(I wrote the first essay here as an assignment for a class in epistemology. It is followed by an essay the original version of which I had written on my own about three years earlier.)

The Gettier Problem

By DOUG SHAVER
April 16, 2009

Gettier's 1963 essay "Is Justified True Belief Knowledge?" demonstrated an apparent conflict between our intuitive understanding of what constitutes knowledge, based on common usage of the word "know," and the formal definition that had previously been widely accepted by epistemologists. The formal definition sought to encode the intuitive understanding. Gettier offered two hypothetical scenarios in which a person's belief in some proposition met the criteria of the formal definition but nevertheless failed on intuitive grounds to constitute knowledge.

From common usage, we take it that a person S does not know a proposition P unless (a) S actually believes P, (b) P is true as a matter of fact, and (c) it is not just by chance that S's belief in P is a true belief. These necessary conditions have been codified by most epistemologists as the JTB — justified true belief — analysis of knowledge, according to which S knows a proposition P iff (a) S actually believes P, (b) P is true as a matter of fact, and (c) S is justified in believing P. The formal definition thus stipulated that the three criteria were not only necessary but sufficient to constitute knowledge.

Gettier's paper seemed to prove that the criteria were not sufficient. In particular, he showed that under certain conceivable circumstances, one's belief in a true proposition could be justified but still not constitute knowledge because in those circumstances, one's belief was true only by chance notwithstanding the justification. Many Gettier cases depend on a transfer of justification from some false proposition $P1$ to a distinct true proposition $P2$ that is logically implied by $P1$. It seems obvious that if $P2$ can be deduced from $P1$, then whatever justifies belief in $P1$ must also justify belief in $P2$. However, in these Gettier cases, the evidence justifying belief in $P1$ is not that which establishes the truth of $P2$. The relevance of the Same Evidence Principle (two people in

the same mental state must necessarily have the same evidence) is therefore illusory, and so a kind of equivocation occurs in the argument that if *S* is justified in believing *P*1, then *S* is justified in believing *P*2.

One solution would be a refinement in the notion of justification that would distinguish between sufficient reason for mere belief in a proposition *P* and sufficient reason to assert that one knows *P*. Let us for the moment call the former "doxastic justification" and the latter "epistemic justification." Such a refinement would, among other desiderata, imply Feldman's notion of essential dependence on a falsehood, in effect requiring a causal connection between the facts actually represented by *P* and *S*'s belief that *P* is a true statement. Thus, doxastic justification would still be transitive across deduction, but it would not thereby become epistemic justification.

<center>* * *</center>

Is it really a problem?

By DOUG SHAVER
June 2006

Edmund L. Gettier challenged the conventional definition of knowledge in a 1963 paper titled "Is Justified True Belief Knowledge?" According to the convention, a person knows a proposition *P* if and only if

- *P* is true in fact
- The person has a sincere belief that *P* is true
- The person's belief is justified.

Gettier cited two variations of the convention. He noted that Chisholm, instead of justification, held that "adequate evidence" was necessary. In Ayer's formulation, according to Gettier, the person must be "sure that *P* is true" and have "the right to be sure."

The debate over what it means to know something goes back at least to Plato and almost certainly predates him by many centuries if not more. Philosophers had settled on the "justified true belief" convention by the 20th century, and it remains the usual definition although the "Gettier problem" is often mentioned as a complication that

today's epistemologists have to deal with somehow.

In his article, Gettier contrived two examples of people justifiably believing a true proposition but, according to common usage, lacking true knowledge of the proposition. Many analogous examples have been presented in the years since the article was published. Here is a comment from the *Internet Encyclopedia of Philosophy*:

> But his article had a striking impact among epistemologists, so much so that hundreds of subsequent articles and sections of books have generalized Gettier's original idea into a more wide-ranging concept of a Gettier case or problem, where instances of this concept might differ in many ways from Gettier's own cases. Philosophers swiftly became adept at thinking of variations on Gettier's own particular cases, and, over the years, this fecundity has been taken to render his challenge even more significant (http://www.iep.utm.edu/g/gettier.htm).

The "Gettier problem," then, is to define knowledge in such a way as to preclude the possibility that the truth of *P* and a person's being justified in believing *P* coincide only by chance. But why is it considered a problem at all? Speaking of chance, the odds against a person's being in a Gettier situation are quite overwhelming. The actual probabilities would be hard to calculate, but one intuitively suspects that one would be far more likely to win the lottery than to be in such a situation. Why, then, are philosophers so concerned about it?

It is because the possibility of a chance correlation between justification and truth would contradict something we wish to be implied by any claim of knowledge. This is reflected in a common lay definition of knowledge, according to which to know something is to believe it with certainty or without possibility of error. If I say I know *P*, I want it to be understood that I don't think it possible for *P* to be false—not necessarily because I consider myself infallible, but because I think my justification for believing *P* is unassailable. It is reflected, too, in the fallacious assertion by some people, "I don't believe it, I know it." In common usage, all beliefs are fallible by definition, since uncertainty is always implied. To say you believe something is to acknowledge the possibility of being mistaken. That which is known, however, is by some presumed definition necessarily true, and therefore whatever is known cannot also be believed.

In the philosopher's lexicon, though, to believe *P* is simply to consider *P* to be a

true statement, without regard to whether it is considered certainly true or only probably true. Since anyone who knows *P* must consider *P* to be true, then by definition knowledge without belief is logically impossible. But while knowledge entails belief (and therefore cannot exist without it), they are not the same thing except insofar as knowledge can be considered a type of belief. We may refer to all other types as mere belief, and we thereby raise the question: What must accompany belief to make it knowledge?

As noted before, the conventional answer is that the belief must be justified and the proposition believed must be actually true. If I say, "I know *P*," but *P* happens to be false, then no matter how good my justification is, I am in error. It is sometimes said that at one time in human history, "everybody knew that the world was flat." It is probably true that everybody believed it, and all things considered the belief was probably justified at the time. The world does in fact look flat, and through most of human history nobody had any good reason to think that the reality differed from the appearance. Nevertheless, it is not true that anybody ever knew that the world was flat. They might have said they knew it, but they were mistaken.

Knowledge is not mere certainty. It is not equivalent to the absence of doubt. Neither is it a simple coincidence of belief with reality. The belief must be rationally justified. If I am superstitious enough to believe that I will win the lottery because I found a four-leaf clover on the day I bought a ticket, and it happens that I do win, I cannot truthfully say that I knew I was going to win the lottery. The certainty with which I expected to win is irrelevant. I had no good reason to believe I would win, and therefore I did not know I would win.

But why not? If I had said after finding the clover, "I know I'm going to win the lottery," did not my winning prove me right? No, it did not, because we want the word knowledge to mean something more than belief that just happens to be true. If I tell you that I know something, you want some assurance that it is an actual fact, and you can have no such assurance unless I can give you a good reason to think it is a fact, and I cannot give you a good reason if I don't have one myself. Of course, if I feel certain about something, then I probably think I do have a good reason, and if I give you that reason I might consider you excessively skeptical if you still don't believe what I tell you. There can be much debate, then, about what constitutes sufficient justification for a claim of

knowledge. The history of science and even of mathematics is littered with knowledge claims that seemed justified beyond debate but nonetheless had to be abandoned.The examples began to pile up with the advent of modern science, but the trend was spotted a thousand years before that by the original skeptics among the Greek philosophers, who simply denied that it was possible for anyone to know anything.

Their argument was simple. A claim of knowledge implies perfect certainty, but since no one is infallible, perfect certainty can never be justified, and therefore nobody can really know anything. That position is called philosophical skepticism. Few people who identify themselves as skeptics nowadays embrace it. They do not dispute the fallibility of all humans. Nor do they dispute the impossibility of perfect certainty. Instead they dispute the usefulness of defining knowledge in terms of such certainty. Their position is that we may properly claim to know that *P* is a statement of fact if our justification for believing *P* is of such a nature as to make it intellectually perverse to disbelieve *P*. As in certain other areas of life, perversity may sometimes lie in the eye of the beholder, but a discussion of specific cases must await another essay. We can proceed for the time being on the assumption that facts do exist and that we do know what some of them are. In other words, for some set of propositions that we are justified in believing, those propositions are actually true, and therefore we are not wrong to say that we know them.

A distinction may be noted between actually knowing something and being justified in claiming to know it. For anyone to know *P*, *P* must be true in fact. But how does one who claims to know *P* confirm that *P* actually is a fact? To cite one's justification is to beg the question. It seems inescapable that as long as we are fallible, we must accept that even our best claims to knowledge could be mistaken. This does not mean we must act as if we had reason to doubt those claims. It does not mean we have to pay serious attention to all who question the claims. It means only that as a matter of principle, we may never say that it is not possible for the claims to be wrong. (What we are discussing here is knowledge about empirical reality. What it might mean to know the laws of logic or to know that a mathematical theorem is true is for another essay.)

While perfect certainty about any proposition is unattainable, we may settle for moral certainty. This is the point at which it becomes intellectually perverse to entertain doubt. We have achieved moral certainty if, notwithstanding our human fallibility, we

have reduced the likelihood of error to a negligible amount. We may believe with a moral certainty that the sun will rise tomorrow, that the law of gravity will remain in effect, and that the stars will not appear rearranged in the night sky. We could be wrong, but we are justified in acting, for the time being, as though that were not possible.

Our certainty about these things is not justified by the fact that "everybody knows" they are so. Nor is it justified solely by a consensus within the scientific community that they are so. Common knowledge is not dependable enough to establish moral certainty, and neither is the unanimity of scientists. The moral certainty, where itexists, is established by the evidence that the scientific community has compiled in support of its consensus. That evidence is voluminous and consistent. There is an immense amount of it, and it all implies the truth of our current understanding of gravity, the earth's motion, and the nature of the stars. No fact contradicts that understanding, and virtually no one says that anything contradicts it. Where evidence is so voluminous, so consistent, and so uncontroverted, the demands of any reasonable skepticism are satisfied and doubt becomes perverse.

(This document was last updated on May 31, 2012.)