

[This is a draft version of a talk given in Nov. 2005, with slight revisions from April 2006. Please do not treat as a definitive statement of my views.]

My question in this talk is “Does Knowledge Matter?” Before I give you my answer—which is “not in itself,” roughly—I need to explain exactly what the question means.

Think of epistemology as studying our beliefs and the process of inquiry by which we arrive at them.¹ There will be many ways of sorting our beliefs, in themselves or with reference to the inquiry that led to them. Some of these won’t be particularly interesting. No one much cares whether a particular belief is the product of an inquiry that began on a Tuesday. We do care about whether a belief arose from an inquiry that was well-conducted (whatever that means); it reflects on how good the believer is as a believer. The question is: Should we care whether a belief counts as knowledge? Here I’m thinking of ‘knowledge’ as we intuitively judge it when we’re not thinking about philosophical concerns; including the judgment that (at least many) people make that knowledge is lacking in Gettier cases. So: When we evaluate a belief, should we care whether it’s knowledge, including the avoidance of Gettier cases?

Mark Kaplan (1985) has given an argument that we shouldn’t care. Suppose someone has a Gettiered belief; it’s justified and true but fails to amount to knowledge because of a false lemma or some such. Kaplan points out that this can’t lead to any criticism of your methods of inquiry. *Ex hypothesi* your belief is justified; the conduct of your inquiry was entirely proper. Nor (though Kaplan does not emphasize this point) is

¹ Compare Pollock and Cruz’s claim (1999) that epistemology should be seen as ‘doxastology’, the study of beliefs.

your belief erroneous. Kaplan's conclusion is that the concept of knowledge does not provide a useful goal for our inquiries.²

John Hawthorne, by contrast, offers a conception of knowledge on which knowledge would be important because it is critical for our practical reasoning. On Hawthorne's conception, it is unacceptable to use p as a premise in your practical reasoning if you do not know that p (Hawthorne 2004, hereinafter KL, p. 30), and vice versa (see the Practical Environment constraint, KL p. 176).³ If this analysis holds, knowledge is obviously important. Few things are more important than whether a belief is a suitable premise for practical reasoning, and on Hawthorne's analysis that question is the question of whether the belief amounts to knowledge.

I will argue that analyzing practical reasoning will not show that knowledge is important in itself. When we consider what beliefs may be used as practical premises, there will be several different standpoints from which we may consider which premises will be acceptable. From one standpoint, it is important that a belief be true if it is to be used as a practical premise; from another standpoint it is important that it be well justified; from another standpoint it may be important that it be non-Gettiered in a certain way. From no single standpoint is it important that the premises of practical reasoning be known. Insofar as knowledge is important to practical reasoning, it is because a belief

² Kaplan points out that if we rule, as Descartes did, that no fallibly justified belief could ever count as knowledge, then knowledge can be a useful guide to our inquiries: What we know will be what we have infallibly justified. On such a conception of knowledge Gettier problems will never arise (Kaplan p. 362); this will not be the kind of knowledge at issue in this essay. Indeed, most know think that this conception of knowledge would inevitably lead to skepticism.

³ Hawthorne does not offer this conception up as motivating the importance of knowledge—he takes it for granted that the puzzles concerning knowledge are of interest in themselves (KL p. 21n49)—but it obviously could be converted into an argument for the importance of knowledge.

that amounts to knowledge will have several of these other characteristics that in themselves are important to practical reasoning.

That is not to say that we should abandon knowledge talk completely. Talk of knowledge, I will argue, is useful in much the same way as a Swiss Army Knife is. Let me explain this analogy with another analogy. This will illustrate the difference between a concept that is important in itself and a concept that is important because it comprises other concepts that are important in themselves.

Suppose that an auto magazine is rating off-road vehicles as to whether they are Colorado-Rally-Worthy (a term I just made up). To be Colorado-Rally-Worthy the vehicle must have a certain mileage per tank, a certain horsepower, a certain cargo capacity, and a certain clearance off the ground. There are two reasons that we might care about whether the vehicle is Colorado-Rally-Worthy.

The first reason is this: Suppose there is a real Colorado Rally. To win the Colorado Rally a driver must drive a certain number of miles without refueling, carrying a certain payload, going up mountainsides that require a certain horsepower, and over roads that will destroy your undercarriage if you don't have a certain clearance. Then Colorado-Rally-Worthiness is important in itself if you have any interest in competing in the Colorado Rally. If your off-road vehicle falls short of Colorado-Rally-Worthiness in any respect you might as well not enter it. There is a particular purpose for which Colorado-Rally-Worthiness is important as such.

Another reason we might care about Colorado-Rally-Worthiness is this: Suppose that there is no single such Colorado Rally, but there are many different rallies (or whatever)

for which it may be important to have a vehicle that gets many miles per tank, or that has a good cargo capacity, or many horsepower, or a high undercarriage. Then there will be no single purpose for which you require Colorado-Rally-Worthiness as such. Sometimes you may be able to do without that many horsepower; other times you may not need such a high undercarriage. (And maybe sometimes you'll wish you had a little extra something.) But we might still care about the magazine's designation of Colorado-Rally-Worthiness, might still seek out vehicles that were so designated, because it's a quick way of summing up a lot of things we do care about. We'd like to have a vehicle that has each of these positive characteristics to a certain degree. If we ask "Is it Colorado-Rally-Worthy?" we can find out four things we think about with one question.

On this second scenario, Colorado-Rally-Worthiness is what I'll call a Swiss Army Concept. There's no particular task that requires a Swiss Army Knife. Tasks require knife blades, screwdrivers, corkscrews, bottle openers, scissors, etc. Some tasks may require more than one, but no task requires that the various tools be in the form of a Swiss Army Knife. Nevertheless, Swiss Army Knives are quite useful. There is a reason why we have Swiss Army Knives instead of carrying around separate tiny knives, screwdrivers, etc.: It's much easier to carry them all around in one package. Similarly, there's a reason why we might care about Colorado-Rally-Worthiness even if it isn't necessary for any particular task that might be accomplished. It's easier to answer "Is this Colorado-Rally-Worthy?" than "What is its horsepower, cargo capacity, undercarriage, and mileage per tank?" The Swiss Army Concept is a concept that is not important in itself, but that provides an economical way of summing up several other concepts that are important in themselves.

So my claim will be that knowledge is a Swiss Army Concept, at least with respect to value for practical reasoning. There are actually several ways for a belief to be a good premise for practical reasoning. If, for some concept C, a belief is (in some way) a good premise iff it is C, then C is important in itself for practical reasoning. Knowledge will not be important for practical reasoning in this way, but ascribing knowledge is an economical way of ascribing several concepts that are important in this way.

To illuminate the multiple ways in which a belief can be a good practical premise, let us look at Hawthorne's account of the lottery problem. The lottery problem is this: We are generally unwilling to ascribe advance knowledge that a particular ticket in a fair lottery will not win, but we may be willing to ascribe knowledge of propositions that entail that this ticket will not win. We may be willing to say that you know that you will not be able to afford to go on an African safari next year even though you own a ticket for a lottery whose prize is more than the cost of a safari.⁴ The lottery problem can stand in for much reasoning about the not quite certain future or present. To use some examples of Vogel's (1990), we may be willing to say that you know where your car is but unwilling to say that you know it is not one of the few cars stolen each day; we may be willing to say that you know where you will be next week but unwilling to say that you will not be one of the few apparently healthy people who will suddenly drop dead before then. Hawthorne's view is that we can explain these ascriptions by defining knowledge in terms of suitability for practical reasoning.

⁴ Compare Lewis's (1996) example of Bill, who we know will never be rich because he spends all his money on lottery tickets.

The idea is this: Define knowledge so that a belief that p does not amount to knowledge in a certain practical environment iff it is not acceptable to use the belief as a premise for practical reasoning in that environment.⁵ Then, in the practical environment in which it is relevant, you will know that you won't be able to afford the safari; and you won't know that your ticket won't win the lottery in the practical environment in which *that* belief counts. In fact, in such an environment you wouldn't know that you can't afford a safari either. Thus Hawthorne's account seems to explain our judgments of knowledge while providing an important role for knowledge in practical reasoning.

To flesh out Hawthorne's argument, consider a practical environment in which you might want to use as a premise the belief that your lottery ticket won't win. Someone offers you a ticket for a 10,000-ticket lottery with a \$5000 prize, at the price of one penny. Let us suppose that, in fact, it will turn out that this ticket loses. Still, you shouldn't reason as follows:

[Argument A] (1) If I buy this ticket, it will lose.

(2) So I'll be out a penny.

(3) So I won't buy the ticket.

This is terrible reasoning; the penny cost is worth the chance that your ticket won't win.

Similarly, Hawthorne points out that it would be "intuitively awful" to reason as follows:

[Argument B] (4) I will not have enough money to go on an African safari next year.

(1) So if I buy the lottery ticket I will lose.

(3) So I should not buy the lottery ticket (KL p. 174; my numbering).

⁵ The Practical Environment Constraint (KL p. 176) provides one direction of the biconditional; Hawthorne provides the other direction on (KL, p. 30). Thanks to Mylan Engel for pointing this out.

Accordingly (1) and (4) are not acceptable premises in this practical environment, and in this environment you know neither that the ticket won't win nor that you won't be able to afford the safari.

Note that, if anything is wrong with these arguments, it is with their premises ((1) for argument A, (4) for argument B). Once we accept the premises, the conclusion follows much as the conclusion of a deductively valid argument follows from its premises. I will consider only arguments with this property, which I'll call *formal acceptability*, the better to focus on the epistemic properties of the premises.⁶

Consider now a practical environment in which you might want to exploit your belief that you won't be able to afford a safari in a more natural way. You have bought the lottery ticket, and you are now in a bookstore buying a guidebook for next year's vacation. Hawthorne argues that it is acceptable to reason as follows:

[Argument C] (4) I will not have enough money to go on an African safari next year.

(5) So I will have no use for a guidebook to Africa.

(6) So I should buy the local destination guide (see KL p. 177).⁷

⁶ Sometimes good practical reasoning may not be formally acceptable, so defined. For instance, one might decide not to buy a guidebook to Africa based on the premise "Probably, I will not be able to afford a safari," even though (4) does not follow deductively and hence the argument to (5) is not formally acceptable. Such arguments may perhaps be treated as enthymematic formally acceptable arguments, with the missing premise "This time, what will probably happen will happen." That premise is akin to a lottery premise, which on Hawthorne's view can be known in an appropriate practical environment. (See note 8, below.)

⁷ It may that, on Hawthorne's account, argument C is proper only if you have forgotten that you have a lottery ticket; if you are thinking about your ticket then you are not in a position to know (4). On the other hand, Hawthorne opposes an account on which invoking (1) in the argument from (4) to (3) destroys knowledge of the premise by making a new possibility salient; he argues that the possibility might not become salient for someone who is sufficiently dogmatic. So Hawthorne might not think that knowledge of (4), if you do have it, is destroyed when you remember the existence of the ticket. Many other ins and outs of this debate are discussed in KL.

Accordingly, on Hawthorne's account, you do know (4) on this occasion. You can know propositions about the future without ruling out lottery-like alternatives, so long as the decisions you are making do not require you to take those alternatives into account.

So Hawthorne argues that our intuitive judgments of knowledge line up reasonably well with cases in which the subject's belief is an acceptable premise for practical reasoning. (There are many complications to this view, but we can leave them aside.)

The question, however, remains: What is it for a premise to be acceptable? When we look more closely at this question, we will see that there is no way of asking it such that the acceptable premises are exactly the known ones.

Here is one possible answer: We care about whether practical reasoning will turn out well for us. So formally acceptable practical reasoning from acceptable premises should turn out well for the reasoner. But in the practical environment in which you have been offered the lottery ticket, the reasoning that will in fact turn out the best for you is the reasoning that leads you to decline the ticket. *Ex hypothesi* the ticket will lose, and if you bought it you would have been out a penny. This produces the uncomfortable result that arguments A and B are both acceptable arguments, and (1) and (4) are both acceptable premises. You can reason from the premise that your ticket will lose or that you will not be able to afford a safari. In fact, the premises that actually yield the best results given formally acceptable practical reasoning are all and only the true premises.⁸ This line of thinking shows that for practical reasoning it is important to have true beliefs.

⁸ With the possible exception of statements about probabilities. Suppose that the ticket in fact wins. Argument A has a false premise, (1), "If I buy this ticket, it will lose" is false, and as expected it does not yield the best result since in fact you would be better off buying the ticket. But if we replace (1) with the premise "If I buy this ticket, it is very likely to lose," we have an argument with a true premise that will not yield the best result.

If we want to avoid the uncomfortable result that arguments A and B are acceptable, it is obvious what we should do. We should say that when we ask whether practical reasoning is acceptable, we are not asking about the practical reasoning that will in fact lead to the best outcome. From this standpoint, we view acceptable practical reasoning as reasoning that is not vulnerable to criticism, that is not feckless or rash or overcautious.⁹

From this standpoint, arguments A and B clearly are vulnerable to criticism. The payoff for winning the lottery is so much higher than the cost of the ticket that you are not entitled to ignore the small chance that the ticket will win. So (1) and similarly (4) are not acceptable premises in this practical environment. Arguments A and B may be criticized even if they in fact turn out to save a penny. In the practical environment of the bookstore, however, you are entitled to use (4) as a premise. It would be feckless to

This argument, however, will not be formally acceptable, as explained in note 6 above. It stands to the formally acceptable argument A as a probabilistic argument from “Probably p” to q (where p entails q) stands to the deductively valid argument from p to q. To make the inductive practical argument formally acceptable, we would have to add a suppressed premise “This time, what will probably happen will happen,” and that premise would be false in the case where the ticket wins.

Other practical reasoning with probabilistic premises will not be enthymematic for formally acceptable arguments. Take the argument from “There’s a small chance that this ticket will win” to “So I’ll buy it, just on the off chance.” This argument will not lead to the best possible outcome, even though the premise is true and it does not rely on a suppressed premise “This time, what will probably happen will happen.” In fact, there is no premise that could be added to the argument to guarantee a good result; “This ticket will win” would guarantee a good result if it were true, but the person who buys a lottery ticket because the odds are favorable is not counting on the assumption that the ticket will win.

Thus, for probabilistic premises, the important property from the standpoint of actual success is not truth. In fact from this standpoint it is never ideal to reason from a probabilistic premise, because one will always be at least as successful reasoning from true non-probabilistic premises. It might be thought that reasoning from true probabilistic premises is appropriate from the standpoint that is concerned with whether one is reasoning rationally (discussed below). If the notion of probability at work here is objective, this will not be true; it may be irrational to reason from the premise that 666 will probably win the lottery even when the objective probability that 666 will win is high (because the lottery is fixed). If probability is defined as probability based on the applicable evidence one has, then it may be true by definition that reasoning from true probabilistic premises is rational. But even if this standpoint establishes the importance of justified true belief in probabilistic propositions, it will not help establish the importance of the concept of knowledge in non-probabilistic propositions. This concept of knowledge will not much resemble our ordinary concept on which we can know non-probabilistic propositions such as “I will not be able to afford an African safari.” (Thanks to A.P. Martinich for a helpful conversation on this topic.)

⁹ Compare Hawthorne’s discussion of the reading of ‘should’ on which it is obvious that a premise like (1) should not be used in practical reasoning with the possible reading of ‘should’ on which what you should have done is what would in fact have led to the best outcome (KL p. 175n33).

refuse to buy the local guidebook because you claimed not to know that you wouldn't be able to afford a safari; this is not the sort of decision that should be thrown into doubt because of a lottery ticket. So this standpoint yields the result that Hawthorne desires: (4) is an acceptable premise in the practical environment of the bookstore but not of the lottery purchase.

The problem is that from this standpoint (4) is *always* an acceptable premise in the practical environment of the bookstore. It is acceptable even when it is false. Suppose that, in the bookstore, you refuse to follow argument C because of the remote chance that you might win the lottery, and then you do go on to win the lottery. Your original reasoning would be as feckless as ever; it would be through luck that your faulty reasoning produced the best outcome for you. Conversely, suppose you reason as in argument C, buy the local guidebook, and go on to win the lottery. Was your original reasoning acceptable? From this standpoint, yes. If argument C is beyond criticism in the case in which you don't win the lottery, it is beyond criticism in the case in which you do. You were not being feckless or dogmatic in thinking that you would not be able to afford a safari. That the right reasoning did not lead to the best outcome in this case is simply epistemic bad luck (though financial good luck).

From the standpoint that concerns itself with whether your practical reasoning can be criticized, what is important for practical reasoning is how well justified your beliefs are. The practical environment matters here: It determines how much justification you need for your belief to be acceptable. Nevertheless, this standpoint does not establish the importance of a factive property of beliefs. Unless the practical environment calls for absolute certainty, it will be the case that acceptable reasoning may proceed from false

premises. *A fortiori*, acceptable reasoning may proceed from premises that are not known. So whether a belief counts as knowledge is not important in itself from either standpoint.

The argument concerning the lottery case can be applied to any practical reasoning that calls for an instantaneous decision. If it is important that the subject's reasoning in fact leads to the best outcome, we should be concerned about whether her premises are true. If it is important that the subject's reasoning be beyond criticism, we should be concerned about whether her premises are well enough justified given her practical situation.

Most decisions, however, are not instantaneous. To accomplish anything we need to be able to make a plan and carry it out over an extended period. In such a case success requires more than just having a true belief at any one point. So when we look at practical reasoning over an extended period of time, properties of the belief other than its truth and justification may be important.

Consider this example of Williamson's (KIL, p. 62): A burglar is ransacking a house looking for a diamond. He knows that there is a diamond in the house, so he continues to look all night even when he fails to find it. If, on the other hand, he had a Gettiered belief that there was a diamond in the house, he might not continue to look all night. Suppose that he inferred that there was a diamond in the house because he had been told that there was one under the bed, when in fact the diamond was in the drawer. Then he would give up after failing to find the diamond under the bed. He has a justified true belief when he

sets out to look for the diamond, but it will not be enough to keep him looking long enough to have a good chance of finding it.

This example might be taken to show that in some cases knowledge *is* the important concept for evaluating practical reasoning.¹⁰ But looking at extended plans will still not reveal any one from standpoint from which the acceptable premises for practical reasoning are exactly the things we know. There is still a split between premises that will actually lead to successful reasoning and premises that lead to reasoning that is beyond criticism, and from neither perspective is knowledge what we need.

Consider what it takes for a belief to serve as a premise for practical reasoning that will produce a plan that actually succeeds. The burglar might reason as follows:

[Argument D] (7) There is a diamond somewhere in this house.

(8) If I burgle the house, I'll get the diamond.

(9) So I'll burgle the house.

For argument D to be formally acceptable, it must be right for the burglar to want to find the diamond (so we must discount the wrongness of burglary itself!), and the value of the diamond must outweigh the costs of burgling the house. Also, both (7) and (8) have to be acceptable premises. We can suppose that (8) entails that the burglar won't get caught (if you like, make it "I'll get the diamond and won't get caught"), so among other things this entails that it is an acceptable premise for the burglar that he won't get caught while he's looking for the diamond. I will focus on the acceptability of (7), leaving aside the issue of when (8) is justified.

¹⁰ Williamson himself uses the example to argue that "The burglar knew there was a diamond in the house" has more explanatory power than "The burglar had a true belief that there was a diamond in the house."

Consider two burglars, Moriarity and Raffles. Both initially have a true belief in (7), so both initially follow argument D. This initial reasoning, however, is not enough to ensure that they get the diamond. The successful burglar must preserve his belief in (7) until he actually finds the diamond. So if Moriarity gives up his belief in (7) once he has looked under the bed, and Raffles will preserve his belief in (7) until he finds the diamond, the Raffles will find the diamond and Moriarity will not. What is important is not only truth but stability of belief.

Note, however, that Raffles will find the diamond if his belief is stable, no matter how unjustified he may be in preserving his belief. Suppose that Raffles' informant told him that the diamond was under the bed, but Raffles persisted in believing that there was diamond somewhere even after he has found that it is not under the bed; or suppose that Raffles never had any justification whatsoever for his belief, but was determined to stick to it as long as possible. It will not matter from this standpoint. So long as Raffles' belief is true and persistent, his plan will actually succeed.

From the other standpoint, suppose that we are concerned with whether the burglars' reasoning can be criticized. To consider whether it can be criticized at the outset is no more than to consider how well justified their beliefs are, as discussed in the previous section. To take the extended view, we should ask whether their reasoning yields a plan that can be completed without exposing them to criticism. This will be so if they not only are justified in believing (7) at the outset, but will remain justified in believing (7) for the duration of the plan. Suppose that Raffles has been told that there is a diamond in the house, and Moriarity has been told that there is a diamond in the house and it is definitely

under the bed.¹¹ Then Raffles ought to follow argument D and stick to his plan of searching the house until he finds the diamond. To do otherwise would be to give up too easily. Moriarity, on the other hand, ought to follow argument D at first; but when he finds no diamond under the bed, he ought to abandon it. His reason to believe its premise (7) has been undercut. To keep searching for the diamond would be stubborn.

From this standpoint, then, an acceptable premise is one that is justifiably believed and that is likely to stay justifiable as new evidence is overturned. This rules out some Gettier cases, as in the burglar who has been told that there is a diamond and it is definitely under the bed. It does not, however, rule out all cases of false belief. Suppose that there is no diamond in the house. For Raffles, premise (7) will remain justified for as long as it takes to search the house; so Raffles' premise will remain beyond criticism until he completes it. Then it will be seen to have failed, through Raffles' bad luck, but from this standpoint an acceptable premise is not necessarily one that will lead to a successful plan. This standpoint on practical reasoning shows stable justification to be important in itself, not knowledge.

There is one standpoint from which it is important to have a belief that is true and also has a justification-like property. We can ask: If we use this belief as a practical premise, will the resulting reasoning succeed, and would it still have succeeded if circumstances had been different? From this standpoint, an acceptable premise is not only true but counterfactually true; where, plausibly, by counterfactually true we mean

¹¹ The idea is that once Moriarity finds that there is no diamond under the bed he should come to doubt his original information. In Williamson's original case it may be that the burglar, finding no diamond under the bed, should conclude that it was moved and that he should look elsewhere in the house. We can stipulate that that is not the case in the current situation.

that whatever the agent believed would be true under different circumstances, not that this actual belief would have been true under different circumstances. So if success depends on the agent's successfully identifying a random animal that is produced for her, and she reasons "That's a duck, so I'll say it's a duck," her premise "That's a duck" could easily have been false (if a chicken had been produced instead), but she would still have been reasoning from a true premise (because then she would have reasoned from "That's a chicken").

Hence, when we are concerned with the counterfactual success of practical reasoning, the important property of belief is safe truth.¹² Williamson (KIL p. 123ff.) and Ernest Sosa (2000) have both posited safe truth as a requirement on knowledge. Nevertheless, it does not seem to me that this standpoint establishes the importance of knowledge in itself for practical reasoning.

For one thing, safe truth is arguably neither necessary nor sufficient for knowledge. Against sufficiency: Knowledge also needs some sort of internal justification; the beliefs of BonJour's clairvoyant (BonJour 1985, ch. 3) are safe, since if the proposition that he believes because of his clairvoyance were false he would not believe it.¹³ Similarly, there may be cases of safe true belief that are nevertheless Gettier cases. Suppose that, in the barn-façade case (Goldman 1976), although the county is full of barn-façades, some law

¹² If we define 'counterfactual success' so as to mean "success in every world up to the nearest world in which the current reasoning would not be successful," then instead of safe truth the important property will turn out to be sensitive truth, as in the analysis of knowledge in (DeRose 1995). But it is hard to see why that particular definition of counterfactual truth would be of practical interest, unless we already assume that practical premises should be epistemically sensitive. And if we make that assumption, we have not succeeded in grounding the importance of knowledge in its role in practical reasoning.

¹³ The clairvoyant's belief is modally unstable in this way: If he were not clairvoyant but merely believed things that popped into his head, his beliefs would be false even though things would seem the same to him. But if worlds in which an actually reliable faculty fails count as close for the purpose of safety, then it seems as though safety must be a matter of internal justification; otherwise worlds in which perceptual faculties fail could also count as nearby, and perceptual knowledge would be considered unsafe.

or physical fact makes it impossible to build a barn- façade in Yoder's field. Yoder's field is so swampy that a barn-façade could not stand up without the additional support provided by a real barn's other walls. Then the belief that Yoder's field contains a barn is arguably safe; in the nearby worlds in which I believe there is a barn in Yoder's field, there is a barn in Yoder's field. Still, the nearby barn-façades keep us from knowing that there is a barn in Yoder's field. Perhaps safety can be defined so that the possibilities of error are near enough to make the beliefs unsafe. We might say that "There is a barn in Yoder's field" belongs to a class of closely related propositions ("There is a barn in Stoltzfuss's field," etc.), such that in nearby possible worlds I falsely believe one of them in place of "There is a barn in Yoder's field." However, there will be a risk of gerrymandering the concept of safety to fit the concept of knowledge.

Against the necessity of safety for knowledge, see Comesana (2005) and Neta and Rohrbaugh (2004). These authors present interestingly different diagnoses of why knowledge can fail to be safe. Comesana argues that safety requires reliable reliability, whereas knowledge merely requires reliability. This suggests that safety provides additional value over and above whatever values are folded into knowledge ascriptions, for surely reliable reliability is even more valuable than mere reliability. Neta and Rohrbaugh argue that knowledge is an important cognitive achievement, and that earned achievements are in general not safe; an achievement that is earned despite the possibility of failure is nevertheless earned (p. 404).

This leads to the more important problem with the idea that knowledge *qua* safe true belief would be important from the standpoint of practical reasoning. It is not clear why safe true belief would be more practically valuable than true belief alone. An

achievement is earned even if it is not earned safely, so it may not matter overmuch that our reasoning be not only successful but counterfactually successful. Surely what we care about most is success in the actual world. If we do point out a flaw in a premise that actually led to success, it is more relevant if the agent could have been criticized for relying on the premise than if the premise could have failed in some counterfactual situation. (Of course these often go together.)

In addition, the safety of the premises does not seem to have anything to do with the practical environment. Premise (4) seems just as safe in the context of Argument B as in the context of Argument C. Whether the decision at issue concerns buying a lottery ticket or a guidebook, the possibility that you can afford the safari is just as remote. We might escape this consequence by defining remoteness in terms of the possibilities that you ought to take into account, given the practical decision you are making; but this definition could only be motivated from a standpoint concerned with the rationality of the practical reasoning, rather than its actual and counterfactual success.

Another way to try to bring truth and justification together is to argue that, in order to be immune from criticism, practical reasoning must proceed from true premises. Practical reasoning from false premises, or even from premises that are not known, is as such open to criticism. Williamson makes an argument that can easily be extended to yield this conclusion: Our evidence is identical to our knowledge, and rationality requires respecting the evidence (KIL ch. 9, revising Williamson 1997). If this were true, then reasoning whose premises were not known would always be rational. Williamson bolsters this position by arguing that we cannot always know what is rational for us.¹⁴ If

¹⁴ Williamson directs the argument toward theoretical reasoning, but it also will apply to practical reasoning.

I see a table before me, it is rational for me to believe in its existence, because it is part of my evidence that I see (and know) that there is a table before me. If I am hallucinating a table, then it is not part of my evidence that I see and know that there is a table, so it is not rational for me to believe in a table. Yet the hallucinator may be internally indistinguishable from the person who sees the table.

Yet we may grant Williamson's analysis of the hallucination and similar cases without accepting his argument that rational belief always proceeds from known or even true premises, if we are not antecedently committed to the importance of knowledge. Consider argument C, the bookshop argument, as it is made by two exact duplicates, one holding a winning ticket and one holding a losing ticket. Is there a sense in which the loser is deliberating rationally and the winner is not? The only difference between them concerns the result of the lottery drawing, an event that will take place after they have acted on their deliberations. In the hallucination case we could argue that the person who saw the table has direct access to the existence of a table, which the hallucinator lacks, and that this distinguishes their reasoning. But in the lottery case it is implausible that the loser has any direct access to the fact that she will not be able to afford a safari.¹⁵ There needs to be some other relevant difference if we are to conclude that the loser deliberated rationally and the winner did not.

It is true that the loser's deliberation begins from a true premise and the winner's deliberation from a false one. But this will not provide a basis for criticizing the winner's deliberation and not the loser's. If we criticize the winner for reasoning from a false

¹⁵ Though I do not have space to explore the possibility here, this case also calls into question Williamson's argument that all our knowledge serves as evidence; if we have inferential knowledge about the future, that will be a good candidate for knowledge that is not itself evidence. (See Brian Weatherson's discussion of inductive knowledge as a counterexample to the knowledge-as-evidence thesis, at <http://tar.weatherson.net/archives/000644.html>.)

premise, she may say, “Yes, but I had every reason to believe it was true. Should I instead have reasoned from the true premise that I would be able to afford a safari? That would have led to a better outcome, but it would have been bad reasoning.” The actual truth of her premise is irrelevant from the standpoint of criticizing her deliberation; to make it relevant, we must adopt the standpoint from which we care about the actual success of her deliberation. And as we have seen, from that standpoint justification does not matter. Similarly, to insist that deliberation is not rational if the premises are not known is to assume the value of knowledge for practical reasoning. It will not help us use practical reasoning to establish the importance of knowledge.

We have seen that, from various standpoints on practical reasoning, it is important that one’s premises have various properties: truth, justification, persistence of belief, stability of justification, safety. There is no single standpoint from which it is important that one’s premises be known. We might wonder, then, if knowledge is not important for practical reasoning, why do we talk about knowledge at all when we are concerned with practical matters? Why not simply talk about the things that are important? Yet in fact people talk about what people know much more often than about what they are justified in believing, for instance.¹⁶

To answer, think of the Swiss Army Knife metaphor. A Swiss Army Knife is useful to carry around when you do not know exactly what task you will be faced with. If you are faced with a task that requires a knife, a screwdriver, a corkscrew, or a bottle opener, you will have what you need; and you will not face the awkwardness of having to carry

¹⁶ One of the criticisms that may be made of the argument of (Kaplan 1985) is that it leaves it mysterious why anyone would ever have thought knowledge important (in its contemporary guise, in which it can be based on fallible evidence).

around four separate tools. Analogously, when evaluating someone's epistemic situation you may not want to know which standpoint you will eventually want to take on their practical reasoning. If you say "S believed truly that p," and it becomes important to figure out whether S's reasoning should be criticized, then you won't have said anything helpful. So it will be useful to have a quick way of expressing all these different concepts that may be useful for the different standpoints on S's practical reasoning. If you say "S knew that p," your audience knows that S should not be criticized for reasoning from p, and S's reasoning from p would actually succeed if it was formally reliable, and (if applicable) that S was in a position to retain her belief that p long enough for her plans to succeed, etc. Even though the concept of knowledge is not needed for answering any one of these questions, it provides an efficient way of expressing an answer to all of them.

This provides a little bit of progress toward the question about why we use a concept of knowledge that rules out Gettier cases. In some Gettier cases, though a belief is justified and true, it does not count as knowledge because the belief or justification is unstable. If, like Williamson's burglar, you are about to discover countervailing but misleading evidence, your belief will not remain immune from criticism long enough that your plan will remain immune from criticism throughout its execution. If you are about to forget a belief, you may not be able to carry out any plans based on it.¹⁷ When we call a belief knowledge, we guarantee it satisfactory from these standpoints.

But this is not quite enough as an account of why we do not count Gettier cases as knowledge. The most important standpoints on practical reasoning are surely whether it leads to actual success and whether it is immune from criticism. So the most important

¹⁷ Marc Moffett has devised examples where temporal instability might make us reluctant to ascribe knowledge; see http://rationalhunter.typepad.com/close_range/2004/05/a_thought_exper.html.

properties of beliefs from a practical standpoint will be truth and justification. If the importance of knowledge derives mostly from the importance of justification and truth, why are there so many cases in which we judge that justified true belief is not knowledge?¹⁸

The answer, I think, comes from extending the Swiss Army Knife metaphor. If we become used to Swiss Army Knives as the way to carry around the tools we need, we may come to see them as valuable in themselves (even though they are not). Then someone who has a knife, a screwdriver, a corkscrew, etc. may still be seen as lacking something important. We will have brought ourselves to care about not only the individual tools, which are what we really need, but also about how they are connected. Separate tools will not seem as satisfactory as the same tools in a single package.

Analogously, when we ascribe knowledge it seems to be a unified concept that may be important for its own sake. Even if knowledge is important primarily because of the importance of truth and justification, it seems as though what is important is that the truth and justification be combined in the right way. Typically a justified true belief is one in which whatever makes it justified is also whatever makes it true. So this will seem to be characteristic of knowledge. A belief that lacks this characteristic, in which justification and truth are somehow mismatched, will be seen as lacking the organic unity that typical knowledge has. Even if justification and truth are independently important, from different standpoints, when we use a single word to ascribe them together what seems important is that they come together in the right way. Hence when justification and truth are mismatched we will have a Gettier case, where we are reluctant to ascribe knowledge.

¹⁸ Weinberg, Nichols, and Stich (2001) have cast doubt on the universality of Gettier intuitions, but what requires explanation is why anyone at all has strong Gettier intuitions.

But in fact there is no standpoint from which the mismatch of knowledge and justification is particularly important for practical reasoning, except insofar as it subverts the temporal or modal stability of belief, truth, or justification. Truth, justification, and stability will be important in themselves for various ways of looking at practical reason. Knowledge is important for practical reason only insofar as it combines these other qualities; and when we demand that these qualities be brought together in an organic whole, this demand does not yield anything that we need for a belief to be a good practical premise.

BonJour, Laurence (1985). *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard University Press). Comesana, Juan (2005). "Unsafe Knowledge." *Synthese* (2005) 146: 395–404.

DeRose, Keith (1995). "Solving the Skeptical Problem." *Philosophical Review* 104, 1-52.

Goldman, Alvin (1976). "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73, 771-791.

Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford University Press, Oxford.

Kaplan, Mark (1985). "It's Not What You Know That Counts." *Journal of Philosophy* 82, 350-363.

Lewis, David (1996). "Elusive Knowledge." *Australasian Journal of Philosophy* 74, 549-567.

Neta, Ram, and Guy Rohrbaugh (2004). "Luminosity and the Safety of Knowledge." *Pacific Philosophical Quarterly* 85, 396-406.

Pollock, J. and J. Cruz. (1999). *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman and Littlefield.

Sosa, Ernest (2000). "Skepticism and Contextualism." In J. Tomberlin (ed.), *Philosophical Issues* 10, 1-18.

Vogel, Jonathan (1990). "Are There Counterexamples to the Closure Principle?" In *Doubting: Contemporary Perspectives on Skepticism*, ed. M. Roth and G. Ross. Kluwer, Dordrecht.

Weinberg, Jonathan, Shaun Nichols, and Stephen Stich (2001). "Normativity and Epistemic Intuitions." *Philosophical Topics* 29, 429-60.

Williamson, Timothy (1997). "Knowledge as Evidence." *Mind* 106, 717-741.

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press, Oxford.