Philosophical Perspectives, 24, Epistemology, 2010

THE MYTH OF KNOWLEDGE

Laurence BonJour
University of Washington

It is my conviction that epistemology has, in recent times, gone seriously astray. The main symptom of this has been ever more complicated and convoluted accounts of the supposedly central concept of knowledge, combined with less and less attention to the question of why knowledge is supposed to be intellectually important and valuable. In effect, epistemologists have been more concerned with preserving the supposed common-sense applications of the term "knowledge" than with making sense of why there should be such a concept and of why it or its application should be topics of central concern to philosophy.

The starting point for these developments was the widespread abandonment of the strong Cartesian conception of knowledge in favor of what is sometimes referred to as the weak conception[1]: the "fallibilist" conception, according to which one can have knowledge that a particular claim is true even though one's justification (evidence, warrant, or supporting grounds) for that claim is less than conclusive — even though the having of that level of justification is compatible with the claim in question actually being false. It is widely believed that the correct specific concept of knowledge, the concept that is assumed to be embodied in common sense and ordinary thought and to be of central concern to epistemology, is one that satisfies this latter conception. My main thesis in this paper is that this whole "fallibilist" view of knowledge is mistaken, that there simply is no well-defined, intellectually significant concept of knowledge fitting this general conception: none that can be genuinely found in common sense or indeed can be constructed or stipulated in a satisfactory way. The supposed weak concept of knowledge in question is, I am suggesting, a philosophical myth.

The great majority of philosophical discussions of knowledge in recent times have assumed such a concept, and in this way are predicated upon this myth. When the myth is abandoned, some of the most widely discussed problems that focus explicitly on the concept of knowledge, I will suggest, simply disappear, with the Gettier problem and the lottery paradox being the most obvious examples.

In Part I of this paper, I will attempt to develop and support these claims. Then, in Part II, I will suggest, rather more tentatively, that the correct concept of knowledge is in fact some version of the strong Cartesian conception, under which knowledge requires a conclusive level of justification, one that guarantees the truth of the relevant claim. This conception is usually rejected, often with very little discussion, for reasons that seem to me to have in the end much less substance than is usually thought, however initially persuasive they may be. These reasons and the underlying issues they involve will be discussed below. But the main concern of the paper is to criticize the view that advocates the concept of knowledge that I regard as mythical rather than to defend this, to my mind, more plausible alternative; and it still seems to me possible that in the end the right thing to say is that there simply is no coherent concept of knowledge to be found in common sense—or at least none that is worthy of any serious philosophical concern.

Part I: Against the Weak Conception

1. Preliminaries

I begin with an account of the conception of knowledge that I am opposing. Here the obvious starting point is what is often referred to as the "traditional" account or definition of knowledge as *justified true belief*, an account that still seems to me broadly correct, albeit obviously in need of much additional clarification and specification. There are many different issues that could be raised in relation to this conception, including issues pertaining to the belief and truth conditions, but for present purposes, four central points will suffice.

First, knowledge in the sense that is allegedly delineated by this conception is supposed to be a supremely valuable and desirable cognitive state, one whose possession marks the difference between full cognitive success and at least some degree of cognitive failure: knowledge is the epistemic *summum bonum*. And it is because knowledge has this status that skepticism about knowledge seems intellectually threatening in a way that raises serious philosophical concerns. Thus, I suggest, any satisfactory concept of knowledge must be such as to make sense of this supposed supreme value; and any features or requirements that pertain to it must be ones that contribute in some intelligible way to this evaluative status. I will assume throughout that this is an absolutely essential condition for any adequate account of knowledge.

Second, whether a person has knowledge of something in this sense is supposed to be an all or nothing matter, not a matter of degree. While beliefs or opinions can be justified or rational to various degrees, intuitively one either has knowledge or one does not, with no room for gradations in between.[2]

Third, as just suggested, the relevant notion of justification, in contrast, is emphatically a matter of degree, allowing for a wide range of variations from very weak to very strong.

Fourth, the prevailing view, which I will adopt with no further discussion here, is that justification in the relevant sense (*epistemic* justification) has to do with likelihood or probability of truth, with stronger degrees of justification making the truth of the relevant belief correspondingly more likely or probable. Here I will think of justification in internalist terms: as a reason or reasons that a person has for thinking that the belief in question is true.[3]

Given this initial picture, the obvious question that needs to be asked is: what specific degree of justification is required to achieve this exalted cognitive state of knowledge? And the most initially obvious answer to this question, one taken for granted by Descartes (and also by Locke and a host of other historical philosophers), is that knowledge requires the highest possible degree of justification: justification that is *conclusive*, that *guarantees* the truth of the claim that is believed, that makes it impossible in relation to that justification that the belief is false. This is the *strong conception of knowledge* referred to above.[4]

In recent times, this strong conception of knowledge has, for reasons to be considered momentarily, almost always been dismissed as untenable. But it is worthwhile to pause for just a moment to reflect on its internal structure and, as it might be termed, its conceptual cohesiveness — on the way in which the three ingredients of the conception relate to each other and fit together to make an initially plausible, well-motivated combination, and one that is obviously of great intellectual and philosophical significance. We have one ingredient, belief, that brings a specific claim or proposition into the purview of the cognitive agent in question. We have an appealing cognitive goal, namely bringing it about that one believe such a claim only if it is true. And, since we are in general unable to directly determine the truth of a belief, we have a less direct means for achieving this goal, namely seeking justification of the epistemic sort for the claim in question and relying on the degree of justification that results as a guide for which beliefs to accept. Given such a picture, it clear that complete cognitive success has been achieved only when the justification attained for the claim is sufficient to establish completely that it is true. And it is thus natural and eminently reasonable to identify the situation of such complete success as the most valuable cognitive state, namely knowledge. No one could reasonably deny the epistemic value of such a state, and both the intuitive rationale for its constituents and the way that they fit together seems entirely clear.

But of course there is an obvious and familiar problem: familiar sorts of considerations show that the strong conception is extremely difficult to satisfy, indeed in many or even most cases apparently impossible to satisfy. Perhaps one can have conclusive justification for simple a priori claims and for claims about some aspects of immediate experience, but, most recent epistemologists would agree, for little if anything else. And that means that under the strong conception, there will be at best very little knowledge and that only in quite

restricted areas. How serious a problem this is and indeed even what exactly the problem is are issues that seem to me to be much less straightforward than is usually thought; I will return to them in Part II of this paper. But the way in which the discussion typically goes is something like this: the standard definition of knowledge is invoked, the strong conception is mentioned almost in passing, and then is dismissed at once as leading to an unacceptable skepticism, as being unreasonably stringent or demanding, as contrary to common sense, and the like. Usually all this is done in a quick sentence or two, with the central idea being that very many of the ascriptions of knowledge that ordinary people seem to make would be false if the strong conception were correct, and that it is implausible (in some way that is usually not very fully specified) that this is so.

The conclusion that is then most commonly drawn is that the correct concept of knowledge is still a version of the traditional definition, but one according to which the degree of justification required for knowledge is one that falls short of conclusive justification but is still fairly strong, substantially stronger than one that merely makes the claim barely more likely to be true than false. This general sort of view is often referred to as the *weak conception of knowledge*, and I will adopt this label for it as well. (In addition to the three conditions specified in the traditional definition, virtually all recent versions of the weak conception also incorporate a further condition aimed at warding off Gettier-style counterexamples. The specification just given should be interpreted as allowing for this sort of addition, though I will suggest below that the need for such a condition is merely one of the unfortunate consequences of adopting the weak conception.)

2. The central problem for the weak conception

Given the account so far of the weak conception of knowledge and how it is dialectally arrived at, the immediate and urgent question should be: what then is this specific level of justification, the attaining of which transforms what is up till then at best merely increasingly probable or likely true belief into the exalted state of knowledge?

Here it is perhaps unreasonable to demand a precise numerical specification of a level of probability and perhaps not even entirely clear that numerical probabilities are the right way to think of degrees of justification. But if we are to suppose that there is a definite concept of knowledge which when satisfied yields the exalted cognitive state in question, it is surely not good enough to say merely, as is commonly said, that the level of justification in question is "strong" or "high" or "adequate" or enough to make it "highly likely" that the belief in question is true, for nothing this vague is enough to specify a definite level of justification and a corresponding definite concept of knowledge.[5] And yet the striking fact is that philosophical discussions that either explicitly invoke or tacitly presuppose the weak conception of knowledge almost never have anything

much more helpful than this to say about what this "magic" level of justification, as I will somewhat tendentiously refer to it, might be — or, even more important, about *why* it has this very special status. Indeed, it is fair to say that nothing like a precise specification of the "magic" level has ever been seriously suggested, let alone more widely accepted.[6]

This failure on the part of those who either espouse or presuppose the weak conception of knowledge to offer anything like a clear specification of the "magic" level of justification or even to give any reasons for thinking that such a specification might in some way be possible seems to me to constitute in itself a very serious objection to that conception.[7] But an even more serious objection is that it is very difficult or, I believe, impossible to see what could give any level of justification that is short of being conclusive the kind of special significance that the weak conception requires it to have.

Perhaps the clearest way to think about this problem is to suppose that there is some issue about which it is important to you to find the truth, and that there is some specific proposition that constitutes the best candidate you can find for the truth about this issue. It is easy to understand how finding higher and higher levels of justification for the claim in question improves your cognitive situation, making it more likely that the corresponding belief is true, and also how finding genuinely conclusive justification, if that were possible, would be the best situation of all. But the claim of the weak conception is that there is some specific level of justification that is less than conclusive but that nonetheless transforms your cognitive situation in a much more radical way than did increases in justification up to that point (or further increases above it). Before this level is attained, you merely have a belief that is more and more likely or probable, but at that point you suddenly have *knowledge*. But *why* does achieving this specific level of justification make such a difference and what exactly is this difference supposed to amount to?

One thing that it clearly does *not* amount to is that further increases in justification cease to matter. If the issue in question is indeed an important one, then as long as the justification you have achieved remains less than conclusive, further increases in justification will still be cognitively valuable, will still further improve your cognitive situation by making the truth of your belief even more likely. Indeed, it is hard to see why such further increases are not valuable *in exactly the same way*, to precisely the same extent, as those that came earlier, before the supposed "magic" level was reached. And this means that the sorts of inquiry that lead to such increases in justification are also no less valuable, no less urgent than they were before. But why then is it supposed to be so supremely important whether or not the "magic" level has been reached? Why are we supposed to care so much about this — or indeed care about it at all? (And of course, if we really do care so much, why is so little attention devoted to determining what that level actually is?)

It is entirely possible, of course, that in relation to a particular issue, you might decide that the level of justification and corresponding likelihood of truth

attained so far is "good enough" and might be willing to rest content with that result. In some cases, when such a satisfactory level has been reached, you might even be moved to say that you *know* the claim in question — where this means roughly, I would suggest, that you are sure enough of it, that your justification *approximates* closely enough to being conclusive, for the issue at stake. (See the further discussion below, in Part II.) But clearly what that "good enough" level might be will vary widely from case to case, depending on the specific issue and the surrounding context, and thus your willingness to say that you know in such a case — or for others to ascribe knowledge to you on the same sort of basis — provides no real reason for thinking that there is one "magic" level of justification of the sort that the weak conception requires.[8]

But what then could give the supposed "magic" level of justification the kind of special significance that it is supposed by the weak conception to have? What could explain why a gradual increase of justification, rather than merely yielding gradually increasing levels of probability or likelihood or rational belief, at some point results in a qualitatively different state that is cognitively exalted in the way that knowledge is supposed to be? I believe that there is simply no answer available to this question, no account of what would give any level of justification that is short of conclusive this very special status. And this is the most fundamental reason for thinking that the weak conception of knowledge is mistaken and that the supposed concept that would satisfy it is a myth. Paradoxically enough, the weak conception, though it seeks to preserve the truth of common-sense attributions of knowledge, makes it impossible to give an account of why those attributions have any real epistemic significance or why they should matter very much for epistemology.

Before turning to other issues, there is one more fairly obvious reason for thinking that there is no non-conclusive level of justification that can play the role that the weak conception of knowledge requires. Suppose that I have attained just the supposed "magic" level of justification, but no more, for two independent claims, P and Q, which I believe on that basis and which are in fact true. And suppose that any needed further condition, such as an anti-Gettier condition — see the next section — is also satisfied, so that according to the supposed weak concept of knowledge, I know that P and know that Q. And suppose that I now infer the conjunction of P and Q, and then validly infer some further claim R from that conjunction (understanding why and how the inference is valid), that I proceed to believe R on that basis, and that R is also true (and that any further required condition is again satisfied). Do I thereby achieve the exalted state of knowledge with respect to R?

The answer, of course, is that if the levels of justification in question can be thought of as levels of probability or as behaving anything like levels of probability, then the level of justification that I thereby acquire for R will in general fall short of the "magic" level, because the product of two probabilities less than 1 is always smaller than either of them. Thus under the weak conception, there is no guarantee that when I know two things and infer even the simplest,

most obvious conclusion from them, I will then have knowledge of the result. (Of course I *may* have such knowledge, if the levels of justification are, contrary to what I have been assuming, sufficiently greater than the "magic" level.) And this seems to be an intuitively unacceptable result: what is the supposed state of knowledge really worth, if even the simplest inference from two pieces of knowledge does not lead to further knowledge?

Here we are, of course, in the vicinity of the much-discussed issue of cognitive or epistemic closure. Reasons have been given for denying that closure always holds. But, though I have no space for a detailed discussion here, those reasons usually pertain to relatively restricted cases: usually to the issue of whether anti-skeptical claims that follow from more ordinary claims must be known if the more ordinary claims are known. The denial of closure that results from the weak conception is, however, far more radical — more radical, I believe, than any that has ever been seriously advocated. And here we have an additional reason for thinking that there is no non-conclusive level of justification that can have the kind of cognitive significance that the weak conception of knowledge requires it to have, and that the supposed specific concept of knowledge that would satisfy that conception is a myth.

In addition to the problem of specifying the "magic" level of justification and of explaining what could possibly give it such a status, the assumption that the correct concept of knowledge satisfies the weak conception also gives rise to a number of collateral problems. Here I will focus on the two best known and most wisely discussed of these: in the next section, the Gettier problem; and, in the following section, the lottery paradox. My suggestion is basically that these by now familiar "problems" are at bottom not really genuine epistemological problems at all. Unlike the more traditional problems of epistemology, these supposed problems do not reflect genuine intellectual issues that require deep philosophical insight and creativity for their solution. Instead, they are entirely artificial difficulties, created by the mistaken assumption that the weak conception of knowledge is correct, and can be easily dissolved by simply rejecting that assumption.[9] The ease with which this can be done seems to me to constitute a further argument, albeit a less direct one, against the weak conception.

3. The Gettier problem

Consider first the Gettier problem. Gettier famously showed that there are easily constructed examples that satisfy the traditional definition of knowledge, construed along the lines of the weak conception, but that nonetheless do not seem intuitively to be cases of knowledge.[10] These cases may be roughly characterized as ones in which a belief is strongly justified (exactly how strongly matters little, as long as the justification is less than conclusive) and also true, but true in a way that intuitively is merely accidental or fortuitous in relation

to the justification in question — true in some unexpected way that is other than the way that the justification seems to point to. Thus in one of the most widely discussed examples, you have strong evidence that one of your colleagues, Professor Nogot, owns a Ford automobile and on this basis come to accept the general claim that at least one of your colleagues owns a Ford; the evidence is in some way misleading or deceptive, and Nogot does not in fact own a Ford; and yet a different colleague, Professor Havit, about whom you have no relevant evidence at all, does in fact own a Ford. Your belief that one among your colleagues owns a Ford is both true (because of Havit's ownership) and strongly justified (because of your evidence pertaining to Nogot), but does not seem intuitively to be a case of knowledge, with the intuitive reason being that Havit's possession of a Ford is not in any way suggested or indicated by your evidence and so is a mere accident in relation to it. As Gettier makes clear in his own original discussion, the possibility of this sort of case rests squarely on the assumption that the level of justification required for knowledge is some level that is less than conclusive and thus presupposes that the correct concept of knowledge is a version of the weak conception.

There is no room in this paper to even begin to recapitulate the extended and extremely convoluted discussion that was provoked by Gettier's argument. Fortunately, what matters for present purposes is not the details but rather the general character of the discussion: Many, many different revisions of the conception of knowledge were proposed, usually but not always taking the form of adding a fourth condition to the standard three. Many of these proposed revisions then fell victim to further counterexamples of different and often more complicated sorts, leading to still further proposals, and so on, and so on. By the time this amazingly voluminous discussion had at least largely run its course, two things were clear. First, though there were several fairly prominent general lines of solution, there was no single solution to the problem that came very close to being generally accepted. This fact by itself would perhaps not differentiate the Gettier problem from most of the standard problems of philosophy. But second, and much more important though less noticed, the discussion yielded almost nothing by way of theoretical insight into the nature of the problem and its implications for the concept of knowledge and for epistemology generally: no insight into how and why such a problem arises in the first place, no account of why the concept of knowledge should involve further conditions of the various kinds proposed, no understanding of how these further conditions might fit together with the remaining conditions to make a cohesive and intellectually significant concept, and thus no real understanding of why any of the modified conceptions of knowledge that emerged from the discussion should be regarded as intellectually important or significant — as worthy of the very extensive philosophical concern that has been lavished on them. Instead, the discussion of the Gettier problem was focused almost entirely, in a shallow, essentially ad hoc way, on piecemeal revisions and specific examples. And thus all of this

discussion, I would suggest, yielded very little in the way of genuine philosophical understanding of knowledge or of anything else.

The concern I am raising here may be illustrated by focusing on one of the most widely proposed fourth, anti-Gettier conditions, the requirement that the person's justification not be *defeasible*, where the justification for proposition *P* is defeasible if there is some true proposition *Q* such that if the person were to become justified in believing *Q*, he or she would no longer be justified (to the requisite degree) in believing *P*.[11] Thus in the example above, there are various choices for *Q*: perhaps some description of the ways in which the evidence about Nogot is misleading, perhaps some stronger evidence showing that Nogot does not own a Ford, or perhaps just the denial of the original claim about Nogot's ownership. And it seems quite plausible that in all of the familiar sorts of Gettier examples, such a true proposition, or usually more than one, can be found.

Thus the addition of such a condition would yield the intuitively delivered result that the various Gettier examples are not in fact cases of knowledge and in this way would "solve" the problem. But does the addition of such a condition really yield any insight into what is going on? Do we understand *why* knowledge, understood as a supremely and uniquely valuable cognitive state, should require or involve such a condition? In relation to the *strong* conception of knowledge, the fact that justification is in this way defeasible would of course show at once that it is less than conclusive and would thus show that the concept of knowledge is not in fact satisfied. But if justification need not be conclusive, why exactly must it not be in this way defeasible? What, other than merely avoiding the counterexamples, could be the rationale for such a requirement? And my suggestion is that without such a rationale, the significance of the supposed concept of knowledge that results from adding the indefeasibility requirement is obscure. (I would suggest that here, and in other places as well, the relevant "intuitions" in fact derive tacitly from the strong conception, which is, like many intuitively compelling ideas, much easier to explicitly repudiate than to fully escape from.)

There is no space here to consider further anti-Gettier conditions, but my suggestion is that the same general sort of objection applies to virtually all of them. In addition to various other problems, there is simply no clear intuitive rationale for them — and so also no clear intuitive rationale for the supposed weak concept of knowledge that results from incorporating them: no reason why the resulting concepts should be thought to be intellectually important and worthy of serious philosophical attention.[12]

Here is my diagnosis of the situation. Part of what gives rise to the supposed problem is obvious enough. If a claim is justified to some less than conclusive degree, then it will obviously be possible for it to still turn out to be false. And if this is possible, then it will also be possible for it to be true in some way other than the way that the justifying evidence or reason suggests and in this way to be accidentally or fortuitously true. In such a case, a person who believes that claim on the basis of the evidence in question will have a true belief, but

in an accidental or fortuitous way. But what is needed to generate a seemingly serious, indeed extremely intractable problem is the further assumption, made by proponents of the weak conception, that there is a definite concept of knowledge, a supremely valuable, all-or-nothing cognitive state, that requires only some less than conclusive degree of justification. For if this assumption is made and a case of the sort just described occurs in which the level of justification is taken (on what basis?) to be at the "magic" level or greater, the intuitive unsatisfactoriness of the situation from an epistemic standpoint will require saying that the exalted state of knowledge has not in fact been achieved. And this in turn will require assuming that some further requirement for that state that has not been satisfied, and we will be off to the unrewarding Gettier races.

Without that assumption, in contrast, there is, I suggest, simply no discernible philosophical problem to be solved. Cases of the sort in question, cases where truth is accidental or fortuitous in relation to the justification available for a claim, may at times be surprising, and it is useful to see how and why they may occur. But without the assumption of the weak conception, there is simply no need to seek a definite, precisely defined condition that will exclude them.[13] Whereas *if the weak conception of knowledge were correct, there would have to be such a condition*, no matter how difficult it is to clearly define it or, more importantly, to assign it a plausible rationale. Thus my suggestion is in effect that in light of the very unsatisfactory results of the Gettier-inspired inquiry, we should apply modus tollens rather than modus ponens to the just italicized conditional.

4. The lottery paradox

There are a number of different versions of the lottery paradox, not all of them directly concerned with knowledge. But the version I am concerned with here can be set out very simply, with the weak conception of knowledge as an essential ingredient.

Assume, following the weak conception, that there is a less than conclusive "magic" level of justification that is, together with belief and truth (and the satisfaction of an anti-Gettier condition — hereafter I will take this for granted), sufficient for knowledge. And now suppose that a lottery is to be held, in which there will be one and only one winning ticket, drawn in a fair and random way. Whatever that "magic" level of justification may in fact be, it seems clear, so long as what is essential to epistemic justification is something in the vicinity of likelihood or probability of truth, that if the number of tickets is large enough, the probability of, and so the justification for, the claim that any particular ticket will lose will reach or exceed the "magic" level. Suppose then that the number of tickets is indeed at least that large (even though we have no clear idea how large that would have to be). It seems to follow that if a person were to believe, on the basis of this high probability and resulting justification, that some particular

ticket will lose, and it will in fact lose, then that person knows, in advance of the actual drawing, that the ticket in question will lose. This already seems counterintuitive: given the weak conception, we may suppose that the person knows that one ticket will win, and the chosen one has as good a chance as any of the others. But, still worse, the same reasoning would obviously apply to all the other losing tickets, so that a person who somehow happened to form the corresponding set of beliefs on the same sort of basis would thereby know of each losing ticket that it would lose. Indeed, such a person would apparently be in a position to infer from this other knowledge, together with the knowledge that there will be a winning ticket, that a certain specific ticket will win, and have knowledge of this result!

Because it is abundantly clear on an intuitive basis that such a person does not genuinely know any of these things, the lottery example may well seem to constitute a simple and decisive *reductio* of the weak conception of knowledge — with the supposed "paradox" deriving only from the untenable assumption that the weak conception is correct. In the end, I believe that this is the correct assessment. But before accepting it as such, we must first ask whether there is any tenable response to the problem that preserves the weak conception: an intuitively plausible way of adding some further condition or requirement (in a way parallel to that which is attempted in relation to the Gettier problem) that will rule out the various claims in the lottery case as satisfying the requirements for knowledge, while still leaving the commonsensical knowledge claims that the weak conception is intended to preserve untouched.

One response is to accept that the person indeed knows the claims about the individual tickets, but to block any further inferences by rejecting the view that one always knows the validly inferred consequences of one's knowledge. As already argued above, such a view seems to me to drastically undercut the supposed value of knowledge.

Another sort of response is to impose some further condition on justification that will rule out lottery cases, while leaving other knowledge claims that involve the weak conception untouched. In fairly striking contrast to the situation with the Gettier problem, however, suggestions for how to accomplish this are pretty thin on the ground. I will look at two possibilities, starting with an earlier attempt of my own, one whose deficiencies seem to me to reveal well enough some of the pitfalls in this area. In *The Structure of Empirical Knowledge* (and an earlier paper),[14] I suggested that what goes wrong in the case of the initial claim that some specific ticket in the lottery will lose is that although this claim can be assumed to be strongly enough justified to satisfy the justification requirement for knowledge (which I was assuming to fit the weak conception), it is also one of a class of relevantly similar, equally justified alternatives, all individually very unlikely, but such that the person in question knows (according to the correct version of the weak conception) that one of these unlikely alternatives will in fact obtain and has no relevant way of distinguishing among them. And thus, I suggested further, we should add a further requirement for knowledge that rules

out knowledge in this specific sort of situation, with the background assumption being that such a requirement will cause no trouble in most other cases where the justification meets the standard of the assumed to be correct version of the weak conception.

There are many problems that can be raised for this proposal. One objection is that the problem posed by the lottery does not really require knowledge that one ticket will win: thus even if there are a few unsold tickets and no one will win if one of these is drawn, the intuitions behind the problem seem unaltered.[15] Another is that the problem seems to persist even if the chances of winning and so the resulting justifications are not precisely equal. I think that the proposed additional requirement could be modified to allow for these more complicated sorts of cases, but will not attempt to do so here, since there are more serious problems ahead.

A further, much more serious objection is that ruling out knowledge in lottery cases turns out to impugn a much wider range of ordinary, common-sense knowledge than might at first have been thought.[16] This is because many cases of common-sense knowledge entail that various lottery-type propositions will not hold and so, assuming closure, these common-sense claims cannot be known unless one can also know the denials of those lottery-type propositions — with the latter knowledge being ruled out by the proposed additional requirement. Thus, to take a rather salient example, I think I know that (*P*) I am going to continue working on this paper tomorrow. But I also may well know (by the standard of the correct version of the weak conception) that I am one of a large group of people, relevantly similar on my information, at least one of which will suffer a fatal heart attack today — and I cannot know, given the proposed requirement, that (*Q*) I will not be one of the victims. But *P* obviously entails *Q*, and so, given closure, I cannot know *P* without also knowing *Q*.[17] How far this sort of problem extends is uncertain, but a little exercise of imagination will suggest that it extends very far indeed, threatening at least a large swath of supposed common-sense knowledge.

But the deepest and most serious objection to the proposed additional requirement, in my view, is that it — and so the proposed version of the weak conception that incorporates it — again lacks any plausible intuitive rationale. One way of putting this is to say that it seems rather *ad hoc*: ruling out the objectionable cases merely by giving a slightly more abstract version of the situation in which they arise.[18] But while this way of putting the point has some intuitive force, it does not seem to me to go to the heart of the matter. What is really objectionable about the suggested additional requirement, as was the case with the first of the anti-Gettier requirements discussed above, is that we have no clear understanding of why the supremely valuable cognitive state of knowledge (understood as fitting the weak conception) should involve such a requirement — or better, turning the same point around, that we have no idea of why a state that satisfies this requirement (along with the others involved in the weak conception) should be regarded as supremely valuable from a cognitive

or epistemic standpoint, why it should be thought to have the exalted evaluative status that knowledge is intuitively supposed to have.

As with the first anti-Gettier condition, this would not be so if it was the strong conception of knowledge that was being assumed. In relation to that conception, a failure to satisfy the proposed requirement would show immediately that the justification was less than conclusive, thus failing to satisfy the justification condition, but this would also mean of course that no further, independent requirement would be required. But if it is the weak conception that is supposed to be correct, why exactly should the failure to satisfy the suggested requirement rule out knowledge? If the justification for knowledge need not be conclusive (so that knowledge is "fallible"), why should the knowledge that this possibility of mistake will in fact realized in one of a set of indistinguishable cases mean that the rest of them cannot still be genuine knowledge? How is this significantly different from the more general fact that some proportion of our claims of knowledge will turn out to be mistaken because the justification was less than conclusive? (Here again, it seems to me that intuitions whose real source is the strong conception are creeping in, making this attempt to accommodate them seem even more obviously ad hoc.)

These last two problems also afflict a different and rather simpler additional requirement proposed by Dana Nelkin (with something like this also being suggested by a number of others).[19] Nelkin's idea is essentially that mere statistical probability, of the sort involved in a lottery or similar example, *never* satisfies the justification condition for knowledge.[20] This proposal will rule out many cases of seemingly ordinary common-sense knowledge in the same way as did the previous one, because they entail lottery-type propositions that could only be justified and known on a statistical basis. Thus, in the specific example offered above, I will be unable to know on the basis of high probability that I will not be the one in the relevant group who suffers a fatal heart attack and hence will also be unable to know that I will continue to work on this paper tomorrow.

And again, the intuitive rationale for imposing this additional condition (apart from merely its success in ruling out all and only the cases we want to rule out) is far from obvious. Why, if the point of epistemic justification is finding truth and if (as the weak conception insists) conclusive justification is not required for knowledge, should high enough statistical probability not suffice? Nelkin suggests that the problem is the lack of a causal or explanatory connection between the probabilistic justification and the fact that makes the claim true (assuming that it satisfies the other requirements for knowledge and so is in fact true) and suggests further (though pretty vaguely) that this points in the direction of some sort of externalist account of justification and knowledge.

Externalist views are beyond the intended purview of this paper, but of course many internalist views (including my own) have appealed to explanatory connections as at least one basis for justification. But the problem with using the idea of explanation as a rationale for the proposed requirement is that the justificatory relevance of explanatory connections seems itself to be that they

make truth more likely or probable. To be sure, the resulting probability or likelihood is not in any obvious way statistical in character and certainly does not in general come with a specific numerical value attached. But I can find no clear reason for thinking that it is fundamentally different in kind from that which results from statistical considerations — and thus no clear way in which the justificatory relevance of explanation provides a good rationale for ruling out statistical probability as also yielding a basis for justification.

All of this seems to me to suggest strongly that the lottery paradox is indeed best viewed as a simple reductio of the weak conception of knowledge. If less than conclusive justification were enough for knowledge, then we would have knowledge in the lottery case, and it is intuitively very clear that we do not. My further suggestion is that the underlying reason for this is that it is utterly clear in a lottery case that the justification is less than conclusive; whereas in other cases that allegedly constitute knowledge under the weak conception, this lack of conclusiveness is less conspicuous and so more easily ignored.[21]

5. Summary of Part I

I have argued that the view that common sense embodies a specific concept of knowledge that both satisfies the weak conception and is of central philosophical importance faces a number of serious objections. There is no apparent way to specify or determine, even to a reasonable measure of approximation, what the required non-conclusive level of epistemic justification might be — and, more importantly, no way to say why any such level should have the sort of significance that the weak conception attributes to it. There is also no plausible intuitive rationale for the further requirements that are needed by the weak conception to solve the Gettier problem and the lottery paradox, and so no plausible account of why the supposed concept that would include these would be intellectually or philosophically important. All this suggests that a reconsideration of the strong conception is in order, which is what the second part of the paper will attempt to do.

Part II: A Tentative Defense of the Strong Conception

1. Introduction

If the weak conception is indeed untenable for the reasons indicated, then there are, as far as I can see, only two possibilities left within the limitations assumed in this paper. One is that the common-sense concept of knowledge is indeed one that fits the strong conception (with the further questions about that concept that this leaves open being of little moment here). The other is that there is after all no coherent, intellectually satisfactory concept of knowledge to be

found in common sense: none that can both satisfy the essential constraints on such a concept and make good sense of ordinary usage and practice. I will return to this last alternative at the very end of the paper. But the main task of this second part will be to see what can be said in favor of the view that the strong conception is, after all, the correct one, with the main task being to reconcile that view with the ordinary usage of the terms "know" and "knowledge."

The problem, of course, is that if the strong conception is indeed the correct account of knowledge, then most of our ordinary claims to know and attributions of knowledge will turn out to be mistaken. The precise extent to which this is so can be debated in relation to many different epistemological issues, but the details of this do not matter very much for our present purposes. It is enough to say that on the prevailing views of these various issues, the proportion of knowledge attributions that can possibly be correct according to the strong conception will be quite small at best.

As already briefly noticed above, there are many who would regard this result by itself as a conclusive refutation of the strong conception, with the idea being apparently that it is either impossible or at least very, very unlikely that common-sense knowledge attributions could be so badly mistaken. But such a conclusion seems to me to be much too quick — especially in the absence of any very appealing alternatives. It is no doubt surprising that most of the employments of a common-sense concept should be mistaken, but it does not seem to me self-evidently impossible that there could be an adequate explanation of why this is so. And this is especially the case where the concept in question is a broadly evaluative one that characterizes a kind of ideal state or situation, one that might often be approached but rarely fully realized. I believe that there are in fact resources available to a proponent of the strong conception that can go quite a long way toward an explanation of this situation — though whether they go far enough remains to be seen.

2. *Explaining false knowledge attributions: some alternatives*

How then might we try to explain, in a way compatible with the strong conception, why ordinary people make so many false attributions of knowledge, both to themselves and to others? I will suggest that there are a number of plausible explanations for individual cases of this sort. No one of these can by itself adequately handle enough of these cases, but when they are taken together, it is at least much less obvious that they do not succeed. I will begin by considering these individually.

(i) Simple epistemological error. Perhaps the most obvious explanation, one that plausibly extends very widely and that seems largely unnoticed by many philosophers, is that an ordinary person may reasonably regard the justification for a belief as conclusive even where deeper philosophical insight shows, or at least seems to show, that it is not. Consider, for example, a case where an ordinary

person seems to himself or herself to be perceiving a standard sort of "medium-sized" physical object at close range and under good conditions, and believes on this basis that such an object is there. Even if the person's justification in such a case is not in fact conclusive (because of subtle philosophical objections having to do with, for example, Cartesian demons or the possibility of being a brain-in-a-vat), it is easy to see how it might nonetheless seem to him or her to be conclusive,[22] leading to a self-attribution of knowledge (and to attributions of knowledge to others whose situation is similar).

This sort of explanation extends plausibly to a great many of the common-sense knowledge attributions that turn out, according to the strong conception, to be false, including many of those that seem intuitively most compelling. Very often the reasons why the justification in such cases is less than conclusive involve philosophical considerations, especially skeptical possibilities of various sorts, which ordinary people cannot be assumed to be familiar with (and whose obviousness is easily exaggerated by philosophers). Cataloging the full range of such cases is too large a task to be attempted here, but it would plausibly include, in addition to beliefs about the existence of physical objects on the basis of direct perception: many simple, perceptually based causal beliefs (e.g., that the stroke of an axe caused a cut in a tree), many beliefs about the mental states of other people (e.g., that the accident victim who is screaming and bleeding profusely is in great pain), many beliefs about the contents and structure of the physical world (e.g., that Chicago is north of St. Louis, or that chickens exist in the world), and so on. The general formula here is that the more complicated, subtle, and unobvious are the philosophical reasons for thinking that a given sort of justification is inconclusive, the more plausible it is that ordinary people may be led to make mistaken attributions of knowledge while still adhering in their thinking to a concept of knowledge that fits the strong conception.[23]

(ii) Exaggeration. As Butchvarov points out in a very suggestive discussion of this issue,[24] certain concepts are often used in a way that involves exaggeration, and it is clear that knowledge is one of these. Butchvarov's characterization is that the concepts in question are ones that pertain to evaluative standards or ideals and that have serious practical significance. His leading example is the concept of romantic love. It is obvious that romantic love represents a kind of ideally perfect situation with respect to a certain sort of interpersonal relation and also that the applicability of this concept has a range of practical implications that are usually viewed as highly desirable. This creates a strong pressure to attribute the state of being in love to oneself and to others even when the strong (though rather elusive) requirements for being in such a state are not clearly met — and even when they are clearly not met. Moreover, unlike some other kinds of exaggeration that are, one might say, less serious, many such attributions of love are not very readily withdrawn (even though there is very often at least a background worry about whether they are correct).

The analogy with knowledge is not perfect, but it does seem to me to be quite helpful. As viewed by the strong conception, knowledge is also an ideal

state of a certain sort and one with even more pervasive important practical significance. (See the discussion, in the next section, of assertion and practical reasoning.) Thus it is easy to see how exaggerated attributions of knowledge might be made in cases where the desirability of achieving this ideal is particularly high, especially when the gap between the actual case and the ideal is, or is viewed as being, relatively small, but sometimes even when this is not the case. Many such attributions will be withdrawn if strongly challenged or if the issue comes to seem especially urgent, but this will again not be done as readily as with other, more casual sorts of exaggeration.[25] (Knowledge may also be attributed in ways that involve more casual sorts of exaggeration, including the case, mentioned earlier, where the degree of justification is viewed as being close enough for the purposes or issues at hand. These attributions will be more readily withdrawn if seriously challenged or if the stakes turn out to be more serious.)

(iii) Close approximation. The third sort of explanation is in effect a special case of exaggeration, but one that is different enough to warrant separate discussion. Sometimes a case may seem so close to an ideal standard as to make it seem unreasonable or needlessly fussy to insist on the difference. It seems plausible to suppose that there are many cases where the strength of a person's justification, though not genuinely conclusive, is so close to being conclusive (or is mistakenly assessed by common sense as so close to being conclusive) that there seems to be no serious intellectual point in refraining from an attribution of knowledge. This is not a matter of being close enough for some specific, contextually relevant purpose; nor is it a matter of meeting some independently specifiable standard that is just short of conclusiveness. Instead it simply reflects a general common-sense impatience with fine hair-splitting. And this, I suggest, accounts for a significant class of mistaken knowledge attributions in cases where the justification really is overwhelmingly strong even though not quite conclusive — and perhaps especially for cases, such as very strongly confirmed inductive generalizations, where it is clearly realized that complete conclusiveness is not possible, even in principle.

(iv) Knowledge claims as conditional. A final sort of explanation, suggested by Richard Fumerton,[26] is that knowledge claims are sometimes, perhaps even often, elliptical. Someone claims to know that *P*, but the proposition he or she really thinks is conclusively justified (or, perhaps, justified in a way that closely approximates conclusiveness or is being claimed in an exaggerated way to be conclusive) is not *P* simpliciter, but rather that if various possibilities that are being assumed to be in some way not worth taking seriously are indeed false, then *P*. One common version of this is one where several possibilities are taken seriously and all but one of those is thought to be excluded by the evidence (conclusively or in one of the other ways just noted), but there are other remote possibilities that are dismissed without any real consideration. An example: Given that the last piece of cake is gone and that Ann, David, and Jennifer were the only ones besides myself in the house, and that I am sure (or close to sure or am claiming in an exaggerated way to be sure) that neither I nor

Ann nor David ate it, I claim to know that Jennifer did. I am assuming, without any real consideration, that various other remote possibilities do not hold: that the piece of cake did not evaporate in some mysterious way; that it was not entirely devoured by a horde of ants that then departed, leaving no trace behind; that a very clever but strangely motivated burglar did not sneak in and eat it; and so on. Fumerton's suggestion is that the knowledge claim in such a case is better represented as a conditional, with the denials of the possibilities that are not being seriously considered conjoined in the antecedent and the proposition contained in the original knowledge claim as the consequent.

3. Two further factors: assertion and practical reasoning

In addition to the alternatives just discussed, there are two further factors that, I suggest, help to explain the prevalence of knowledge attributions that do not accord with the strong conception of knowledge, even though (as I am provisionally assuming in this part of the paper) that conception is correct as an account of the common-sense concept of knowledge. Both of these have the effect of increasing the pressure to make knowledge claims even where the standard for knowledge is not fully met, especially if the gap is or is perceived to be relatively small. (The two factors in question have in fact often been taken, sometimes only implicitly, to support the weak conception, but I will argue that this way of construing them is less plausible, given the other problems with the weak conception, than the one I will suggest.)

(i) Much has been written recently about the idea that knowledge is the "norm of assertion": that is, roughly, that assertions are properly made only where the assertor has knowledge of the asserted proposition (so that an assertion is in effect an implicit claim to have knowledge).[27] There is something that seems intuitively right about this view: it is, after all, odd to say something of the form *P but I do not know that P*. But if I am right that the strong conception is correct and is indeed the only really intelligible conception of knowledge that we have, and assuming that justification is indeed at best rarely conclusive, we would be led to the result that assertions (without some probability qualification or the like) can almost never be properly made.

But such a result is obviously intolerable from a practical standpoint. We have urgent needs for communicating information — even uncertain, perhaps only apparent information — to each other, and it would be extremely burdensome to have to couch virtually all assertions in probabilistic terms (in addition to being almost pointless where, as is usually the case, we have no way of determining the specific degree of probability with any real precision). Thus there is a strong practical pressure to make unqualified assertions even where this is not really (fully) warranted by the epistemic situation (and to back them up, when challenged, with attributions of knowledge). Sometimes this leads to loose assertions that are unacceptable in a really patent way, but it also helps

to explain why knowledge attributions falling under categories (ii), (iii), and (iv) of the previous section are so prevalent. (Such a further explanation is, of course, not needed for category (i), where the person believes, mistakenly, that the standard for knowledge, and so for correct assertion, is fully met.)

(ii) A further connected but still distinguishable pressure in the direction of knowledge attributions that do not fully meet the standard of the strong conception is provided by the role of knowledge in practical reasoning. Much of the reason that we seek true beliefs about the world is so that we can use them in deciding what to do in various practical situations. Here it is obvious that what we really need and want is knowledge that fits the strong conception: knowledge that involves conclusive justification and so an assurance of truth. For only knowledge of this sort can provide a secure basis for determining the course of action that is really best in relation to our desires and purposes. And this makes it at least plausible that, as some have suggested, knowledge is also the norm for what can be properly used in practical reasoning.

But given our assumption that the strong conception is correct, the resulting problem is very obvious. Lacking very much genuine knowledge and still faced with the unavoidable need to deliberate and act, what are we to do? It would perhaps be possible in principle to proceed for a time in probabilistic terms, and Bayesians, among others, have tried to say how such reasoning might work. But the problems with this are obvious and overwhelming: we rarely know probabilities with enough precision to arrive at anything very definite on this basis; and in any case, we ultimately have to actually do some particular thing and not just probably do it.

Thus there is again a practical pressure in the direction of reasoning from and acting upon claims whose justification does not fully meet the standard for knowledge — and then in effect buttressing such reliance by attributing the corresponding knowledge to oneself and to others. And this again can further explain the attributions of knowledge that fall under categories (ii), (iii), and (iv) above — though once more no such further explanation is needed for category (i).

Thus, I am suggesting, the intuitive connection between knowledge, on the one hand, and assertion and practical reasoning, on the other, rather than supporting the weak conception, can instead contribute to an explanation of why knowledge attributions that are incorrect according to the strong conception are so commonly made (even though the strong conception is, as I am presently assuming, correct). The fundamental point with regard to both assertion and practical reasoning is that while the intuitive connection of each with genuine knowledge is strong, this by itself in no way guarantees or even makes especially likely that there will in fact be enough genuine knowledge to meet the practical pressure to make assertions and the even stronger practical pressure to engage in practical reasoning. Rather than inferring from the existence of obvious cases of assertion and practical reasoning to a correspondingly weak standard for knowledge that is satisfied for the relevant claims, it is, I submit, at least as

plausible in itself — and more plausible given other considerations — to suppose that the practical needs involved lead to many instances of (less than fully proper) assertion and (less than optimum) practical reasoning where the standard for genuine knowledge is not satisfied for the claims involved — and hence, given the intuitive connection, to corresponding inaccurate, exaggerated attributions of knowledge.

4. Summary and tentative assessment

My suggestion is that the elements in the previous two sections, taken together, go a long way toward explaining how the view that the strong conception of knowledge is correct might be reconciled with the range of knowledge attributions that ordinary people actually make. There are strong pressures, both intellectual and practical, that favor relatively loose and generous attributions of knowledge to others and, especially, to ourselves; and these lead to many exaggerated knowledge attributions, especially but not only where our epistemic situation closely approximates, or seems to us to closely approximate, the correct standard. They also may lead to elliptical claims of knowledge, where the claim should really be conditioned on the falsehood of a set of possibilities that are being assumed not to hold. And, it must be emphasized again, there is also a wide range of cases where though the justification is, on fairly subtle philosophical grounds, inconclusive, this cannot be assumed to be apparent to ordinary people, who make knowledge attributions accordingly.

It will be obvious that all of this is still highly schematic and impressionistic at best, needing to be filled out and amplified by a detailed consideration of many more examples than can be offered here. But the question that still needs to be asked, albeit in a tentative and provisional way, is whether this sort of explanation, thus amplified, would be good enough. Even without the needed amplification, such an explanation is surely good enough to call into question the quick, almost casual dismissals of the strong conception that are so prevalent, but is it (or can it be made) good enough to strongly support the view that the strong conception of knowledge is indeed the correct one, the one that is operative in ordinary thought?

I have to admit to a good deal of uncertainty about the correct answer to this question, and thus the result of this paper will itself be, alas, inconclusive. What I want to insist on, however, is that this uncertainty gives no real support to the weak conception, which still remains impossible to flesh out in a tenable way (or, I think, to views like contextualism). The alternative is rather that there simply is no coherent conception of knowledge to be found in common sense and ordinary thought — or at least none that has the kind of intellectual and philosophical significance that such a concept is supposed to have.

To many, this last suggestion will surely seem intolerably puzzling or even paradoxical: how could a widely used common-sense concept be simply

incoherent? But in light of the discussion here, I think we can see at least the outlines of a kind of quasi-historical explanation of how this might be so. Putting the point very roughly and impressionistically, common sense thought in this area began with the strong conception of knowledge, which is both intuitively compelling and admirably fits our practical needs. But the increasing realization that this conception is very difficult to satisfy led, not to the substitution of an ersatz weak conception that would be impossible to specify in a meaningful, intuitively satisfying way, but rather to an essentially irresolvable discord between the basic concept of knowledge and the attributions of knowledge that we want and need to make. The result is a situation in which there is simply no conception of knowledge that can make coherent sense of all the things we think and want to say about it.[28]

It will be obvious that the difference between the situation just described and the one suggested earlier (in which the strong conception is correct but widely misapplied) is anything but sharp and that the two possible situations may well shade into each other in ways that make it hard to draw clear lines. This is perhaps the main reason for my own uncertainty about which conclusion is correct.

5. *The implausibility of skepticism*

If the strong conception of knowledge is indeed the correct one, then the result is clearly a fairly strong form of skepticism — not, I think, total skepticism, but still a skepticism according to which a very high proportion of ordinary knowledge claims and attributions are in fact false.[29] As noted at the beginning of this paper, such a result is often dismissed as too implausible or even absurd to be accepted. One objection appeals directly to the ordinary uses of the terms "know" and "knowledge," and it is this objection that I have been mainly dealing with in this part of the paper.

But I want in this section to consider what seems to be a somewhat different objection, one that appeals directly to the alleged "intuition" that we do have knowledge of many things that the strong conception would seem to exclude. I have in mind here what is sometimes referred to as the "Moorean" objection, named after its perhaps most famous proponent, G. E. Moore. Moore simply asserted that we do have knowledge of such things — his famous example is his knowledge, as he holds them out in front of himself, that he has hands — and that any view that denies this must be rejected. (Part of the idea is that any premises upon which the denial of such knowledge is grounded will inevitably be less certain than the knowledge attributions themselves.)[30]

The intuition to which Moore is pointing is indeed a powerful one, not easily dismissed. But before concluding on this basis that the strong conception of knowledge must, after all, be mistaken, I think it is important to ask, more carefully than is usually done, just what the content of Moore's intuition really

is. Is Moore claiming merely that it is correct to apply the words "know" and "knowledge" to such cases, even if this is done on the basis of low epistemic standards of the sort that the contextualist claims to often apply? Is he perhaps claiming only that the justification for the claims in question, while less than conclusive, exceeds some "magic" level that can be neither specified nor explained as to its significance? Or is the intuition not rather that the justification in such cases *really is conclusive*, even though various philosophical arguments and possibilities might seem to show that it is not?

Those who have such an intuition will have to ask themselves what its content really is, but my suggestion is that it is the last of these three alternatives that is correct. And in this case, the Moorean intuition does not really count against the strong conception of knowledge and provides no genuine motivation for alternatives like the weak conception and contextualism. Instead it is a direct challenge to the philosophical grounds for thinking that justification is only very rarely conclusive. Though intuitively very powerful, it seems to me that this challenge can be met by appeal to the various arguments in question. But the important point here is that the correctness of the strong conception is not really at issue in the resulting discussion.

6. Conclusion

I have tried to show in this paper that the weak conception of knowledge is untenable and that the specific common-sense concept that is often thought to realize it is a myth.[31] My tentative conclusion is that the correct common-sense concept of knowledge is indeed a version of the strong conception, though a complete case for this result would require a much more extensive survey of examples than has been possible here. A further implication, I believe, is that focusing on knowledge and especially on preserving common-sense attributions of knowledge yields little that is of deep epistemological and philosophical interest. What this further suggests is that the really important philosophical issues pertain to justification and not to knowledge.

Notes

1. "Conception" rather than "concept," because a specific concept would have to say more precisely just what this lesser level of justification is — something that I will claim cannot satisfactorily be done. A *conception*, as I will use the term here, is thus a general outline or characterization which one or more specific concepts might fit. The weak *conception* of knowledge is, of course, perfectly real: what I am claiming to be a myth is the idea that that there is a specific common-sense concept that satisfies it. (The earliest use of the terms "weak" and "strong" to distinguish two conceptions of knowledge — or senses of "know" and "knowledge" — seems to be in Norman Malcolm's paper "Knowledge

and Belief," reprinted in Malcolm, *Knowledge and Certainty* (Englewood Cliffs, N.J.: Prentice-Hall, 1963). Malcolm's actual account of the distinction is more complicated and a bit idiosyncratic, having to do with something like the intentions of the person using the term.)

2. I will assume for the most part that the degree of justification required for knowledge does not vary from context to context (as claimed by contextualists) or from subject to subject (as claimed, in effect, by John Hawthorne, *Knowledge and Lotteries*, Oxford: Oxford University Press, 2004). I have no space here to explore these recently popular views (apart from a very brief discussion in a later footnote). But I also think it obvious that an account that avoids complexities of this sort is preferable if it can be made to work.
3. Thus the paper is mainly directed at those who share the internalist point of view — or at least agree that knowledge involves a substantial internalist element. But I believe that the main line of argument would apply with little if any modification to an externalist view like Goldman's reliabilism, where degrees of justification correspond to degrees of reliability. I also think that it would apply to many if not all alternative externalist and quasi-externalist accounts of the "third" condition for knowledge (the one that differentiates knowledge from mere true belief); but there is no room here for a discussion of these further alternatives.
4. Sometimes it is said that the justification must *entail* the truth of the claim in question, but I do not want to assume that anything that can justify belief must be capable of standing in an entailment relation (and so would presumably have to be propositional in nature). Little needs to be done to turn this conception into a specific concept: clarification of the relevant concept of truth and perhaps a specification of a requisite degree or strength of belief. But it is worth noting that this conception of knowledge assumes that the concept of justification, and the further concepts needed to explicate justification, are conceptually prior to the concept of knowledge and can be understood independently. This assumption seems to me to be correct, but I have no space for an independent defense of it here. For a contrary view, see Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), chapter 9; and for an objection to the strong conception based on this point, see Hawthorne *op. cit.*, pp. 137–40.
5. Discussions that say something close to this are legion. Included among them is my book *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985).
6. Perhaps the only philosopher whom might be viewed as having seriously tried to define the degree of justification required by the correct realization of the weak conception is Roderick Chisholm, in the three editions of his *Theory of Knowledge* (Englewood Cliffs, N.J.: Prentice-Hall, 1966, 1977, 1989), together with other works written in the same period. Chisholm's final account, in the 3rd Edition, is that a proposition p is *evident* (his term for the level of justification required for knowledge) for a subject S just in case "For every proposition q, believing p is at least as justified for S as is withholding q [that is, suspending judgment with respect to q]." But after all of Chisholm's rather painstaking discussion in these works, neither what the resulting level of justification really amounts to nor, more importantly, the rationale for that level being the right one is at all clear.

7. It could, I suppose, still be insisted that common sense somehow tacitly involves or incorporates a definite concept of knowledge that involves such a specification, even though there is no apparent way to determine what it is or how it is thus embodied. But it is hard to see why such a claim should be taken seriously.
8. Here, of course, we see one of the motives for contextualism. I will have a little to say about contextualism in a later footnote, though, as already noted, a full discussion is impossible here.
9. I would also say something similar about, for example, the Preface paradox.
10. Edmund Gettier, "Is Justified True Belief Knowledge?" *Analysis*, vol., 23 (1963), pp. 121–23.
11. This is a very simple version of such a condition; there are also more complicated versions.
12. A rather different sort of problem pertains to a much simpler anti-Gettier condition that has sometimes been offered: the requirement that it not be an *accident* (in relation to the justification) that the original belief is true. As discussed above, the idea that the truth of the relevant belief is merely an accident in relation to the justification is a good intuitive characterization of the various Gettier-type examples, so this condition also might seem to "work" in the sense of ruling out counterexamples — though it may seem rather more obviously ad hoc, amounting as it does to little more than the requirement that the case in question not be a Gettier case. But a deeper difficulty for this condition is that as long as the justification is less than conclusive, it isn't clear why it isn't *always* to some degree an accident that the belief is true. To say that the justification is less than conclusive is to say that there are some possible ways that the claim in question could have been false even though that justification obtained. But then why isn't it to at least some small degree accidental, in relation to that justification, that one of these alternatives didn't actually occur? And to say merely that the truth of the belief must not be *too* accidental, that there must be a high though not complete degree of non-accidentality pertaining to it, is to introduce a problem parallel to that which we saw concerning the degree of justification required for the weak conception, and one that seems no more tractable. (In relation to the anti-defeasibility requirement, we might also ask, somewhat more speculatively: if justification is less than conclusive, won't there always be truths which if carefully selected could defeat the person's justification — especially if it is just barely at the "magic" level? I do not know how to show that this is so, but the suggestion strikes me as quite plausible. And if it is so, then the "solution" in question, for which we could find no clear rationale anyway, doesn't seem after all to work.)
13. Without the assumption of the weak conception, there is also no need to ignore the pretty obvious possibility that the way in which truth can be accidental or fortuitous may very well be a matter of degree, rather than the all or nothing situation that an anti-Gettier condition has to assume.
14. *Op. cit.*, pp. 235–36, n. 21.
15. See Keith DeRose, "Knowledge, Assertion, and Lotteries," *Australasian Journal of Philosophy*, vol. 74 (1996), p. 570.
16. This was originally pointed out by Jonathan Vogel, in his paper "Are There Counterexamples to the Closure Principle?" in M. Roth. and G. Ross (eds.),

Doubting: Contemporary Perspectives on Skepticism (Dordrecht: Kluwer, 1990). The point is usefully elaborated by Hawthorne, *op. cit.*, pp. 3–7.

17. This is a slight variant of an example offered by Hawthorne, *op. cit.*
18. For an elaboration of such a criticism, see Dana Nelkin, "The Lottery Paradox, Knowledge, and Rationality," *Philosophical Review*, vol. 109 (2000), p. 385.
19. *Ibid.*, section 5.
20. Thus this proposal is perhaps best regarded as a modification of the justification condition rather than as a separate requirement, but this makes no real difference for the issues that matter here. (Nelkin also sometimes puts the point by saying that the person's belief must not be based on an inference from high statistical probability to truth (what she calls a "P-inference"), but it is at least not obvious that a person who adopts a claim on the basis of its high statistical probability must be making such an inference to truth, rather than simply adopting the claim because it is so very probably true.)
21. Contextualism is in a way an alternative to the weak and strong conceptions and also offers a response to the lottery paradox, but I have no space here for an extended discussion of that, to my mind, rather puzzling view. The central contextualist claim is that the level of justification required for an attribution of knowledge to be true varies from context to context (where the relevant contextual factors pertain to the person attributing knowledge rather than to the person to whom the knowledge is attributed). Thus in an ordinary context, where issues of skepticism have not been raised, the required level of justification will be rather low, low enough that many or most ordinary knowledge claims will be true. In contrast, the explicit mentioning of skeptical possibilities allegedly has the effect of dramatically raising the epistemic standard, so that a much higher level of justification, which many or most ordinary knowledge attributions will fail to satisfy, comes to be required. And the solution to the lottery paradox is roughly that the explicit appeal to statistical evidence makes it salient that at least one ticket that is very unlikely to win will in fact win and thereby has the effect of again raising the epistemic standard required for knowledge. Contextualism is not a version of the weak conception, but it does resemble the weak conception in holding that there are many cases (even though these must be contextually specified) where some less than conclusive level of justification is adequate to satisfy the requirement for knowledge (that is, for the correct ascription of "know" and "knowledge"). It thus seems to me to face the same problems, already discussed in relation to the weak conception, of saying more precisely just what those standards are and, much more importantly, of explaining why they are intellectually significant: of explaining just what epistemic difference reaching or exceeding such a standard (in the relevant context) is supposed to make and why it is supposed to matter. As far as I can tell, the only answer that is suggested by the contextualist view is simply that when the contextually determined level of justification is reached, the terms "know" and "knowledge" come to correctly apply — in effect that the "knowledge" flag can correctly be waved. And the question is why anyone should care about knowledge (or "knowledge") if no more than that is involved. (It does no good to say that the difference is that the claim in question can now be properly asserted and properly used in practical reasoning, since these facts are intuitively supposed

to be consequences of some prior, purely epistemic difference.) I believe that a similar objection can also be made to the "sensitive moderate invariantist" view suggested by Hawthorne, *op. cit.*, chapter 4.

22. Indeed, on a Moorean direct realist view, it presumably is conclusive, so that the corresponding knowledge attribution is not in fact mistaken.
23. It should be added that this general sort of explanation may also combine with the others to come: epistemological ignorance may make an exaggeration seem less extreme and so more acceptable; may make the situation seem like a close approximation to the correct standard when it is not; or may make it seem some possibilities are unlikely enough to be disregarded or that some remaining alternative has been conclusively established or ruled out, where this is not the case. (Some of the other kinds of explanation may also combine with each other in analogous ways, but I will not bother to specify these explicitly.)
24. Panayot Butchvarov, *The Concept of Knowledge* (Evanston, Ill.: Northwestern University Press, 1970), pp. 54–58. I have restated the point in terms of concepts rather than words.
25. It is also worth noting that the a view like the one being suggested here can invoke the salience of various possibilities of error, not to account for shifting standards for the attribution of knowledge or for knowledge itself (as is done, in different ways, by contextualism and by Hawthorne's "sensitive moderate invariantism"), but instead to explain when and why exaggerated knowledge attributions are likely to be withdrawn. It could also appeal in a similar way to the "practical environment" considerations discussed by Hawthorne.
26. Richard Fumerton, *Epistemology* (Oxford, Blackwell, 2006), pp. 23–24.
27. The main proponent of this view is, of course, Timothy Williamson. See his *op. cit.*, chapter 11.
28. Relevant here is Williamson's intricate, elegant, often brilliant, but to my mind also rather puzzling view of knowledge in his *ibid.*, though I have no space in this paper for a detailed discussion. Williamson rejects the idea that knowledge can be defined or analyzed along the lines of the traditional definition or in any other way. But his view seems nonetheless to amount to a version of the weak conception, in that he takes knowledge to be achieved in very many cases where the person's reasons or evidence yield only justification that is less than conclusive, including cases of perception, memory, and inference to the best explanation. What is extremely puzzling is that he then combines that view with the further views that one's body of evidence is constituted by one's knowledge and that evidence must be assigned a probability of 1. This seems to amount to saying that a claim may become part of one's knowledge on less than conclusive grounds, but then is subsequently to be treated as though those grounds were conclusive — so that its probability or likelihood is thus mysteriously ratcheted up (and then perhaps used as the basis for further justification and similar ratcheting up of other claims). Though the problem that Williamson is grappling with here is perfectly real (how to reconcile inconclusive justification with the subsequent employment of the claim in question, with no probabilistic qualification, as evidence for further claims), his solution still strikes me as incoherent: in effect, a curious and untenable amalgam of the weak and strong conceptions.

29. The view that the concept of knowledge is simply incoherent would most plausibly mean that such claims and assertions are neither false nor true, but instead make no coherent claim.
30. G.E. Moore, "Proof of an External World," reprinted in his *Philosophical Papers* (London: Allen & Unwin, 1959).
31. I have also suggested, though have not attempted to argue in detail, that contextualism is not an attractive alternative. See footnote 20.