Philosophy 5340 – Epistemology

Topic 2: The Problem of Analyzing the Concept of Knowledge

Analyses of the Concept of Knowledge

Part 1: An Overview

1. The Justified True Belief Analysis

S knows that *p* = def.

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

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

1. Necessary conditions versus sufficient conditions.

2. Is the conjunction of necessary conditions that have so far been discovered a sufficient condition?

3. Testing analyses: Searching for counterexamples.

3. Objections to the Justified True Belief Analysis of Knowledge

Some types of possible counterexamples:

1. The Gettier counterexamples.

2. Absent causal connections: an apple and the holographic image case.

3. Abnormal causal connections: a second apple, a causal connection, and the holographic image case.

4. The indiscriminability problem: Barns versus barn facades, and identical-twin cases.

5. The existence of undermining evidence that one doesn't possess – e.g., the friends who think it is a holographic image case, and Harman's two 'Tom Grabit' cases.

A different type of objection:

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:

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

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*

"JB*p*" means that *S* is justified in believing that *p*

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

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

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

(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

1. The relation of knowledge to justification doubt. Doesn't knowledge entail justified belief? But if so, what level of justification? 100% justification, or just anything greater than 50% justification, or something between 50% and 100%?

2. Does the closure principle hold for knowledge or not?

3. If closure does hold, then, for example, (Kp and Kq) ⇒ Potential-$K(p \,\&\, q)$. But if knowledge entails justified belief, and the level of justification can be less than complete certainty, how can (Kp and Kq) ⇒ Potential-$K(p \,\&\, q)$ possibly be true?

5. Possible Responses to Counterexamples That Appear Sound

1. Strengthen at least the third clause of the justified true belief account.

2. Supplement the justified true belief account by adding a fourth clause.

3. Jettison the third, 'justification' clause of the justified true belief account, and add one or more new clauses. Introduce an externalist clause involving, e.g., reliability.

6. The First Type of Response: The Strengthening Strategy

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

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?

1. On one interpretation of what it is to be sure, this analysis entails that one has virtually no knowledge.

2. On other, less demanding analyses of what it is to be sure, it would seem that the Gettier counterexamples would again be applicable.

Plus Features?

1. On the strong interpretation of what it is to be sure, this analysis entails what intuitively appears to be the right relation between (*Kp* and *Kq*) and Potential-*K*(*p* & *q*).

6.2 William W. Rozeboom and the Strengthening Strategy

1. Rozeboom's account is as follows:

S knows that *p* = def.

(1) *S* is completely certain, subjectively, that *p* is the case.
(2) It is true that *p*,
(3) *S* is justified in being completely certain that *p* is the case.

2. Rozeboom notes that a consequence of this account is that it may well be that no one ever has any knowledge.

3. Rozeboom suggests that this is not as troubling as it first appears, since we can use the term "knowledge" to refer to **close approximations** to knowledge, strictly understood.

Objections/Possible Problems?

1. Is this what we are actually doing? Are we really saying only that we have close approximations to knowledge?

2. How does this explain the line that we intuitively draw that places the Gettier cases outside the scope of the term "knowledge", given that the probability that the belief is true in a Gettier case can be just as high as in non-Gettier cases that we classify as knowledge?

7. The Second Type of Response: Supplementation Strategies

Various proposals have been advanced for adding a fourth clause to the knowledge-is-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. Among the more important 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 are discussed at some length in section 7 of "Epistemology – Topic 2 – Part 2".)

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 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**?

4. The case where part of one's evidence that is 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.

9. Richard Feldman's Counterexample

Richard Feldman, in his article "An Alleged Defect in Gettier Counter-Examples" (*Australasian Journal of Philosophy*, volume 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 *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 <u>internal</u> state of oneself – either experiences or other justified beliefs – that justify one in believing *p***.

Here the most important alternatives are as follows:

1. Alvin Goldman's Causal Connection approach;
2. The "Discrimination and Counterfactuals" strategy;
3. Nozick's "Knowledge as Tracking" strategy;
4. Harman's "Inference to the Best Explanation" account.

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

1. The account that Alvin Goldman offers is as follows:

"*S knows that p* if and only if
the fact that p is causally connected in an 'appropriate' way with S's believing p.

'Appropriate' knowledge-producing causal processes include the following:

(1) perception

(2) memory

(3) a causal chain, exemplifying either Pattern 1 or Pattern 2, which is correctly reconstructed by inferences, each of which is warranted (background propositions help warrant an inference only if they are true)

(4) combinations of (1), (2), and (3)." (459)

2. Goldman defines Patterns 1 and 2 as follows:

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**.

Objections/Possible Problems?

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.

3. Should concepts such as those of **perception** and **memory** be part of an analysis of the concept of knowledge? Shouldn't it be a **non-trivial** result that perception and memory can generate knowledge?

4. There are two aspects of this definition that, because of vagueness, tend to shield this account from criticism. First, there is the idea of "appropriate" knowledge-producing causal processes. To see why this is problematic, consider a variant on the

apple/holographic image case, in which there is a holographic image only if the device is triggered by the presence of a real apple. Now there is a causal process that runs from the apple through the holographic image to the perceiver, but one would not count this as a case of knowing that an apple is present. If Goldman rules this out by holding that the causal process is not an appropriate one, then, since he has offered no definition of "appropriate causal process", the term appears to allow him to accept or reject causal processes as needed to avoid objections.

5. The other place where there is vagueness in the account is in connection with the "correctly reconstructed by inferences" requirement. For consider the following statement: "Though he need not reconstruct *every* detail of the causal chain, he must reconstruct all of the important links" (454). Here the problem is that it is vague what counts as an important link. Consider, for example, perception. What are the important links here? Does the causal process that runs from experiences to beliefs about external objects contain "important links"? If so, and if they have to be reconstructed by **inferences**, then a direct realist account of perception will be ruled out.

What seems to me important is simply that whatever inferences are present be ones that are justified. I cannot see how one can make any **independent** judgments about the importance of causal links, and then check to see whether all of the important causal links have been reconstructed by inferences.

6. **Explicit** references to causal connections appear unnecessary, since, at least in the case of **nonbasic** or **inferential** knowledge, inferences of a non-deductive sort will only be justified if it is reasonable to believe that the relevant sates of affairs are connected causally – or, at least, either causally or nomologically. In short, it looks as if something like the following thesis is true:

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.

7. One of the fundamental points about Goldman's analysis is that, in jettisoning the requirement that a necessary condition for a belief to be a case of knowledge is that the belief be justified, Goldman is opting for a thoroughgoing externalist account of knowledge, and it appears to be true, on Goldman's account, that one can know something without being justified in believing it. For suppose that John acquires, without knowing it, the power of telepathy, and he finds himself having the thought that Bruce is in some particular mental state. Would John be justified in believing the proposition in question? It would seem not, for someone else – Mary – might find herself having precisely the same thought about Bruce, purely by accident. Surely Mary would not be justified in believing the proposition in question. But if she is not, then how could John be justified, given that he and Mary could be in precisely the same internal state?

Conclusion: Given Goldman's proposed analysis of the concept of k knowledge, one could know that *p* without being justified in believing that *p*.

15. The "Discrimination and Counterfactuals" Strategy

This second 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 (justified) true belief? The answer is that one needs to ask whether, if the situation **had been different** in certain 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 that 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. But one traditional answer is that 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 ask, 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, 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. Robert Nozick's "Knowledge as Tracking" Strategy

Nozick suggests that the concept of knowledge can be analyzed as follows:

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 is an interesting account of the concept of knowledge, but it has at least one consequence that seems rather counterintuitive – namely, it entails the 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 upon 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 (3). 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, 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 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.

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 more or less **immediate consequences** of one's **analysis** of the concept of knowledge?!

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. 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 $\frac{2}{3}$. (A proof of this entailment is given in the appendix.) 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*.

Appendix

Introduce the following abbreviations:

'P_n' = 'Object *n* has property P'

'Q_n' = 'Object *n* has property Q'

'E' = '$(P_1 \& Q_1) \& (P_2 \& Q_2) \& \ldots \&(P_n \& Q_n)$'

'Prob(q/p)' = 'The logical probability of *q* given *p*'.

What we want to prove is that

$$\mathbf{Prob(Q_{n+1} \vee \sim P_{n+1}/E) \geq Prob(Q_{n+1}/ P_{n+1} \& E).}$$

Proof

We can prove this by proving the following general result:

$$\mathbf{Prob(}r \vee \sim p/q\mathbf{) \geq Prob(}r/p \ \& \ q\mathbf{).}$$

This can be proved as follows. First of all, the following is a theorem of probability theory:

(1) $\mathrm{Prob}(r/p) = \mathrm{Prob}(q/p) \times \mathrm{Prob}(r/q \ \& \ p) + \mathrm{Prob}(\sim q/p) \times \mathrm{Prob}(r/\sim q \ \& \ p)$

Now replace 'r' by '$r \vee \sim p$', so that we have:

(2) Prob(r v ~ p/q) = Prob(p/q) x Prob(r v ~ p/p & q) + Prob(~p/q) x Prob(r v ~ $p/$~p & q)

But

(3) Prob(r v ~ p/p & q) = Prob(r/p & q), since the only way that 'r v ~ p' can be true if 'p & q' is true is by 'r' being true

Also

(4) Prob(r v ~ $p/$~p & q) = 1.

Substituting in (2) using (3) and (4) then gives one:

(5) Prob(r v ~ p/q) = Prob(p/q) x Prob(r/p & q)+ Prob(~p/q)

Since all probabilities are equal to or less than one, we have that

(6) Prob(r/p & q) ≤ 1.

It then follows from (5) and (6) that

(7) Prob(p/q) x Prob(r/p & q)+ Prob(~p/q) ≥

Prob(p/q) x Prob(r/p & q)+ Prob(~p/q) x Prob(r/p & q)

It then follows from (5) and (7) that

(8) Prob(r v ~ p/q) ≥ Prob(p/q) x Prob(r/p & q)+ Prob(~p/q) x Prob(r/p & q)

But

(9) Prob(p/q) x Prob(r/p & q)+ Prob(~p/q) x Prob(r/p & q)

= Prob(r/p & q) x[Prob(p/q) x + Prob(~p/q)]

Then, since

(10) Prob(p/q) x + Prob(~p/q)] = 1

it then follows from (9) and (10) that

(11) Prob(p/q) x Prob(r/p & q)+ Prob(~p/q) x Prob(r/p & q)

= Prob(r/p & q)

Finally, it follows from (8) and (11) that

(12) Prob(r v ~ p/q) ≥ Prob(r/p & q)

Given this general result, replace 'r' by 'Q_{n+1} 'q' by 'E, and 'p' by 'P_{n+1}'. This gives us the result that we want:

Prob(Q_{n+1} v ~ P_{n+1}/E) ≥ Prob(Q_{n+1}/ P_{n+1} & E)