

- Williams, B. (1981) *Moral Luck*. Cambridge: Cambridge University Press.
- (1995) 'Internal Reasons and the Obscurity of Blame', in Williams's collection *Making Sense of Humanity and Other Philosophical Papers*. Cambridge: Cambridge University Press, 1995.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

7

Practical Reasoning and the Concept of Knowledge

Matthew Weiner

Epistemologists have devoted much more attention to the question 'What is knowledge?' than to the question 'Why does knowledge matter?' There is a tremendous amount of work that attempts to determine which beliefs count as knowledge, less that argues that we should care whether our beliefs count as knowledge. Yet without a positive answer to the second question, the first question lacks interest.

The problem is particularly acute in the post-Gettier area. In the *Meno* Plato addresses the question of why knowledge is more valuable than mere true belief. But after Gettier, an account of the value of knowledge must explain something stronger: why knowledge is more valuable than mere *justified* true belief. As Kvanvig puts it, the question 'is whether and how knowledge has a value exceeding the sum of its parts' (Kvanvig 2003: p. x).

In this essay I will consider a version of the view that knowledge is valuable because of its role in practical reasoning. This view is based on a conception of knowledge presented by John Hawthorne (2004), on which it is constitutive of knowledge that what is known can serve as a premise for practical reasoning. Hawthorne's view seems to provide a simple account about why we should care about whether we know ordinary statements (as opposed to probabilities, which I will discuss in Section 5). I will argue, however, that adverting to practical reasoning will not establish that knowledge is valuable in itself. This may seem a minor result, refuting a single argument for the significance of knowledge (and not one that Hawthorne himself puts forward). Yet I will argue that considering this view illuminates the true nature and significance of knowledge. When we examine which beliefs are suitable practical premises, we see that knowledge is not valuable in itself over and above the value of its components. Nevertheless the concept of knowledge is valuable, for it gives us an economical way of summing up many properties of beliefs, each property valuable in itself.

In Section 1 I explain the conception of knowledge that I will argue for in terms of a Swiss Army Knife, which has no value in itself over and above the

value of its components, but which is nevertheless valuable in that it provides an economical way to carry its components around. In Section 2 I begin to consider the argument that knowledge is valuable because the propositions that are known are the propositions that are acceptable premises for practical reasoning. I take up different standpoints on practical reasoning in order to show that there is no single standpoint from which the acceptable premises for practical reasoning are exactly those propositions that are known; from one standpoint the acceptable premises are exactly those propositions that are true, from another standpoint the acceptable premises are exactly those that are justified. Sections 3 and 4 consider further standpoints; though, from these standpoints, acceptable premises have some property that may go beyond mere justification or mere truth, still there is no single standpoint from which the acceptable premises are exactly those that are known. From each standpoint a different property of beliefs is shown to be valuable. Section 5 considers reasoning from probabilistic premises; it may be that the acceptable probabilistic premises are exactly those probabilistic propositions that are known, but this will not establish the value of knowledge in any way that will satisfy proponents of the traditional conception of knowledge, because on the conception of knowledge motivated by this view of practical reasoning we only have knowledge of probabilities. Section 6 returns to the Swiss Army Knife metaphor, in order to explain the utility of a conception of knowledge that integrates the properties of beliefs that have been shown to be valuable from the various standpoints on practical reasoning, and to account for why we might have developed a conception of knowledge on which Gettier cases do not count as knowledge.

1

On the conception of knowledge Hawthorne (2004) proposes, it is unacceptable to use p as a premise in your practical reasoning if and only if you do not know that p .¹ If this conception 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.

Put another way: If we solve for X in 'A belief is a suitable premise for practical reasoning iff the belief has property X ,' then X is certainly a valuable epistemic property. But is X knowledge? I will argue that we may take several

¹ The Practical Environment Constraint (2004: 176) provides one direction of the biconditional; Hawthorne provides the other direction on 2004: 30. 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 (2004: 21 n. 49)—but it obviously could be converted into an argument for the importance of knowledge. (Thanks to Mylan Engel for pointing out Hawthorne's explicit statement of both directions of the biconditional.)

different standpoints on practical reasoning. Each of these standpoints yields a different solution for X , and each different X is thereby a property that is epistemically valuable in itself. But none of these X s is knowledge; there is no single standpoint from which the suitable practical premises are those that are known.

This is not to say that the concept of knowledge has no value. Talk of knowledge is useful because of knowledge's relation to these epistemically important concepts. The concept of knowledge is analogous to a Swiss Army Knife, which is never needed qua Swiss Army Knife, but which is useful because it contains many individual blades that may be needed for their own sake. As the Swiss Army Knife provides an economical way of carrying blades that are valuable in themselves, so the concept of knowledge provides an economical way of expressing several other concepts that are valuable in themselves.

I will explain the Swiss Army Knife analogy with yet another analogy, which shows how a concept may be useful as an economical way of expressing other concepts. Let us consider two different scenarios in which an auto magazine might rate off-road vehicles as to whether they are Colorado-Rally-Worthy (my own imaginary term). In each scenario Colorado-Rally-Worthiness requires a certain mileage per tank, a certain horsepower, a certain cargo capacity, and a certain clearance off the ground. But only in the first scenario is Colorado-Rally-Worthiness valuable in itself; in the second scenario it is valuable only because of the value of its components.

The first scenario is this: There is one major road rally in Colorado. 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 want to compete in the Colorado Rally. If your off-road vehicle falls short of Colorado-Rally-Worthiness in any respect, you might as well not run the race. There is a particular purpose for which Colorado-Rally-Worthiness is important as such.

The second scenario is this: Suppose that there is no single such Colorado Rally, but many different road rallies. In different rallies it is 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 there are so many rallies, possibly indefinitely many, that it is not practical to look for a vehicle that is suited to every single one.) But we might still care about whether our vehicle is one that the magazine calls Colorado-Rally-Worthy. We might seek out vehicles that are so designated because the designation is 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?' that one question tells us four things we want to know.

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 tool, 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; it provides an economical summation of the concepts that turn out to be important for practical reasoning from the various standpoints.²

2

To illuminate the multiple standpoints on practical reasoning, 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.³ 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

² I suspect that knowledge will turn out to be a Swiss Army Concept for other purposes as well. For instance, Kaplan (1985) argues that knowledge is not important to inquiry, because once we have determined that a belief is justified we have settled that the believer's inquiry is exempt from criticism. We might take this as one standpoint on inquiry. From another standpoint, we might care about whether the inquiry is successful; and then we will care about whether the knowledge results in true belief. These will be seen to resemble the first two standpoints on practical reasoning that we will discuss, with knowledge providing an economical combination of truth and justification.

³ Compare Lewis's (1996) example of Bill, who we know will never be rich because he spends all his money on lottery tickets.

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.

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. 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 is relevant. 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 \$5,000 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 lottery ticket, it will lose.
- So,
- (2) I'll be out a penny.
- So,
- (3) I should not buy this ticket.

This is terrible reasoning. You shouldn't dismiss out of hand the possibility that the ticket will win when you're considering whether to buy it; that undermines the whole point of deliberating about whether to buy the ticket. What makes the reasoning terrible is not simply that in this case it leads to a conclusion that does not maximize expected utility, but that it rules out a possibility that ought to be taken into consideration in this deliberation.⁴

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.
- So,
- (1) If I buy this lottery ticket it will lose.

⁴ Thanks to a referee for a clarification here.

So,

- (3) I should not buy this ticket (2004: 174 (my numbering)).

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 there is something wrong with these arguments, it is with the premises, (1) or (4) respectively. For Argument A, given that the ticket costs a penny and that it is worse to lose a penny than not to, the conclusion follows as inexorably as a practical conclusion can follow. For Argument B, given the auxiliary premise that a winning lottery ticket yields enough money to go on an African safari, the conclusion is similarly inexorable. In each case, we may assume that the auxiliary premises are beyond question.⁵ For the most part, I will consider only syllogistic arguments whose only possible flaw is in the acceptability of a certain premise, the better to focus on the epistemic properties of those premises.⁶ For convenience's sake, I will call such arguments *proper* arguments. If an argument is proper and the premise in question is acceptable, then the conclusion follows. (In Section 5, I will briefly consider probabilistic practical arguments and the idea that practical reasoning should not take the form of a practical syllogism at all.)

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.

So,

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

⁵ Thanks to Alan Millar for pointing out the necessity of these auxiliary premises.

⁶ It might be thought that Argument B has another flaw, in that it would be an odd way of arriving at the decision not to buy the lottery ticket. No one will actually start with the premise that she will not be able to afford the safari and conclude that her ticket won't win. But this does not amount to a flaw in the argument itself, any more than a theoretical deduction becomes invalid merely because it is strange to carry it out in particular circumstances.

Possibly one could give an account on which (4) is an acceptable premise and Argument B is proper, but on which nevertheless the action prescribed by the argument should not be carried out. On such an account one would not always be able to exploit the conclusions of proper practical reasoning with acceptable premises. This would be the practical analogue of denying deductive closure for knowledge, so that one cannot always gain knowledge by performing valid deductions from true premises. However, as an explanation of the problem with Argument B, this approach seems less promising than an account on which (4) is unacceptable.

So

- (6) I should buy the local destination guide (2004: 177).⁷

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. It is not merely that winning the lottery is improbable enough that the local guide has a higher expected utility than the guidebook to Africa; in this case, we might think, winning the lottery is a remote enough possibility that you need not consider it at all in your deliberations. (See Section 5 for a consideration of the probabilistic version of this deliberation.)

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 remains, however: What is it for a premise to be acceptable? When we look more closely at this question, we will see that there is no single 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 proper 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 proper practical syllogisms are all and only the true premises. This line of thinking shows that for practical reasoning it is important to have true beliefs.

There is an obvious way to avoid the uncomfortable result that Arguments A and B are acceptable. We can 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

⁷ Hawthorne may hold that 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 his 2004. In my text we may stipulate that you are not thinking about your lottery ticket.

reasoning as reasoning that is not vulnerable to criticism, that is not feckless or rash or overcautious.⁸

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 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. Failing to follow Argument C 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.⁹ The practical environment matters here: It determines how

⁸ Compare Hawthorne's contrast 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 (2004: 175 n. 33).

⁹ Exactly what conception of justification is at issue will depend on what notion of criticism of the practical reasoning is at issue. If criticism is warranted only when the believer has made some sort of culpable mistake in evaluating the evidence, then the relevant notion of justification will consist in not having made a culpable mistake in evaluation; if criticism of the reasoning is warranted whenever one's premise is not adequately supported by the evidence, then the relevant notion of justification will consist in the evidence's support for a proposition, without regard to the procedure through which one arrived at the proposition. For instance, if one blamelessly makes a mistake in calculation, the resulting practical reasoning can be criticized in the second sense but not in the first, since the reasoning is based on a premise that the evidence does not support, but the mistake made in evaluating the evidence was not culpable. And the premise one arrives at is justified in the second sense but not the first, since it is not supported by the evidence but there is

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.

3

The argument concerning the lottery case can be applied to any practical reasoning that calls for an instantaneous decision. If it is important whether 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 whether the subject's reasoning is 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 (2000: 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 standpoint from which the acceptable premises for practical reasoning are exactly the things we know. There is still a split

no culpable mistake in the procedure used to arrive at it. (Thanks to a referee for helping to clarify this point.)

¹⁰ Williamson himself immediately 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.' Later (Williamson 2000: 9), he argues that knowledge is more stable than true belief, with reference to the *Meno* question about the value of knowledge.

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.

So,

(9) I should burgle the house.

For Argument D to be proper, 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. Then, for (9) to follow, both (7) and (8) must 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. Let us focus on the acceptability of (7), assuming for the sake of argument that (8) is acceptable from whatever standpoint.

Consider two burglars, Moriarty 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 Moriarty happens to give up his belief in (7) once he has looked under the bed while Raffles will preserve his belief in (7) until he finds the diamond, then Raffles will find the diamond and Moriarty 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's informant told him that the diamond was under the bed, but Raffles persists 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 evidence for believing there is a diamond in the house, but has nevertheless got the idea into his head and will not give it up until he has thoroughly searched the house. It will not matter from the standpoint of actual success that Raffles does not know (7). So long as Raffles's belief is true and persistent, he will find the diamond.¹¹

¹¹ Compare Kvanvig's objection to Williamson in terms of beliefs that are fixed by nonevidential factors (Kvanvig 2003: 15).

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 Moriarty has been told that there is a diamond in the house and it is definitely under the bed.¹² 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. Moriarty, 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 comes in. This rules out some Gettier cases, as in the Moriarty case just described. 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 his reasoning will remain beyond criticism until he completes it. Then Raffles will be seen to have failed, through bad luck. But from the standpoint of criticism an acceptable premise is not necessarily one that, combined with a proper argument, will lead to a successful plan. This standpoint on practical reasoning shows stable justification, not knowledge, to be important in itself.

4

There is one standpoint from which the acceptable premises for practical reasoning always have both truth and a justification-like property. We can ask: Would this successful practical reasoning still have succeeded if circumstances had been different? From this standpoint, an acceptable premise is not only true but counterfactually true.¹³ Hence, when we are concerned with the counterfactual

¹² The idea is that once Moriarty finds that there is no diamond under the bed he should come to doubt his original information. Williamson's original description of the case does not seem to rule out that the burglar, finding no diamond under the bed, should conclude that it was moved and look elsewhere in the house. We can stipulate that in the current situation that is not the case.

¹³ On the most plausible reading of the questions that this standpoint is concerned with, what matters is not the counterfactual success of reasoning from the exact same premise but counterfactual success of reasoning from whatever premises the agent would have come to believe in different

success of practical reasoning, the important property of belief is safe truth, where a belief that p is safe iff in nearby possible worlds the agent believes that p only if p is true (Pritchard 2005: 71).¹⁴ Besides Pritchard, Williamson (2000: 123 ff.) and Sosa (2000) have 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, for the beliefs of BonJour's clairvoyant (BonJour 1985: ch. 3) are safe. In every nearby world in which the clairvoyant arrives at his belief by the same method, it is true.¹⁵ 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 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.

This may not be a definitive refutation of the sufficiency of safety; 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 consider these propositions

circumstances, by the same method that she could come to believe the premise in question in the actual circumstances. So if success depends on the agent's successfully naming 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 'counterfactual truth' should be taken to mean counterfactual truth of the premise arrived at by the same method.

¹⁴ Sosa (2000: 14) defines safety in terms of the counterfactual $B(p) \rightarrow p$, where $B(p)$ stands for belief that p . Since the antecedent of the counterfactual is true in the actual world, this is more clearly understood in terms of nearby possible worlds. 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): If p were false, then the agent would not have believed that p . 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.

¹⁵ 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.

and believe them falsely. Yet in other cases, it does not destroy knowledge if in nearby possible worlds we falsely believe closely related propositions. If I always identify geese as ducks, I can still know that there is a swan in a lake, even though in a nearby world in which the swan is replaced by a goose I would falsely believe the related proposition 'There is a duck in the lake.' To analyze knowledge as safe true belief, we would need a principled reason that one of these alternative propositions should count against safety and the other not. The prospects for such a reason seem poor, as do the prospects of analyzing knowledge as safe true belief, and thus of establishing the importance of knowledge from the standpoint from which we care about the safe success of practical reasoning.

Against the necessity of safety for knowledge, see Comesaña (2005) and Neta and Rohrbaugh (2004). These authors present interestingly different diagnoses of why knowledge can fail to be safe. Comesaña 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 whatever value mere reliability has, reliable reliability has more of it. Neta and Rohrbaugh (2004: 404) 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.

This leads to the more important problem with the attempt to establish the practical importance of knowledge qua safe true belief. 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 whether our reasoning is 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. Hence even if there is a standpoint from which we value safe success in our practical reasoning, the practical premises that are acceptable from this standpoint will not be just the known ones, unless we do know that our lottery tickets will lose. We might escape this consequence by defining remoteness and safety in terms of the possibilities that you ought to take into account, given the practical decision you are making. But this definition builds in the importance of your reasoning's being immune from criticism, which is motivated from a different standpoint from the one that might motivate safe truth.

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 (2000: ch. 9, revising Williamson 1997). If this were true, then reasoning from premises that were not known would always be irrational. Williamson bolsters this position by arguing that we are not necessarily in a position to know what it is rational for us to believe. If 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. This argument can easily be extended from theoretical to practical rationality.

Even granting Williamson's analysis of the hallucination and similar cases, it will not follow that rational reasoning always proceeds from known or even true premises if we are not antecedently committed to the importance of knowledge. Consider Argument C, the bookstore 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.¹⁶ 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 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

¹⁶ 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.org/2003/08/26/evidence-and-knowledge/> [accessed Sept. 9, 2006].)

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.

5

So far the model of practical reasoning discussed has been the practical syllogism. The premises of these syllogisms are categorical statements, and the conclusion is an action that is called for given those premises. But this is not the only way to approach practical reasoning. Perhaps correct practical reasoning proceeds from premises that assign probabilities to certain outcomes to conclusions that maximize expected utility given those probabilities (or respond to those probabilities in some other way). Hence when presented with the lottery ticket one could reason thus:

Argument E

- (10) The chance that this ticket will win is 1 in 10,000.
 - (11) The ticket costs a cent and is worth \$5,000 if it wins. The expected value of the ticket is 50 cents (\$5,000 divided by 10,000).
- So,
- (12) if I buy the ticket I'll have an expected gain of 49 cents (the ticket's expected value minus its cost).
 - (13) So I should buy the ticket.

And in the bookstore one should reason as follows:

Argument F

- (14) I'll only be able to afford an African safari if this ticket wins the lottery.
 - (10) The chance that this ticket will win is 1 in 10,000.
- So,
- (15) there's only a 1 in 10,000 chance that the African destination guide will be any use, and its expected value is quite small.
 - (16) There's a 9,999 in 10,000 chance that the local destination guide will be of use, and its expected value is reasonably large.
- So
- (17) I should buy the local destination guide.

The question now is, what property must the critical premise (10) have in order to be an acceptable premise for this practical reasoning?

From the standpoint of actual success, (10) is *never* the best premise. Your practical reasoning will actually succeed if you assign probability 1 to whatever turns out to be true. Assigning any lower probability will sometimes lead you to forgo some course of action that would have turned out to be successful. We've already seen this; in the case in which your ticket loses, Argument E leads you to lose a penny, whereas if you assign probability 1 to 'This ticket won't win' you won't buy the ticket and you'll save a penny.

What about the standpoint that is concerned with whether your reasoning can be criticized? Here it matters how we interpret probability. If (10) means that there is an objective chance of 1 in 10,000, then it will not be the right kind of premise to guarantee immunity from criticism. You may be immune from criticism for employing Argument E or F even if the lottery is rigged so that there is no objective chance that your ticket will win (or lose, respectively), so long as your evidence indicates that the lottery is fair. And you may be criticized for using these arguments even if the objective chance that your ticket will win is 1 in 10,000, if your evidence indicates that the odds are different (for instance, if an ordinarily reliable source tells you that there are more or fewer tickets in the lottery).

If probabilities are interpreted more subjectively, however, then it may be that reasoning from true probabilistic premises will be immune from criticism. I do not mean the view of subjective probabilities on which any assignment of credences to propositions is permissible so long as it is consistent with the axioms of probability. (On such a view, the question of which premises about probabilities were acceptable would not arise.) Rather, I mean a conception on which the available evidence determines some probability for a proposition, which will be the credence that a believer ought to give that proposition based on that evidence.¹⁷ Suppose we interpret (10) to mean that on the available evidence the appropriate credence in 'This ticket will win' is 1 in 10,000. If you follow Argument E or F because the evidence indicates that there is a 1 in 10,000 chance that your ticket will win, you will be beyond criticism, for then you are conforming your actions to the evidence. The 'because' is necessary here; if, without evaluating the evidence, you guess that your ticket has a 1 in 10,000 chance of winning, you can be criticized even if the evidence does support that. Believing (10) because the evidence supports it means being justified in believing it, so it appears that (10) is an acceptable practical premise whenever it is true and believed with justification.

¹⁷ Other subjective interpretations of probability are possible. However, if on these interpretations (10) does not amount to 'On the available evidence, it is appropriate to give a credence of 1 in 10,000 to this ticket's winning,' then reasoning from (10) will not be immune from criticism. The interpretation discussed in the text is the friendliest to the idea that knowledge of probabilities is valuable for practical reasoning.

This may even be a standpoint from which the acceptable premises for practical reasoning are all and only the (probabilistic) propositions that are known. On the subjective interpretation, if (10) is true and believed because it is true, then we cannot criticize any proper practical reasoning that uses it as a premise. Arguably in all of these cases (10) will be known. If the conception of justification at issue is that a belief is justified whenever the evidence supports it, and the conception of probability at issue is that a proposition's probability depends on the degree to which the evidence supports the proposition, then whenever the odds of the ticket's winning are 1 in 10,000 the evidence will support belief in those odds, and thus belief in (10) will be justified.¹⁸ Since this justification for believing (10) entails the truth of (10), any belief in (10) based on this justification will amount to knowledge. Thus from this point of view the acceptable practical premises might be only those propositions about probability that are known.

However, this would only be a small victory for proponents of the intrinsic value of knowledge. To begin with, from this standpoint knowledge is not uniquely valuable, since justification, knowledge, and truth are so intimately connected. Not only does having a justified belief about a probability entail knowing what the probability is, but whenever a probability statement is true a justification is available for believing it. You need only properly evaluate the evidence, and you will with justification believe that the probability is what it actually is.

More important, this argument restricts the domain of knowledge. If knowledge is important because known premises can be used in probabilistic practical reasoning, then only probabilities can be known. Very rarely does the evidence make it appropriate to reason as though a proposition had probability 1. If our concept of knowledge is the one that is important for this kind of practical reasoning, we will almost never empirically know categorical propositions. Certainly we will not know anything about the not quite certain future or present, as in the lottery case or Vogel's car theft case. In addition, this knowledge is essentially knowledge of justification, since it has to do with the credence that the evidence supports. Probabilistic practical reasoning will not establish the importance of knowledge of ordinary facts. The critical concept here is much closer to the concept of justification that we have already seen to be important in itself.

¹⁸ On other conceptions of justification, justified belief in (10) can come apart from its truth. For instance, if justification means arriving at belief responsibly and blamelessly, it may be possible to responsibly and blamelessly arrive at an incorrect belief concerning the extent to which the evidence supports a proposition, for instance by an error in calculation. (See n. 9 above; thanks to a referee for raising this point.) In this case there will be no single standpoint from which acceptable practical premises need be both justified and true, and thus the acceptable premises will not be those propositions that are known. As with different interpretations of probability (see n. 17 above), the interpretation of justification discussed in the text is the friendliest to the idea that knowledge of probabilities is in itself valuable for practical reasoning.

We have seen that there is no single standpoint on practical reasoning from which knowledge of categorical propositions is important. Instead, from various standpoints, it is important that one's premises be true, justified, persistent, stably justified, and safe. 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 important things? Yet in fact people talk about what people know much more often than about, for instance, what they are justified in believing.¹⁹

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 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 convenient to have a quick way of expressing all these different concepts that may be important from 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 that any proper argument S made with *p* as a premise succeeded, 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, a justified true belief 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 for your plan to 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

¹⁹ 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).

based on it.²⁰ When we call a belief knowledge, we guarantee it to be 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 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?²¹

My answer 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 say 'S knows that *p*,' knowledge 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.

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 reasoning only

²⁰ 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 (accessed Sept. 11, 2006).

²¹ 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.

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.²²

REFERENCES

- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Comesaña, J. (2005). 'Unsafe Knowledge,' *Synthese*, 146: 395–404.
- DeRose, K. (1995). 'Solving the Skeptical Problem,' *Philosophical Review*, 104: 1–52.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge,' *Journal of Philosophy*, 73: 771–91.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kaplan, M. (1985). 'It's Not What You Know that Counts,' *Journal of Philosophy*, 82: 350–63.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, MA: Cambridge University Press.
- Lewis, D. (1996). 'Elusive Knowledge,' *Australasian Journal of Philosophy*, 74: 549–67.
- Neta, R., and Rohrbaugh, G. (2004). 'Luminosity and the Safety of Knowledge,' *Pacific Philosophical Quarterly*, 85: 396–406.
- Pollock, J., and Cruz, J. (1999). *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman and Littlefield.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Sosa, E. (2000). 'Skepticism and Contextualism,' in J. Tomberlin (ed.), *Philosophical Issues*, 10: 1–18.
- Vogel, J. (1990). 'Are There Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.) *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
- Weinberg, J., Nichols, S., and Stich, S. (2001). 'Normativity and Epistemic Intuitions,' *Philosophical Topics*, 29: 429–60.
- Williamson, T. (1997). 'Knowledge as Evidence,' *Mind*, 106: 717–41.
- (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

²² Earlier versions were presented at the *Central Division Meeting of the American Philosophical Association* in April 2005 in Chicago, with commentary by Mylan Engel, and at the *Epistemic Value Conference* at the University of Stirling in August 2006, with commentary by Igor Douven. Thanks to the commentators and to members of the audience on those occasions, particularly to Berit Brogaard, Christian Piller, Dennis Whitcomb, Stephen Green, and Duncan Pritchard. Thanks also to Ram Neta, Elijah Milgram, Matt McGrath, Keith DeRose, A. P. Martinich, Alan Millar, and two anonymous referees for helpful comments on this and related matters; apologies to anyone I have neglected to include. This research was supported in part by a Texas Tech New Humanities Faculty Grant.

8

Pragmatic Encroachment and Epistemic Value

Pascal Engel

1. PRAGMATIC ENCROACHMENT AND THE VALUE OF KNOWLEDGE

When in the *Meno* (97a–c) Socrates asks whether knowledge is more valuable than true belief, the notion of value which he has in mind seems clearly to be that of *practical value*. In 97a he asks whether 'good men' or 'men of value' (*agathoi androi*) will be 'useful' (*ophelimoí*) and when in 97c he raises his famous question about the difference between having a true belief about the road to Larissa and knowing the road to Larissa, this question is directly couched in terms of what is correct with respect to practice (*pros orthothèta praxeós*) and the example is clearly meant to ask something about the respective roles of belief and knowledge in guiding our actions. Contemporary approaches to the problem of the value of knowledge, however, have investigated other senses of the notion of value or worth of knowledge, in terms of various notions of epistemic virtue (Sosa 2007; Greco 2002; Zagzebski 1996), or in terms of moral appraisal (Brady 2006). Although many of these accounts involve the idea that the worth in which knowledge consists has to do with some kind of practical achievement or success, most of them accept the traditional view that knowledge is an epistemic good, and that its value is mostly of a theoretical, not of a practical nature.

Let us, following Duncan Pritchard (2007), distinguish the *primary value* problem for knowledge—the *Meno* problem of whether and why knowledge is more valuable than true belief—from the *secondary value* problem—the problem of whether knowledge is more valuable than any proper subset of its parts, on the assumption that the components of knowledge are true belief and justification, that is, whether knowledge is more valuable than justification, or true belief, for instance. The claim that some subparts of knowledge can be valuable because of their role in action surfaces in the writings of some virtue epistemologists. Thus Jonathan Kvanvig writes: 'Belief is valuable because it is action guiding' and he tells us that true beliefs are valuable because they lead to actions that 'are successful in satisfying desires and in achieving purposes' (Kvanvig 2003: 30). All