Question: What is one to say about this handling of the Gettier problem?

There are, I think, two main issues that need to be considered. The first concerns the question of how much knowledge we turn out to have, given Rozeboom's analysis of the concept of knowledge. For it certainly looks as if the scope of our knowledge is going to be, at the very least, **very** restricted. First, the evidence one has concerning the occurrence of **past** events is always, it would seem, compatible with those events' not having taken place, so it would seem to follow that one never has any knowledge of the past. Secondly, if one's beliefs about the external world are, as many philosophers have held, inferential beliefs based upon knowledge of one's sensory experiences, then neither will one have any knowledge of the external world, on Rozeboom's account of knowledge, since one's sensory evidence does not entail the existence of any external state of affairs. But even if some beliefs about the external world are non-inferential, is one justified in being absolutely certain that they are true? Is there not some non-zero chance that they are false?

Rozeboom, however, is well aware of this consequence. Indeed, he is inclined to think that, given the account of knowledge that he has defended, **no one ever knows anything**. He contends, however, that there is not really anything unsettling in that conclusion. In support of that, he cites our everyday use of words like "spherical". Such words may be useful, even if there is no object that is strictly spherical in shape: what we are doing is applying a term to things that **approximate** to being perfectly spherical. And this is, he suggests, what we are also doing when we speak about knowing various things: we are referring to cognitive states that approximate, to differing degrees, the impossible ideal of knowledge.

If Rozeboom is right, then one might try to argue – as Rozeboom himself does – that the project of attempting to get an analysis of the concept of knowledge that includes just those cases that one normally classifies as knowledge, and leaves out cases that one does not normally so classify, is a rather dubious undertaking: the difference is only a matter of differing approximations to an ideal, along a continuous scale, and the precise place that one draws the line along that scale can hardly have much significance.

One of Rozeboom's main conclusions, accordingly, is that the project of offering an analysis of the concept of knowledge is best set aside, so that attention can be devoted to the important issues in epistemology. As he says at the conclusion of his article:

"With problems of 'How strongly should X believe *p*?' lying dark and unfathomed before us, we stand to profit from continued epistemological

preoccupation with the nature of 'knowledge' to just about the same extent as would psychology from a return to study of the 'soul'."

I think that there is much to be said for Rozeboom's view that analysis of the concept of knowledge is not really crucial to epistemology. The major epistemological issues can all be framed, I believe, in terms of justified belief. But – and this is the second main issue that I think should be raised concerning Rozeboom's account – one can still ask whether Rozeboom's strengthening strategy provides an answer to the Gettier problem. The answer, it seems to me, is that it does not. For while the Gettier cases are not cases of knowledge if one adopts Rozeboom's stringent account, they are disqualified only at the cost of disqualifying the vast majority of ordinary knowledge claims that we normally view as sound. On the other hand, if one invokes the idea that ordinary knowledge claims are acceptable only if one thinks in terms of approximations to the ideal case where one is justified in being completely certain, it would seem that one then has no basis for not also classifying the Gettier cases as cases of knowledge as well, since, **considered simply as approximations to the ideal of absolutely certain beliefs**, it would seem that the Gettier cases are no further from the ideal than cases that one does classify as knowledge. For the epistemic probability, for example, that either Jones owns a Ford or Brown is in Barcelona, may be as close to the value of one as one wishes. So it would seem, in short, that strengthening alone cannot explain the Gettier cases: some supplementation is also needed.

7. The Second Type of Response: Supplementation Strategies

Various proposals have been advanced for adding a fourth clause to the knowledge as justified true belief analysis of knowledge –

(1) *S* believes that *p*
(2) It is true that *p*
(3) *S* is justified in believing that *p*

– in order to arrive at a satisfactory analysis of the concept of knowledge that handles Gettier-cases, and also other types of cases that give rise to objections to the traditional account of knowledge as justified true belief. Among the most important suggested ways of supplementing the traditional analysis are the following:

1. Michael Clark's "No False Intermediate Belief" strategy;

2. Roderick Chisholm's "No False Belief Is Justified" strategy;

3. Various versions of a "No Defeaters"/"No Undermining Evidence" strategy.

The latter strategies come in a number of different versions:

(1) No potential undermining evidence at all;
(2) No potential undermining evidence whose denial one is actually employing;
(3) No potential undermining evidence that one is justified in believing not to exist;
(4) No potential undermining evidence of a certain, difficult to specify sort.

Comments

1. Version (1) is mentioned by Gilbert Harman in his book *Thought* (Princeton: Princeton University Press, 1973, pages 120-72), but rejected because of an argument

that he offers there, and which is considered below. Version (2) appears to be just an alternative description of Approach 1 – the "No False Intermediate Belief" approach. Version (3) is the approach advanced by Lehrer and Paxson, while version (4) is the one favored by Harman.

2. Given that (2) is equivalent to the "No False Intermediate Belief" approach, the question is whether there is any reason to accept versions (1), (3), or (4) of the No Defeaters"/"No Undermining Evidence" approach.

3. The answer will depend on whether there is reason to think that in at least some of the Tom-Grabit-style cases one's knowledge is undermined by evidence that one does not possess. If there is, then which of (1), (3), or (4) should be adopted will depend upon what view should be taken of various-Tom Grabit-type cases.

4. Gilbert Harman holds that the view that one knows that something is the case only if no undermining evidence at all exists for that proposition is untenable. If so, then version (1) is to be rejected. Is Harman right about this?

Harman's Argument

The argument that Harman offers is based upon considering a proposition of the form *k* or *not-h*, where *k* is some true, but antecedently very improbable proposition, and *not-h* is the denial of the justified true belief whose status as knowledge is being considered. Harman's argument then runs as follows:

(1) *k* or *not-h* is true, since, by hypothesis, *k* is true.

(2) If the true proposition *k* or *not-h* were added to one's evidence, then it together with the fact that *k* is antecedently very improbable would make it reasonable for one to conclude that *not-h* was probably true.

(3) Hence, the true proposition that *k* or *not-h* would undermine one's justification for accepting *h*.

(4) Since it is always possible to find such a proposition *k* that is true but antecedently very improbable, the justification for any belief could always be defeated if any true proposition can serve as a defeater.

A Possible Objection to Harman's Argument

Suppose I think that I have excellent reason for thinking that *h* is true. If I suddenly am presented with evidence that *k or not-h* is true, but where the evidence in question is not in itself evidence against the truth of *h*, does the fact that *k* is antecedently very improbable necessarily give me grounds for **continuing** to believe that *k* is not likely to be true? Why am I not justified in concluding that, given that I have excellent reason for believing *h*, the proposition that *k*, though antecedently improbable, must also be true, since I have just learned that it is likely that *k* or *not-h* is true?

****

5. Notice the following, apparent peculiarity of the Lehrer/Paxson view – that is, version (3) of the "No Defeaters"/"No Undermining Evidence", according to which what is crucial is potential undermining evidence that one is justified in believing not to exist. According to this view, evidence that would not otherwise be undermining becomes undermining if one is completely justified in believing the denial of that evidence.

Thus, consider the Tom Grabit case. If one simply doesn't know whether Tom's mother has said that Tom didn't do it, then the fact that she has said this does not, according to Lehrer and Paxson (*Epistemology – Contemporary Readings*, page 466, paragraph 3) undermine the claim that one knows that Tom stole the book. But if one is completely justified in believing that she did not say this, then the fact that she did say it defeats one's justification for believing that Tom stole the book. This seems like an odd combination of views. (The last sentence of the third paragraph on page 466 suggests that Lehrer and Paxson are muddling together (1) being completely justified in believing that *e* is false and (2) **making use** of the assumption that *e* is false in the reasoning that supports the belief whose status as knowledge is being considered.)

6. Gilbert Harman discusses two Tom Grabit cases. In the one case, Tom's mother testifies that Tom did not steal the book and, although she is lying, no one at the hearing knows that she is lying. In the other case, no one at the hearing believes her, since she is known to be a pathological liar. Harman claims that in the first case, one would not know that Tom stole the book, but that in the second case one would. Is Harman right about this? If he is, one is confronted with the problem of saying, **in a general way**, **exactly what** it is that makes the difference between the two cases – a **crucial** problem with which Harman himself does not bother to grapple.

7. One way of trying to capture the idea that evidence that one doesn't possess, and that would, if one had it, render the belief in question improbable, **sometimes** undermines one's knowledge claim, and **sometimes doesn't** is by employing the idea of "total evidence". Thus, one might say that if the belief that *p* is justified on the basis of the belief that *q*, then *e* is undermining evidence if and only if (1) given both *q* and *e*, the belief that *p* would no longer be justified, and (2) given *t*, where *t* is the total evidence that is actual – and where *t* thus contains both *q* and *e* – *p* would no longer be justified.

In short, the idea would be that *e* must not only undercut one's justification, but it must do so even when conjoined with the rest of the totality of evidence.

Notice, however, that the line that is drawn in this way does not agree with the one that Harman wants to draw between the two Tom Grabit cases. For if the relevant test is the effect of total evidence, then in neither of the Tom Grabit cases that he discusses will it be true that one does not know, since in both of those cases the proposition that Tom's mother is lying will be part of the total evidence.

8. Can a claim to knowledge **ever** be undermined by evidence (a) that one does not possess, and (b) whose denial one does **not** have to assume in justifying the belief in question?

The idea behind this question is that if one thinks that one's knowledge claim is undermined in some Tom Grabit-style cases, perhaps the explanation is that if one is to be justified in believing that Tom Grabit stole the book, one must be justified in believing either (a) that there is no evidence that would render the belief in question unreasonable when conjoined with the evidence that one has, or – perhaps more plausibly – (b) that the belief would not be unreasonable given the totality of the evidence that one might have.

8. Michael Clark's Supplementation Strategy: True Belief Not Based upon False Belief

S knows that *p* = def.

(1) It is true that *p*,
(2) *S* believes that *p*,
(3) *S* is justified in believing that *p*, and
(4) *S*'s justification for believing that *p* does not go through any false beliefs.

Objections/Possible Problems?

1. Richard Feldman's counterexample, described below.

2. This analysis doesn't entail what might ***appear*** to be the right relation between (*Kp* and *Kq*) and *K*(*p* & *q*).

3. Does it handle the apple/holographic image case in a satisfactory way **if direct realism is true**? It would seem that it does not, since if one's belief that there is an apple on the table is non-inferentially justified, one is not making any use of a false belief.

4. The case of evidence that is partly false, but where the false part can be jettisoned. (However, Clark's account can be easily modified to avoid this objection.)

Plus Features?

1. This analysis blocks the Gettier counterexamples, and in a very natural way.

9. Richard Feldman's Counterexample

Richard Feldman, in his article "An Alleged Defect in Gettier Counter-Examples" (*Australasian Journal of Philosophy* 50), responds to the claim that the Gettier counterexamples are defective on the ground that they assume that false propositions can justify other propositions. He does so, not by arguing that false propositions can justify other propositions, but by offering a variant on Gettier's case that does not involve any reasoning that goes through false propositions. Here is Feldman's counterexample:

(1) Mr. Nogot gave Smith very strong evidence for the proposition that he, Mr. Nogot, is in the office, and owns a Ford.

(2) Smith believes, and justifiably, the following proposition:

> (*r*) Mr. Nogot gave him, Smith, very strong evidence for the proposition that he, Mr. Nogot, is in the office, and owns a Ford.

(3) Smith concludes, and justifiably:

> (*s*) Someone gave me, Smith, excellent evidence for the proposition that he is in the office and owns a Ford.

(4) Smith also concludes, and justifiably:

> (*t*) Someone gave me, Smith, excellent evidence for the proposition that there is someone in the office who owns a Ford.

(5) Smith then forms the belief:

> (*u*) Someone in the office owns a Ford.

The final belief is true, and justified, and Smith hasn't gotten to it via any false beliefs, since (*r*), (*s*), and (*t*) are both true. (Notice that (*r*), (*s*), and (*t*) merely say that Smith was given certain evidence, and are compatible with its being the case that the

evidence involved some statements that were themselves false. So in arriving at (*u*), Smith has not made use of any false propositions.)

10. The General Idea Underlying Feldman's Counterexample

Let *e* be John's evidence that Smith owns a Ford, let *p* be the proposition that Smith owns a Ford, and let *q* be the proposition that Brown is in Barcelona. The basic idea involved in Richard Feldman's counterexample is that one can arrive at the belief that (*p* or *q*) by a different, and more unusual route, but one that involves perfectly sound reasoning, and which does not go through any false intermediate beliefs. Here is a slightly expanded exposition of the basic pattern, in which John arrives at the belief that)*p* or *q*) not by the natural, Gettier-example route, by, rather, the following alternate route:

(1) John knows that *e*.

(2) John knows that *e* entails (*e* or *q*).

(3) John believes (*e* or *q*), and he does so because of (1) and (2).

(4) John knows that it is a theorem of probability theory that if the probability of *B* given *A* is equal to *k*, then the probability of (*B* or *C*) given (*A* or *C*) must be equal to or greater than *k*.

(5) John concludes that if *e* provides good support for *p*, then (*e* or *q*) provides good support for (*p* or *q*), and he does so because of (4).

(6) John knows that *e* provides good support for *p*.

(7) John concludes that (*e* or *q*) provides good support for (*p* or *q*), and he does so because of (5) and (6).

(8) John believes that (*p* or *q*), and he does so because of (3) and (7).

John's belief that (*p* or *q*) is surely justified. For given that, by hypothesis, he knows that *e*, and that he knows that *e* provides good support for *p*, he *would* be justified in believing that *p*, if he did so. But then, given that, in view of (3), he *knows* that (*e* or *q*), it follows from (5), together with the fact that he would be justified in believing that *p*, that he must be justified in believing that (*p* or *q*). So we have a case of a justified, true, belief that is not a case of knowledge. But this justified true belief has not been arrived at by any inferences that go through false beliefs. So Michael Clark's "“No False Intermediate Belief” analysis of knowledge fails in the face of Feldman's modified, Gettier counterexample.

11. A "Chisholm-Inspired" Analysis of the Concept of Knowledge

The conceptual framework that Chisholm uses involves some concepts – and, in particular, the concept of a proposition's being **evident** – that we have not considered. But the following is an account that is suggested by Chisholm's discussion, both in *Theory of Knowledge* (Englewood- Cliffs: Prentice Hall, 1966, page 23, footnote 22), and in *Foundations of Knowing* (Minneapolis: University of Minnesota Press, 1982, pages 45-9):

S knows that *p* = def.

(1) It is true that *p*,
(2) *S* believes that *p*,

(3) *S* is justified in believing that *p*, and
(4) *S* has a justification, *j*, for believing that *p* such that *j* does not justify any false belief, *q*.

Objections/Possible Problems?

1. Lehrer and Paxson suggest that "it seems reasonable to suppose that every statement, whatever epistemic virtues it might have, completely justifies at least one false statement" (page 470), but they do not offer any support for this claim. If they are right, then Chisholm's analysis entails that we have no knowledge. But are Lehrer and Paxson right?

The claim that it is reasonable to suppose that every statement "**completely** justifies" (emphasis added) at least one false statement seems very implausible.

But one might shift to the weaker claim that it seems reasonable to suppose that every statement, whatever epistemic virtues it might have, **justifies** at least one false statement, which, if true, shows that Chisholm's analysis is unsatisfactory. But even this weaker claim – which we'll return to later – is far from unproblematic.

2. This analysis doesn't entail what intuitively appears to be the right relation between (*Kp* and *Kq*) and *K*(*p* & *q*).

Plus Features?

1. This analysis blocks the Gettier counterexamples.

2. This analysis also handles Richard Feldman's counterexample.

3. It also handles both of the apple/holographic image cases **even if direct realism is true,** since one can argue that whatever it is that justifies one in believing that there is an apple on the table also justifies one in accepting the false proposition that one's visual experiences are **caused** (in the normal way) by an apple – or, alternatively, the false belief that one is **seeing** an apple.

12. Keith Lehrer and Thomas Paxson's Account: Nonbasic Knowledge as Undefeated, Justified True Belief

1. Rather than offering an account of the concept of knowledge in general, Lehrer and Paxson offer separate accounts of **basic knowledge** and **nonbasic (or inferred) knowledge**.

2. The definition of **basic knowledge** that Lehrer and Paxson offer is as follows:

"We propose the following analysis of basic knowledge: *S* has basic knowledge that *h* if and only if (i) *h* is true, (ii) *S* believes that *h*, (iii) *S* is completely justified in believing that *h*, and (iv) the satisfaction of condition (iii) does not depend on any evidence *p* justifying *S* in believing that *h*." (464)

3. The definition of **nonbasic knowledge** that Lehrer and Paxson offer is as follows:

"Thus we propose the following analysis of nonbasic knowledge: *S* has nonbasic knowledge that *h* if and only if (i) *h* is true, (ii) *S* believes that *h*, and (iii) there is some statement *p* that completely justifies *S* in believing *h* and no other statement defeats this justification." (465-6)

4. A crucial notion in the account of nonbasic knowledge is the idea of **defeasibility,** which they initially define as follows:

"The following definition of defeasibility incorporates this proposal: when *p* completely justifies *S* in believing that *h*, then this justification is defeated by *q* if and only if (i) *q* is true, (ii) *S* is completely justified in believing *q* to be false, and (iii) the conjunction of *p* and *q* does not completely justify *S* in believing that *h*." (467)

5. Lehrer and Paxson say that this account is "basically correct", but falls prey to a technical problem. The technical problem arises from the fact that that if *S* is **not** completely justified in believing *q* to be false, but *q* is such that the conjunction of *p* and *q* does not completely justify *S* in believing that *h*, one can define a new proposition *q** such that it will be true both that *S* is completely justified in believing *q** to be false, and the conjunction of *p* and *q** does not completely justify *S* in believing that *h*. The trick is choose any proposition *r* that is **irrelevant** to the justification of *h*, but which is such that *S* is completely justified in believing *r* to be false, and then to define *q** as the conjunction of *q* and *r*.

6. Thus they are led to offer the following, revised account:

"We propose the following definition of defeasibility: if *p* completely justifies *S* in believing that *h*, then this justification is defeated by *q* if and only if (i) *q* is true, (ii) the conjunction of *p* and *q* does not completely justify *S* in believing that *h*, (iii) *S* is completely justified in believing *q* to be false, and (iv) if *c* is logical consequence of *q* such that the conjunction of *c* and *p* does not completely justify *S* in believing that *h*, then *S* is completely justified in believing that *c* is false." (468)

Objections/Possible Problems?

1. Lehrer and Paxson make use of the concept of **'complete justification'** in their accounts of both basic knowledge and nonbasic knowledge, but offer no explanation of what complete justification is.

2. If the concept of complete justification is interpreted strongly – namely, as justification that enables one to be completely certain that the proposition in question is true – then their account entails that we have very little knowledge. On the other hand, if it is not interpreted strongly, then their account doesn't entail what intuitively appears to be the right relation between (Kp and Kq) and $K(p \ \& \ q)$

3. On the Lehrer/Paxson account, a true proposition *q* that would undermine one's justification if one knew that it was true counts as a defeater **only if** one is completely justified in believing that *q* is false. Perhaps this is right, but other philosophers – such as Gilbert Harman – have thought that a true proposition that would undermine one's justification if one knew that it was true would count as a defeater even if one had no justification – and certainly no complete justification – for believing that the proposition was false. So what are the grounds for thinking that Lehrer and Paxson, rather than Harman, are right on this matter?

4. Lehrer and Paxson think that, in Harman's first Tom Grabit case – where Tom's mother is not believed by others to be a pathological liar – one is not completely justified in believing that it is false that Tom's mother said that Tom was out of town, and that Tom has an identical twin who stole the book. This certainly seems true. But if it is true, then isn't there some significant possibility that Tom's mother did say that, and if that is so, how can one be **completely** justified in believing that Tom stole the book, given that there is a significant possibility of an occurrence that is such that, if one knew of that occurrence, one would no longer be justified in believing that Tom stole

the book? (Whether this objection can be sustained may depend, I think, upon exactly what Lehrer and Paxson mean by 'complete justification'.)

5. Consider, again, the first Tom Grabit case. *S* knows nothing at all about Tom's mother having said anything, and according to Lehrer and Paxson, *S* in that case knows that Tom stole the book. Now suppose that *S* then acquires evidence that makes it likely that Tom's mother did **not** say that Tom was out of town, and that Tom has an identical twin who stole the book. As this evidence increases, *S* will at some point presumably be completely justified in believing that it is false that Tom's mother said that Tom was out of town, and that Tom has an identical twin who stole the book, and, so at that point, according to Lehrer and Paxson, *S* will **no longer** know that Tom stole the book. Does that seem right? If one knew before one acquired evidence for the false proposition that Tom's mother did **not** say that Tom was out of town, and that Tom has an identical twin who stole the book, how can the evidence for that false proposition undermine the knowledge that one previously had?

6. The analysis that Lehrer and Paxson offer of **basic** knowledge does not appear to generate the correct result in the apple/holographic image cases **if direct realism is true**, since, provided the direct realist holds that the belief in question is completely justified, their analysis appears to entail that the person in question has basic **knowledge** that there is an apple on the table.

This objection could be avoided by adding the "no defeater" requirement to the definition of **basic** knowledge. Alternatively, one might also try to answer this objection by maintaining that a direct realist view of perceptual knowledge is false, but I do not think that that is a promising avenue, since surely an analysis of knowledge should be neutral on the issue of whether direct realism is true.

Plus Features?

1. This analysis blocks the Gettier counterexamples.

2. It handles the apple/holographic cases **if indirect realism is true**, since one does have a false, justified belief about the presence of a causal connection.

3. It handles Richard Feldman's counterexample, and does so while being less restrictive than Chisholm's analysis

13. The Third Type of Response: Jettisoning the Justification Condition

The third main way of responding to the Gettier and other counterexamples to the knowledge-is-justified-true-belief analysis of the concept of knowledge is to jettison the justification requirement, and to add one or more new clauses. This will generally result in a thoroughly externalist account of knowledge, according to which one can know that *p* **without having access to any internal state of oneself – either experiences or other justified beliefs – that justify one in believing that *p***.

Here the most important alternatives are as follows:

1. Alvin Goldman's "Causal Connection" approach;

2. Nozick's "Knowledge as Tracking" strategy;

3. The "Discrimination and Counterfactuals" strategy;

4. Harman's "Inference to the Best Explanation" strategy.

14. Alvin Goldman's Causal Analysis of the Concept of Knowledge

14.1 Goldman's Basic Approach

The basic idea behind the causal connection approach is that what is going wrong, in the Gettier-type cases, is that the following two things are not connected in the right way:

(a) The evidence that justifies the belief in question;

(b) The state of affairs in the world that makes the belief in question true.

And what is the right sort of connection? Goldman's answer is that it is certain sorts of **causal connections**.

Illustration: In the "Either Jones owns a Ford or Brown is in Barcelona" example, what makes this sentence true is that Brown is, as a matter of fact, and unbeknownst to Smith, in Barcelona, whereas what makes it reasonable for Smith to believe that the sentence is true is evidence that makes it likely that Jones owns a Ford, and these two states of affairs are not causally connected in any relevant way. Similarly, consider the missing-causal-chain sort of case. If the relevant causal chain were **present**, then what makes it true that there is a piece of chalk on the table would be the cause of one's evidence that there is a piece of chalk on the table.

What sorts of causal connections are needed? Goldman suggests that there are two crucial patterns:

Pattern 1: The state of affairs that makes the belief in question true is **a cause** of the evidence that one has in support of the belief.

Pattern 2: The state of affairs that makes the belief true and the evidence that makes the belief reasonable have **a common cause**.

The basic idea is that pattern 1 applies in the case of perceptual knowledge and memory knowledge, whereas pattern 2 applies in the case of knowledge of future events.

14.2 Criticisms of the "Causal Connection" Strategy

There are a number of possible criticisms of the causal connections approach:

(1) Even where pattern 1 obtains, one may still not have knowledge.

Illustration: The modified apple case, where the laser light operates only if a sensor detects the presence of an apple. Or compare the – rather more controversial – barn case.

(2) Neither pattern 1 nor pattern 2 seems to provide an account of one's knowledge of **laws of nature** – both causal laws and non-causal laws. For it doesn't seem to be true either that the state of affairs that makes it the case that something is a law causes the evidence that we have for the existence of the law, or that the former state of affairs and the evidence have a common cause.

It is, in part, this problem that leads Goldman to appeal to the idea that one can combine causal connections with logical connections, and view the combination as still classifiable as a causal connection. And it is also this problem that leads Harman to

advocate replacing references to causal connections by references to what he calls "inferences to the best explanation".

(3) As both Goldman and Harman point out, it looks as if the causal connections thesis concerning knowledge can be **derived** from an account that doesn't refer to causal connections, along the following lines.

Consider something that one cannot directly observe. How can one have knowledge of the existence of such a thing? Most philosophers who accept a foundationalist view of knowledge would argue that one can have such knowledge only if the entity is connected with things of which one can have direct knowledge. (Michael Huemer's Principle of Phenomenal Conservatism, however, allows one to have noninferentially justified beliefs about unobservable objects.) But I think that this claim can be replaced with a less contentious claim, as follows. Suppose that one has inferential knowledge of object S that is based upon knowledge of T, and where one's knowledge of T is **either** inferential **or** noninferential. Then isn't it plausible to think that such inferentially based knowledge is possible only if there is some sort of connection between objects S and T? But if so, what forms can such connections take? A natural answer, I suggest, is that the connections must be either causal or nomological – that is, a matter of laws, but possibly either causal laws, or else non-causal laws of co-existence.

Illustrations: Principle of inference to a common cause. Case of two properties that are always found together – such as, perhaps, unit negative charge, and a certain mass.

But if it is in virtue of such connections that things not directly observable are knowable, then isn't it plausible that one can have such knowledge only when one **knows** that the relevant causal or nomological connections exist? If so, one has the following thesis:

Thesis Concerning Inferential Knowledge

One can have inferential knowledge of some entity, S, only if the knowledge is based upon the knowledge that S is connected, either causally or via laws of nature, with some entity T of which one can have knowledge, either inferential or noninferential.

Comment: Notice that this thesis has been formulated so as to be **neutral** with respect to the choice between a foundationalist view of knowledge and a coherentist view.

The point here is not restricted, however, to the case of **knowledge**. Thus, consider some object about which one cannot have **noninferentially** justified beliefs. How can one have any justified beliefs about such an object? Isn't it plausible that one can have justified beliefs about it only if it is connected with things concerning which one can have justified beliefs – either inferentially justified beliefs or noninferentially justified beliefs? But if so, what forms can such connections take? The natural answer, once again, is that the connections must be either causal or nomological.

But if it is in virtue of such connections that one can have justified beliefs about things about which one cannot have noninferentially justified beliefs, then isn't it plausible that one can have such justified beliefs only when one is **justified in believing** that the relevant causal or nomological connections exist? If so, one has the following thesis:

Thesis Concerning Inferentially Justified Beliefs

One can have inferentially justified beliefs about some entity, S, only if one arrives at the belief in question on the basis of a justified belief that S is connected, either causally or via laws of nature, with some entity T about which one can have justified beliefs – either inferentially justified beliefs or noninferentially justified beliefs.

Given this latter thesis, one can then appeal to the first of the supplementation strategies – i.e., the "No False Intermediate Belief" approach – to conclude that whenever one has knowledge of things that one is not directly observing, appropriate causal or nomological connections must obtain. The argument in question runs as follows:

(1) Assume Mary knows that *p*, and that her knowledge is inferential, rather than direct.

(2) If Mary has inferential knowledge that *p*, then Mary's belief that *p* is inferentially justified.

(3) Given the above Thesis Concerning Inferentially Justified Beliefs, it follows that Mary cannot be inferentially justified in believing that *p* unless she arrives at that belief on the basis of a belief that the state of affairs that makes *p* true is connected, either causally, or via laws of nature, with some state of affairs concerning which she can have justified beliefs – either inferentially justified beliefs or noninferentially justified beliefs.

(4) But if Mary's belief that *p* is justified **on the basis** of a belief that certain states of affairs are connected, either causally, or via laws of nature, with the state of affairs that makes *p* true, then the latter belief is an **intermediate conclusion** in the process of reasoning that she uses to arrive at the belief that *p*.

(5) But according to the first response to the Gettier counterexamples, Mary knows that *p* if and only if she has a justified, true belief that *p*, and her justification does not go through any false intermediate conclusions.

(6) It follows from (4) and (5) that Mary's belief that certain states of affairs are connected, either causally, or via laws of nature, with the state of affairs that makes *p* true, being an intermediate conclusion, must itself be true.

(7) And so it follows that if Mary is to have inferential knowledge that *p*, certain relevant states of affairs must be connected, either causally, or via laws of nature, with the state of affairs that makes *p* true.

In short, the thesis that inferential knowledge presupposes the existence of causal or nomological connections is a **derived thesis** – following from the "No False Intermediate Belief" approach together with the above thesis concerning inferentially justified belief.

A Minor Comment: Notice that the argument just set out does not make the assumption that the inferences involved are "inferences to the best explanation". For in some cases, a belief will be justified not because it is the best explanation of something else, but because it is a likely effect. (Consider, for example, the justification of one's belief that the sun will rise tomorrow.)

15. The "Discrimination and Counterfactuals" Strategy

This second externalist approach – which can take either a partially internalist form or a thoroughly externalist form, depending upon whether it is formulated in terms of justified true belief, or simply in terms of true belief – rests upon the contention that in the barns and barn facades cases, one does not have knowledge because one does not – by hypothesis – have the ability to distinguish between the barn that one is actually seeing and a mere facsimile that one might have been seeing.

When, then, does one have knowledge according to this second approach? What is required in addition to true belief, or justified true belief? The answer is that one needs to ask whether, if the situation **had been different** in certain ways – ways such that the belief in question would have been false – one **would** have noticed the difference, and **would**, as a consequence, **not have had the belief in question**.

What one has to consider, then, is whether **certain counterfactuals** are true or false. What is a counterfactual? Basically, it's an if-then statement which implies that the antecedent, the "if" clause is false, and which makes an assertion about how the world **would have been different if – contrary to fact – the "if" clause had been true**.

Illustration: Consider some salt that's not in water, and a piece of chalk that's not in water. One can ask what **would** happen if each **were** now in water. And the answer is that if the salt were in water, it would be dissolving, whereas the piece of chalk would not be dissolving.

What determines whether a given counterfactual is true or false? That's a complicated question, and a variety of answers – some of them quite different – have been offered. One answer, which was advanced by Robert Stalnaker and David Lewis – and which is still popular, though it is open to decisive refutation – is in terms of similarity relations among possible worlds. But let us focus upon a different answer, according to which what counterfactuals are true is generally a matter of what causal laws there are. On this view, what makes it the case that if a certain piece of salt were in water, it would be dissolving is, first, that salt has a certain molecular structure, and secondly, that there are laws that entail that anything with such a molecular structure will dissolve when in water.

Given the present approach to the analysis of the concept of knowledge, the idea, then, is that one has to consider possible ways in which the situation might have been different and such that the belief in question would have been **false**, and then one asks, of each, whether one would then have noticed the difference, and, as a consequence, not have acquired the belief in question.

Consider, for example, Henry and the barn. Instead of a barn, there could have been a mere facade. If so, it would have been false that Henry was seeing a barn. Would Henry have noticed the difference? If the answer is that he would not have, then, according to this "Discrimination and Counterfactuals" approach, in the case where Henry was actually seeing a barn, Henry **did not know** that he was seeing a barn. For the following counterfactual, rather than being true, would have been **false**:

"If it had been, not a barn, but a facsimile, then Henry would not have believed that there was a barn there"

Possible Objections/Problems

1. One problem with this approach, at least as stated to this point, is that it would seem that it might **always** be the case that there is some way in which the situation could have been different which is such that one couldn't have detected the difference – one could have been hallucinating, or been a brain in a vat, or confronted with a holographic image, or a facsimile, etc. So it would seem that if ordinary knowledge claims are to be preserved, one has to restrict in some way the **range of alternatives** that are taken into account when one considers how things **might** have been different. Not all logical possibilities can be considered, nor even all possibilities that are compatible with the laws of nature that there are in this world. (Being a brain in a vat certainly seems to be a possibility that is allowed by the laws of nature.)

In addition, we're also confronted with the problem of the possibility of a range of cases starting with ones that differ only marginally from nearby cases, but which range through situations that differ only slightly, but that end up with radically different situations: e.g., lots of facsimiles in Henry's immediate vicinity versus a facsimile off on a planet orbiting around a star in a distant galaxy. Where is the line to be drawn, and in virtue of what underlying principle?!

2. Secondly, there is the fact that, rather than there being something approximating to general agreement that Henry does **not** know in the barn-case that he is seeing a barn, at least quite a fair proportion of people hold that Henry **does** know that he is seeing a barn.

3. The latter intuition connects up, moreover, with a picture of knowledge that seems fairly appealing. According to this picture, there are only three sorts of facts that are relevant to the question whether a person – Anthony – has knowledge of some object A. First, there are facts about Anthony's **internal states** – what beliefs he has, what processes of reasoning he goes through, etc. Secondly, there are facts about the **objects of his beliefs**. Thirdly, there are facts about the **connections** – causal and nomological – between the objects of his beliefs and his internal states. Once these three things are fixed, it seems natural to think that it is also fixed whether Anthony does or does not have knowledge in the case in question, and that how the rest of the world is – that is, the world aside from his internal states, the objects of his beliefs, and the connections between the two – does not affect things one way or the other. Such further facts can neither make it the case that he has knowledge, nor make it the case that he does not.

4. The final, and, I believe, the most fundamental comment that I have to make regarding this approach to the analysis of the concept of knowledge is that if the basic claim involved in it is true, then it seems to me it is true **because** it follows from a different account – that is, from the approach that appeals to the idea of undermining evidence that one does not possess.

Why do I think this is so? Consider Henry and the barn. If all the other barn-like things in the vicinity are barn facades, then the appeal of the idea that Henry does not know that there is a barn in front of him is at its strongest. Now consider how things are as the ration of barns to barn facades becomes greater. Doesn't the appeal of the view that Henry doesn't know there is a barn in front of him become correspondingly less? Or imagine that the barn facades, rather than being in the immediate vicinity, are further away. Once again, doesn't the appeal of the view that Henry doesn't know there is a barn in front of him become correspondingly less? If so, then one needs to

explain that, and the "Discrimination and Counterfactuals" approach fails to do so. By contrast, it seems to me that the idea of undermining evidence may well do so, since the sequence of situations that I've just mentioned are described by propositions that form a sequence of propositions that range from propositions that have strong evidential relevance to the proposition that Henry is seeing a barn – where there are many barn facades in the immediate vicinity – to propositions that have weak evidential relevance – where there are only a few barn facades, some distance away. In short, there are variations in the strength of the potential undermining evidence that correlate with the strength of the appeal of the idea that Henry does not know that there is a barn in front of him.

16. The "Knowledge as Tracking" Strategy

The third thoroughly externalist strategy that I want to consider is related, in certain respects, to the previous approach, and shares with it the use of counterfactuals. This account is essentially that set out by Robert Nozick, and can be summed up in the slogan that **knowledge is belief that tracks truth**.

What is meant by "belief that tracks truth"? First, a belief cannot track truth unless the belief is true. But this by itself is not enough. It must also be the case that – and here's the counterfactual element – that **if** the proposition in question **had not been true**, then the person in question **would not have believed it**.

So, though this is not quite the view that Nozick himself advances, one might put forward the following proposed analysis:

S knows that *p*

means the same as

(1) *S* believes that *p*;
(2) It is true that *p*;
(3) *S* is justified in believing that *p*;
(4) If *p* **had not** been true, then *S* **would not** have believed that *p*.

The clause added – clause (4) – formulates the tracking condition, and it is that clause that is intended to deal with problematic cases, especially the Gettier-type cases.

Consider, then, a Gettier case. In view of clause (4), one has to ask, for example, whether Smith **would** have believed that either Jones owns a Ford or Brown is in Barcelona **if** that proposition, rather than being true, **had been false**. That proposition could have been false in various ways, of course, but the idea here is to imagine the world being changed in some minimal way. Perhaps Brown leaves Barcelona a little earlier, so that he isn't in Barcelona at the time in question. If that **had** been the case, would Smith still have believed that either Jones owns a Ford or Brown is in Barcelona? The answer, surely, is that he would have – since his belief is based upon evidence concerning Jones's owning a Ford, and there is no reason why that would be affected by Brown's leaving Barcelona a bit earlier. So Smith's belief that either Jones owns a Ford or Brown is in Barcelona does **not** track the truth of that proposition. He would still have believed that, even if it had been false.

Comment

This is an interesting account of the concept of knowledge, but it has at least one consequence that seems rather counterintuitive – namely, it entails that falsity of what has been called the "closure condition" for knowledge.

The Closure Condition for Knowledge

The closure condition can be formulated as follows.

Suppose:

(1) *S* knows that *p*;
(2) *p* entails – that is, logically necessitates – *q*;
(3) *S* knows that *p* entails *q*;
(4) *S* comes to believe that *q* because *S* believes both that *p*, and that *p* entails *q*.

Then:

(5) *S* knows that *q*.

Why does the knowledge-as-tracking account entail that the closure condition for knowledge is false? Consider, first, the question of whether you can know, given the tracking account of knowledge, that you are **not** a brain in a vat having precisely the experiences that you are now having. The problem is that even if you have a justified, true, belief that you are not a brain in a vat, the tracking condition will **not** be satisfied. For the question one has to ask is whether the following counterfactual is true:

> "If the proposition that you are not a brain in a vat having precisely the experiences that you are now having were not true – so that you were in fact a brain in a vat **having precisely the experiences that you are now having** – then you would **not** believe that you were not a brain in a vat."

And the answer is that since, by hypothesis, all of your experiences and apparent memories would be just as they are now, you would still believe that you were not a brain in a vat. So the belief that you are not a brain in a vat having precisely the experiences that you are now having would not track truth in the way required by condition (4). So on the tracking account, you do not know that you are not a brain in a vat having precisely the experiences that you are now having.

Secondly, consider whether you can know that you are now seeing a table in front of you. Let us assume that you believe that you are now seeing a table in front of you and that that belief is both true and justified. The question is then whether **your belief tracks truth**. So one has to ask whether the following counterfactual is true:

> "If you had not been seeing a table in front of you, then you would not have believed that there was a table in front of you."

And the answer is that this counterfactual is true, for in evaluating it, one considers worlds in which it is false that you are seeing a table in front of you, but which differ **as little as possible** from the actual world. This means that one does not consider worlds in which you are a brain in a vat, or a pure spirit being deceived by a naughty angel, and where none of the physical things that you take to exist really exist. One considers, instead, worlds such as ones where someone has removed the table from the room a bit earlier.

So the situation is as follows:

You know that you are seeing a table in front of you.

You do not know that you are not a brain in a vat who is not really seeing a table.

But if you are seeing a table, then it follows necessarily that you are not a brain in a vat who is not really seeing a table. The conclusion that you can know that the former is the case while not knowing that the latter is the case – together with appropriate additional assumptions – means that the closure condition is not satisfied by the knowledge-as-tracking account.

The account that Nozick offers is different from the account we have just been considering, in that clause (3), which refers to a justified belief, is dropped and a counterfactual clause about one's belief is added. Here is Nozick's account:

S knows that *p* = def.

(1) It is true that *p*,
(2) *S* believes that *p*,
(3) If *p* were not true, then *S* would not believe that *p*, and
(4) If *p* were true, then *S* would believe that *p*.

Objections/Possible Problems?

1. This, too, is an interesting account of the concept of knowledge, but like the account just considered, it has the counterintuitive consequence that it entails the falsity of what has been called the "closure condition" for knowledge.

2. A second possible objection is that Nozick's account entails that the skeptic is right about some crucial claims. In particular, it follows from Nozick's knowledge-as-tracking account that

(1) One cannot know that one is not a brain in a vat;

(2) One cannot know that one is not dreaming.

Now it is not out of the question that these things are true. But is it plausible that they should be a more or less **immediate consequence** of one's **analysis** of the concept of knowledge? This seems to me very implausible.

17. Harman's "Inference to the Best Explanation" Strategy

This final approach is not so much a self-contained strategy as an idea that can be combined with other approaches, and especially with either of the first two approaches. Thus, it can be shown, I think, that when this account of inference is combined with either the "No False Intermediate Belief" approach or with Chisholm's approach, one can derive the conclusion that either appropriate causal connections or appropriate nomological connections are essential if one is to have inferential knowledge – a fact that has to be simply postulated on the "causal connections" approach.

It is also possible to combine Harman's inference-to-the-best-explanation account of knowledge with the "No Undermining Evidence" view. This is what Harman himself wants to do, since he thinks that in at least some Tom Grabit-type cases one fails to have knowledge because of the existence of undermining evidence that one is not aware of.

18. Summing Up: An Overview of the Alternative Supplementation Strategies

The various supplementation strategies can, I think, usefully be classified in terms of their acceptance or rejection of the following theses:

Thesis 1: Knowledge = Justified belief, plus the truth of relevant **beliefs**.

(The idea here is that while, in view of Gettier's counterexamples, it is not just the truth of ***p*** that is relevant in determining whether one's justified belief that *p* is a case of knowledge, the relevant truths are restricted to propositions that one **believes**.)

Thesis 2: In determining whether a justified true belief is a case of knowledge, the truth of propositions that one does **not believe** may also be relevant.

Thesis 3: The right sorts of **causal connections** are also crucial to whether a given justified true belief is a case of knowledge.

Thesis 4: The truth-values of relevant **counterfactual statements** are also crucial to whether a given justified true belief is a case of knowledge.

19. Interrelations Between the Above Four Theses

19.1 The Derivation of Thesis 3 from Thesis 1

I have suggested that there is considerable appeal in the idea – found in both Gilbert Harman and Alvin Goldman – that Thesis 3 can be derived from Thesis 1, when the latter is conjoined with a plausible thesis concerning justified beliefs about what cannot be immediately or directly or non-inferentially known, namely:

Thesis Concerning Inferentially Justified Beliefs

One can have inferentially justified beliefs about some entity, S, only if one arrives at the belief in question on the basis of a justified belief that S is connected, either causally or via laws of nature, with some entity T about which one can have justified beliefs – either inferentially justified beliefs or noninferentially justified beliefs.

19.2 The Derivation of Thesis 4 from Thesis 2

In the case of Thesis 4, I have considered two accounts that appeal to **counterfactual** statements. On the one hand, there is the "Knowledge as Tracking" type of account, which I have suggested should be rejected on the grounds, first, that it violates the Closure Condition on Knowledge, and, secondly, that it entails in an immediate fashion that certain skeptical claims are true. On the other hand, there is the account that appeals to counterfactuals connected with abilities to discriminate between situations in which the belief in question is true, and those in which it is false. What I now want to argue is that Thesis 4, under that interpretation, can be derived from Thesis 2.

In the Henry and the barn case, it is assumed that Henry is well acquainted with barns. Suppose that is not so. Indeed, suppose that, though it has been explained to Henry both what a building is, and what a movie set, building facade is like, Henry, having lived an unusually sheltered life, has been exposed to neither. Henry is now

exposed to his first building – a barn. Does Henry know that it is a building, rather than a movie set, a building facade?

It seems to me very plausible that Henry does not, on the grounds that he has no basis for believing that his present visual experiences are more likely to be caused by a building than by a backless facsimile. If this is right, then, in the original Henry case, it would seem that a crucial piece of evidence that justifies Henry in believing that he is seeing a barn is the evidence that barns are much more numerous than movie sets – or at least, in his part of the world.

But even if evidence for the belief that barns are more frequent than barn facades in his part of the world is not crucial, it is certainly true that **he would not be justified in believing that he was seeing a barn, and so would not know that he was, if he had the evidence concerning the low frequency of barns, as compared with facsimiles, in his vicinity**. It therefore follows that if Thesis 2 were correct, and one could fail to know because of undermining evidence that one did not possess, then one would be able to **explain** why Henry doesn't know in the original barn case by appealing to the fact that most of the relevant objects in his immediate environment are not barns, but mere facsimiles – a fact that, were he to know it, would undermine his justification for believing that what he now sees is a barn.

In addition to providing an explanation, this account would have the further virtue that it will answer the question – which the fourth approach on its own apparently cannot answer – of **where the line is to be drawn**. Evidence of the existence of facsimiles on another planet will generally not be such that, when combined with the evidence one has concerning the relative frequency of barns versus facsimiles here on earth, one's justification for believing that one is seeing a barn will be undermined. In short, the line is drawn on the basis of the question: "Precisely what sort of evidence would be undermining evidence?"

The derivation of the fourth approach to the analysis of the concept of knowledge from the second approach depends, then, upon the following basic assumption:

The Undermining Evidence Criterion

The cases where the inability to perceptually discriminate is relevant are those cases where evidence concerning one's inability to discriminate would be <u>undermining evidence</u>.

Notice that the strategy involved in this derivation of Thesis 4 from Thesis 2 is different in a certain respect from the derivation of Thesis 3 from Thesis 1. For in the latter derivation, the idea was to appeal to a principle – the Thesis Concerning Inferentially Justified Belief – whose appeal is supposedly independent of which strategy one favors for revising the traditional analysis of the concept of knowledge. In the case of the derivation of Thesis 4 from Thesis 2, however, the idea is slightly different, in that one is appealing to a principle – namely, the Undermining Evidence Criterion – that will appeal, I think, to most people who are attracted to Thesis 4, in part because that principle has the virtue of providing an answer to a question that is otherwise very difficult.

My conclusion, accordingly, is that Thesis 4 appears to be less fundamental than Thesis 2, and, if so, that it owes whatever plausibility it has to any plausibility that Thesis 2 has. This, in turn, means that if the Tom Grabit cases turn out **not** to be

acceptable counterexamples to the traditional analysis of knowledge, then Thesis 4, as well as Thesis 2, seems likely to fall by the wayside.

19.3 Concluding Remarks on the Four Theses

If this is right, then, first of all, Theses 3 and 4 are less basic than Theses 1 and 2 respectively, and, secondly, there would therefore seem to be no reason not to base an analysis of the concept of knowledge either upon Thesis 1, or upon Thesis 2, rather than upon either Thesis 3 or Thesis 4.

This line of thought leaves one with the question of whether to accept Thesis 2. If, as I have suggested, there is no general agreement concerning the Tom Grabit-style cases that are needed to support Thesis 2, it seems to me preferable simply to go with Thesis 1, rather than Thesis 2, on grounds of simplicity, since I think that one should accept a more complex account only if there are clear-cut reasons for doing so.

20. My Own Proposed Analysis of the Concept of Knowledge

The analysis advanced by Michael Clark is a very natural response to a number of counterexamples to the original, tripartite analysis, but it is exposed to Richard Feldman's objection. The analysis advanced by Chisholm avoids Feldman's objection, but it may very well be true, as Lehrer and Paxson suggest, *but do not prove*, that for any justified belief, *p*, there is always some false proposition, *q*, that is justified by *p*.

The proof of this claim does not appear trivial, and it may be that it is not true. The way in which I would attempt to prove it, however, would involve a generalization of the following argument:

Suppose that one thing with property *P* has been observed – call it *A* – and has been found to have property *Q*, where *Q*, rather than belonging to a family of two or more positive properties – such as the family of color properties – is a property that something can only have or not have.

According to Laplace's rule of succession, the probability that any other given thing that has property *P* also has property *Q*, given the evidence that there are *n* things with property *P*, all of which have property *Q*, is equal to $\frac{n+1}{n+2}$. So given the evidence that *A* has property *P* and also property *Q*, the probability that that any other given thing that has property *P* also has property *Q* is equal to $\frac{1+1}{1+2}$, or $\frac{2}{3}$.

It follows from this that, for any other object *B*, the probability that *B* **either lacks property *P* or has property *Q*** must be equal to or greater than $\frac{2}{3}$. Consequently, if there is, anywhere, at any time, some object *B* that has property *P* but not property *Q*, then the proposition that ***B* either lacks property *P* or has property *Q*** will be a false proposition that is confirmed by the proposition that *A* has property *P* and also property *Q*.

Generalizing this argument does not appear to be entirely trivial. But even if the generalization is false, I think that the type of case I've just described can serve as the basis of a decisive objection to Chisholm's analysis.

My idea, then, is to formulate an analysis that, like Chisholm's analysis, is more demanding than Clark's analysis, but less demanding than Chisholm's. Here is my proposal:

***S* knows that *p* = def.**

(1) It is true that *p*,
(2) *S* believes that *p*,
(3) *S* is justified in believing that *p*, and
(4) *S* has a justification, *j*, for believing that *p* such that there is no false belief, *q*, such that (a) *j* justifies *q*, and (b) *q* is such that if *S* were to become justified in any way in believing that *q* is false, *S* would no longer be justified in believing that *p* is true.

Notice that in Feldman's case, Smith is justified in believing that Mr. Nogot owns a Ford, that that belief is false, and that if Smith were to become justified in believing that that belief was false, he would no longer be justified in believing that someone in the office owns a Ford. By contrast, in the case that I just described, where one is justified in believing that object *A* has both property *P* and property *Q*, and where that justifies a false proposition that object *B* either lacks property *P* or has property *Q*, one's coming to be justified in believing that the latter proposition is false would not undercut in any way one's justification for believing that object *A* has both property *P* and property *Q*.