

A DEFENSE OF STABLE INVARIANTISM

Baron Reed
b-reed@northwestern.edu

A true belief and a belief that counts as knowledge can both be practically useful. But the important difference between them, as Plato tells us in the *Meno*, is that only knowledge has a “tether” so that it cannot run away. The general sense that knowledge is fairly stable has been an enduring part of the epistemological tradition—until recently. Some currently popular views—contextualism and interest-relative invariantism—have the consequence that whether someone can correctly be said to have knowledge can change abruptly, even when there is no change in her evidence or reliability. These views are embraced in large part because they are thought to provide solutions to old epistemological problems. But, as I shall argue below, the cure is worse than the disease. The innovative aspects of these views either do not work or undermine our interest in knowledge. A more satisfying view, which I shall call *stable invariantism*, offers solutions to these problems derived simply from the basic commitment to *fallibilism* that virtually every epistemologist shares.¹

I. CONTEXTUALISM

The traditional argument for skepticism is intuitively compelling, especially when we vividly imagine ourselves in a skeptical scenario. But the force of the argument is impossible to see in some of the best known anti-skeptical arguments.² We are left to wonder why, if it could be so easily dismissed, skepticism has been a perennial worry for so many philosophers. Other anti-skeptical strategies concede the apparent forcefulness of the skeptical argument, but argue that skeptical possibilities are—all appearances to the contrary—irrelevant to ordinary knowledge claims. Such a response depends on the rejection of the *closure* principle, which holds that a proposition competently inferred from known premises is itself known.³ For many philosophers, abandoning the closure principle is hardly better than giving way to skepticism. So, where some responses to the skeptic are too

dismissive, others grant too much. The challenge is to find a proper balance between these two extremes.

According to contextualism, *semantic ascent*—focusing, not on knowledge, but on the use of the term ‘knowledge’—allows us to find such a balance. The fundamental thesis of contextualism holds that ‘knowledge’, ‘know’, etc., as well as their cognates, are context-sensitive in an interesting way: the truth-conditions for assertions involving these terms are determined in part by the conversational context in which the assertions are made, where the conversational context fixes the epistemic standards relevant to the knowledge claim.⁴

This thesis makes possible the following response to skepticism. In ordinary conversational contexts, the relevant epistemic standards are relatively lax. So, e.g., when Maria has read the syllabus for my class and I say, “Maria knows when the final exam will be,” I have said something true.⁵ But, when we begin to talk about Descartes’s skeptical arguments, the conversational context shifts. The epistemic standards now relevant are much more stringent. It would now be correct for me to say, “Maria does not know when the final exam will be.”⁶ Because the epistemic standards have shifted, though, this does not contradict my earlier attribution of knowledge to Maria—just as my assertion on Monday that Jim is tall (when I see him standing next to his younger brother) is not contradicted by my assertion on Tuesday that Jim is short (when I see him next to his mother). The net result is that the skeptic is right when she denies that we have knowledge—and, in granting that, we have acknowledged the forcefulness of her argument—but also we are right when we say, in ordinary contexts, that we have knowledge. So, contextualism gives a little ground to the skeptic but ultimately allows us to feel secure in our ordinary knowledge-attributions.

The success of the contextualist response depends on its ability to restrict the skeptic’s challenge to contexts where high epistemic standards are operative.⁷ This is what is supposed to render skepticism irrelevant to our ordinary knowledge-attributions. But, as I shall argue, the

contextualist cannot make this restriction stick. The skeptic can simply construct an argument that parallels the traditional one but is unaffected by the contextualist thesis.

Let us say that the property to which variable epistemic standards apply is *strength of epistemic position* (SEP).⁸ Note that, for the contextualist, SEP does *not* vary from one context to another.

Rather, it is what remains constant, though one's standards governing the level of SEP sufficient for a correct attribution of knowledge may change. Let \sim BIV be the proposition that I am not a brain in a vat, and let H be the proposition that I have hands. The argument, then, proceeds as follows.

- 1) My SEP with respect to the proposition that H entails \sim BIV is excellent.
- 2) If my SEP with respect to the proposition that H entails \sim BIV is excellent, then my SEP for H cannot be much higher than my SEP for \sim BIV.
- 3) My SEP with respect to \sim BIV is poor.

Therefore,

- 4) My SEP with respect to H is poor.

The first premise says that I am in an excellent epistemic position for the proposition that my having hands entails that I am not a handless brain in a vat. Given that the entailment is obvious, the first premise is quite plausible. The third premise says that I am in a poor epistemic position with respect to the proposition that I am not a brain in a vat. Although some philosophers would deny this, the contextualist should not: if I were to already have the ability to reject the possibility that I am a brain in a vat, I would not need the contextualist response to skepticism. The second premise, which connects them, is an application of the following closure principle:

(CP-SEP) If my SEP with respect to the proposition that P entails Q is excellent, then my SEP for P cannot be much higher than my SEP for Q.⁹

This is similar to the closure principle for knowledge over known entailment:

(CP-K) If I know P and I know that P entails Q, then I know Q.

Assuming I know that P entails Q, (CP-K) says that if I know P then I know Q and, equivalently, if I do not know Q then I do not know P. Similarly, assuming my SEP for the proposition that P entails Q is excellent, (CP-SEP) says that if my SEP for P is excellent then my SEP for Q has to be in the neighborhood of excellent, and if my SEP for Q is poor then my SEP for P has to be in the neighborhood of poor. It is necessary to talk about the *neighborhood* of one's SEP because one might have a little room for doubt about the entailment running from P to Q, given that one's SEP for that proposition is excellent rather than perfect. Intermediate values for SEP—e.g., good or fair—could figure in similar closure principles.¹⁰ I leave them aside merely because they play no role in the objection to contextualism. I should note that the contextualist cannot deny the second premise without also calling into question the closure principle for knowledge. (CP-K) holds, essentially, that different items of knowledge in a single rational mind are epistemically significant to one another. Surely these different items of knowledge are connected in virtue of their epistemic grounding—which is what (CP-SEP) holds. Moreover, the closure principle for knowledge cannot be rejected without obviating the need for contextualism—once we reject closure, the way is open to accepting the anti-skeptical response given by Dretske and Nozick.

The conclusion says that my SEP for the proposition that I have hands is poor. The fundamental thesis of contextualism says only that the truth conditions of 'know' assertions will be

determined by the epistemic standards fixed by the conversational context. Though it is assumed that most of our beliefs meet the relatively lax standards operative in mundane conversations, it is compatible with contextualism that they do not. The skeptic has just produced an argument that shows they do not. What is worse, the contextualist is badly positioned to respond to the argument. Because SEP does not vary from one context to another, there is no chance that the contextualist response to skepticism will work here.¹¹ And, because the contextualist has already conceded that the original skeptical argument (framed in terms of knowledge) cannot be directly refuted—hence the need to limit its forcefulness to certain conversational contexts—it doesn't look like the contextualist has any hope of there being a refutation of *any* kind for the new skeptical argument that parallels the old one.¹² In the end, then, contextualism leaves us more, not less, vulnerable to skepticism.

II. INTEREST-RELATIVE INVARIANTISM

Interest-relative invariantism has recently emerged as one of the major rivals to contextualism. The view is characterized by its rejection of the contextualist claim that the semantic value of 'knowledge' varies from one context to another and by its rejection of what is called *purism* or *intellectualism*: the claim that whether a subject has knowledge is determined by purely epistemic considerations.¹³ In rejecting the latter claim, the interest-relative invariantist holds that features of the subject's practical situation may play a role in determining whether she has knowledge.

The interest-relative invariantist's account rests on examples like the following:¹⁴

Train Case 1. You are in Boston and are about to board one of the trains that ultimately go to Providence, which is your destination. You are in no hurry but ask the passenger next to you out of curiosity whether this train is the express (and goes

straight to Providence) or the regular service that first stops in Foxboro. He says that the ticket agent told him that it is the regular service. You have no reason to doubt him.

Train Case 2. You are in Boston and are about to board one of the trains that ultimately go to Providence. Your destination, however, is Foxboro, and it is vitally important that you get there as soon as possible. You ask the passenger next to you whether this train is the express (which goes directly to Providence and does not stop in Foxboro) or the regular service (which does make that stop). He says that the ticket agent told him that it is the regular service. You have no reason to doubt him.

Most would agree that, in the first case, you know the train will stop in Foxboro. In the second case, however, the interest-relative invariantist argues that you ought to make sure the train will stop in Foxboro, given that it is so important for you to get there. If this is true, he concludes, you don't already know that the train will stop in Foxboro in the second case. But you have exactly the same evidence in the second train case as you do in the first. The only difference between them has to do with the practical consequences of acting on the basis of that evidence. Therefore, the practical consequences determine (in part) whether or not you know that the train will stop in Foxboro.

The general principle underlying this argument is that if you know that p , then you should be able to act as if p .¹⁵ In other words, if you shouldn't act as if p , you don't know that p . Along these lines, Fantl and McGrath offer the following general argument:¹⁶

(1) S knows that p .

(2) S knows that if p , then A is the best thing she can do.

Therefore,

(3) S knows that A is the best thing she can do.

Therefore,

(4) S is rational to do A .

From this argument, Fantl and McGrath derive the principle (PCK) that S knows that p only if S is rational to act as if p .¹⁷

But this principle cannot be correct. Consider the following counterexample:

Punishment/reward case 1. You are participating in a psychological study intended to measure the effect of stress on memory. The researcher asks you questions about Roman history—a subject with which you are well acquainted. For every correct answer you give, the researcher will reward you with a jelly bean; for every incorrect answer, you are punished by an extremely painful electric shock. There is neither reward nor punishment for failing to give an answer. The first question is: when was Julius Caesar born? You are confident, though not absolutely certain, that the answer is 100 BC. You also know that, given that Caesar was born in 100 BC, the best thing to do is to provide this answer (i.e., this course of action will have the best consequences—you will be one jelly bean richer!).

In this case it seems clear that, when you weigh the meager reward against the prospect (however unlikely) of excruciating pain, it would not be rational to attempt an answer to the question. Does this show that you don't *know* when Caesar was born? Not at all. Your shortcoming is merely a lack

of certainty when only something in that neighborhood would do. There is nothing odd about continuing to think of yourself as possessing knowledge in this situation, even though it would not be rational for you to do what you know would have the best consequences. Bear in mind that, when the action in question is the one that is the best thing to do, this means that the action will have the best consequences *given* that p . But, if fallibilism is true and our knowledge is grounded in something less than certainty, it is never a given for us that p , even when we know it is true.¹⁸ So, we can take ourselves to know that p while still recognizing that there is a chance that it is false that p . When that chance carries with it very bad consequences if it were to become actual, as in the above case, it may be most rational *not* to act as if p .

If this is correct, it means that (3) does not entail (4): one can know that an action is the best thing to do even though it is not rational to do it. Fantl and McGrath argue that the entailment does hold because (a) we tend to cite knowledge in the judgments we make about whether actions are reasonable and (b) there is no explanation of this fact that allows for a link weaker than entailment between knowledge and practical rationality. In particular, they argue that the link between knowledge and practical rationality cannot be cancelled and thus is not merely a conversational implicature.¹⁹ But this rationale for the argument clearly does not hold with regard to the punishment/reward case above. It would be perfectly natural for you to say to the researcher, “I *do* know this one, though it’s certainly not worth risking a shock.” In this sort of case, our judgment about the reasonableness of the action is guided by a very high grade of knowledge, approaching certainty; the fact that you have surpassed the minimal threshold for a lower grade of knowledge is recognized as irrelevant in this situation. Note that this is not merely a *theoretical* point: it is not just that epistemologists would describe the situation in this way. Rather, the fact that it would be natural for a subject to say, “I know this one but I’m not going to risk an answer,” shows that our ordinary language allows us to disconnect knowledge and practical rationality.

There are two further problems with Fantl's and McGrath's argument that there is a relation of entailment between knowledge and practical rationality. First, premise (3) says that *S* knows that *A* is *the best thing she can do*. By this they mean that *S* knows that *A* will have the best consequences, given that *p*. But it is not clear that this is what ordinary speakers mean when they say that an action is the best thing that can be done. As I argued above, if fallibilism is correct, we are never really *given* that *p*—that is to say, our knowing that *p* will mean that the probability that *p* is high but less than 1. It seems to me that when ordinary speakers say that an action is best, they very often mean that the action has the highest expected utility of all the alternatives under consideration. In the reward/punishment case above, the expected utility of answering the question may very well be lower than the expected utility of not answering, given the horrible consequences of answering incorrectly. In *this* sense of 'best thing to do', answering will not be the best thing to do; moreover, in this sense of 'best thing to do', it *does* seem to be the case that a subject's knowing *A* is the best thing to do entails that it is rational for the subject to do *A*. Given this, it is clear why it sounds absurd to say things like, "I know the best thing to do is to answer, but it is not rational to do it." But this is *not* the sense of 'best thing to do' that Fantl and McGrath are using. So, the data of ordinary language use cannot be taken as supporting their argument.

Second, Fantl and McGrath argue that the link between (3) and (4) is one of entailment because the implicature from knowledge to rationality cannot be cancelled. Thus, e.g., if I were to say, "I know that the best thing to do is to answer, but I am not rational to answer because I don't know for sure that the answer is right," my attempted cancellation would result in absurdity. The cancellation would entail, for example, the absurd conclusion that I am rational to do something that I know is worse than an alternative action.²⁰ But whether one regards the entailed conclusion as absurd depends, of course, on whether one thinks that knowing an action is the best thing to do (in their sense) really does entail that it is rational to perform that action. And this is the very issue at

stake. It is obviously question-begging to argue that the link in question is one of entailment rather than conversational implicature on the grounds that a denial of it leads to absurdity.²¹

Setting aside these objections to the argument, there is a more significant reason to reject interest-relative invariantism: if it were correct, knowledge would be trivialized. Consider the following:

Punishment/reward case 2. As before, a researcher is asking you questions about Roman history. But now there are two punishment/reward scenarios. The first is as before: for a correct answer, you receive a jelly bean, and for an incorrect response, you get a very painful electric shock. In the second, a correct answer results in a \$1000 reward, while an incorrect one leads only to a very mild slap on the wrist. In both scenarios, there is neither reward nor punishment for failing to answer.

Although you must consider both scenarios simultaneously, you are not bound to give the same answer in each.

Most people, I suspect, would refuse to answer in the first scenario but not in the second. A theory that incorporates a practical condition on knowledge would presumably have to say that such a person would both know and not know when Caesar was born.²² To avoid contradiction, these knowledge claims would have to be relativized to practical context.²³ Hawthorne admits that “Allowing such a mechanism [i.e., one driven by practical factors] will make knowledge come and go with ease.”²⁴ But surely—when one can be in two different practical contexts at the same time, and thus simultaneously both know and not know the same proposition—this would make it *too* easy.

There is a further problem. Even if we relativize knowledge to practical context in this way, there will still be something common to both contexts that allows for the different evaluations in

each (i.e., knowing in one but not in the other). Again, we can call this *strength of epistemic position*. Notice that SEP does *not* vary from one practical context to another. Interest-relative invariantists are not arguing that the subject's practical concerns affect the quality of her evidence or of her reliable faculties, or of anything else that contributes to SEP.²⁵ What we find, then, is a view quite similar to contextualism. SEP is a context-invariant feature of the subject. The threshold for SEP sufficient for correctly attributing knowledge to the subject will go up or down depending on non-epistemic factors—in the case of contextualism this will be determined by the speaker's conversational context, while for interest-relative invariantism it will be the subject's practical context. We have seen that, with respect to contextualism, this is an unsuccessful view. The problems it is meant to avoid can simply be reframed in terms of SEP. As we shall see, the same can be said of interest-relative invariantism; the philosophical work knowledge is supposed to do, on this account, ends up being done by SEP instead.

To see this, let us return to the second train case above: you are trying to decide whether to check if the train stops in Foxboro because it is extremely important that you get there as quickly as possible. You have not yet decided whether it is rational to check if the train makes that stop. One of the relevant factors in your decision would presumably be an answer to the question, do you know the train will stop there? But the interest-relative invariantist has expressed the purported link between ignorance and irrationality as a necessary condition on knowledge: S knows that *p* only if it is rational to act as if *p*.²⁶ Consequently, there is no fact of the matter as to whether you know that *p* until it has first been determined whether it is rational for you to act as if *p*.²⁷

When we pursue the matter at a deeper level, and inquire what makes the action reasonable, there is only plausible candidate remaining: SEP. But once you have determined whether your SEP is sufficiently good for it to be rational to check if the train stops, the question as to whether you

know the train stops has no explanatory power. Knowledge thus becomes, for interest-relative invariantism, a mere formality—and, what is worse, one that will “come and go with ease.”

I shall offer one final objection to interest-relative invariantism. If the view were correct, it would lead agents to behave in irrational ways. Consider the following:

Death penalty case. The jury for a capital murder trial knows, without being certain, that the defendant is guilty. The jurors have no ethical reservations regarding the death penalty in general, and, in murder cases as gruesome as the one they are deciding, they feel it would be the most appropriate punishment. So, they also know that, given the defendant’s guilt, the best thing to do would be to impose the death penalty. Therefore, they know that the best thing to do would be to impose the death penalty, in the sense that this would have the best consequences.

Nevertheless, because they are not completely certain that the defendant is guilty, it would not be rational for them to impose the death penalty and they decline to do so.

Notice that, if interest-relative invariantism were true, we would have to conclude, from the fact that imposing the death penalty is not rational, that the jurors do not know that it is the best thing to do. That is, not only would the jurors lack complete certainty that this is the best thing to do, they would not *in any sense* know that it is the best thing to do. But they do know that the following conditional is true: if the defendant is guilty, the best thing to do is to impose the death penalty. From this (via the closure principle for knowledge), the interest-relative invariantist would infer that the jurors also do not in any sense know that the defendant is guilty. From the fact that it is not rational for them to impose the death penalty, it would then follow that they should acquit the defendant. It seems

that there is no room for anything less than the maximum punishment. Surely this is the wrong result.

Now, it might be objected that the jury is not interested in determining whether they *know* the defendant is guilty. Their concern, rather, is whether it has been *proven beyond reasonable doubt* that the defendant committed the murder in question. So, reasoning of the sort just sketched—which traces the entailment linking the irrationality of imposing the death penalty to the jurors’ lack of knowledge regarding the defendant’s guilt—would be beside the point.

This would be a dangerous reply for the interest-relative invariantist to make, for it is hard to deny that the jurors are very much concerned with their cognitive situation regarding the defendant’s guilt. If the jurors do not attempt to describe their cognitive situation using ‘knowledge’-discourse, it shows only that ‘knowledge’-discourse is inadequate in expressing all of what we want to say about knowledge.²⁸ So, when interest-relative invariantists attempt to defend their view by appealing to purported linguistic evidence regarding our use of ‘knowledge’, their focus is far too narrow. Even if there were a relation of entailment holding between sentences like, “I know the train is on time,” and sentences of the form, “I am rational to act as if the train is on time,” the relevance to epistemology of this kind of entailment would be suspect.²⁹ As the death penalty case makes clear, there are other ways of describing a subject’s cognitive situation where the subject does have the knowledge in question but is not warranted in acting on it. Insofar as epistemology is properly concerned with knowledge, and not merely ‘knowledge’-discourse, we are left with little reason to accept interest-relative invariantism.

III. STABLE INVARIANTISM

I turn, finally, to an alternative account of knowledge—one that does not make knowledge “come and go with ease.” What I will call *stable invariantism* is the view that the basic resources of fallibilism

can be used to answer many of the epistemological problems that originally motivated contextualism and interest-relative invariantism. In what follows I shall briefly sketch an answer to the lottery paradox and provide an explanation of the intuitive appeal of skepticism.

1. *The Lottery Paradox*

Fallibilism is essentially the view that a subject's SEP for a given belief makes that belief probably true.³⁰ Consequently, a fallibilist must recognize that knowledge comes in degrees ranging downward from absolute certainty. In this sense, fallibilistic knowledge is *probabilistic*.³¹ As we shall see, it is in virtue of this aspect of fallibilism that the lottery paradox arises.

Suppose that Millie knows a 1000 ticket lottery will have only one winning ticket chosen, and suppose that the winner is ticket 1000. For each losing ticket, it seems that Millie can know it will lose. After all, the odds of losing, for each ticket, are overwhelming.³² Using basic inference rules, Millie can then begin to combine her individual instances of knowledge (e.g., knowing that ticket 1 will lose, knowing that ticket 2 will lose, etc.) into knowledge of conjunctions (knowing that tickets 1 and 2 will lose, knowing that tickets 1, 2, and 3 will lose, etc.). As she adds the last conjunct to a conjunction with 999 components, Millie realizes that, since she knows the first 999 tickets will lose, the winner must be ticket 1000. In fact, it will be the winner—but surely she cannot know that it will win on this basis.³³

The easiest way to respond to the lottery paradox is to deny that we can know, for each losing ticket, that it will lose. Although this is an easy solution, it is not terribly satisfying, given that so many of our ordinary beliefs can be made to fit the lottery pattern.³⁴ A fairly widespread skepticism lies at the end of that path.³⁵ If we wish to avoid it, we will have to find a solution that retains our knowledge, for each losing ticket, that it will lose, while preventing these different pieces of knowledge from being aggregated into the problematic conjunction.

Recall that a closure principle, such as (CP-K) in section II above, is plausible to the extent that we take bits of knowledge in a rational mind to be epistemically relevant to one another. If fallibilism is true, there will be limits to how far this relevance extends—the inferential use of enough merely probable premises (even where the probability is high) will ultimately lead to a conclusion that is improbable. Closure principles that cover only two highly probable premises, are acceptable; the same cannot be said for principles that bring multiple premises into inferential contact, where these premises are less than certain.³⁶

This restriction on closure provides a solution to the lottery paradox. Millie can know, for each individual losing ticket, that it will lose. Perhaps she can also know, for any collection of, say, five of these tickets, that all of them will lose. She cannot, however, know that the first 500 tickets will lose. Where to draw the line, between inferences that produce knowledge and those that do not, may be difficult. But it should nonetheless be clear that the inferences that are most obviously problematic will be ruled out by plausible restrictions on closure.³⁷

One might object to this restriction on closure because it is at odds with an intuitively plausible principle linking epistemic possibility with knowledge:³⁸

(EP-K) If S knows that p , it is not possible for S that $\sim p$.

If (EP-K) is true, there is not the slightest danger in combining separate bits of knowledge into a single conjunction and thus no reason to think that closure needs to be restricted. Hawthorne connects this principle with the Humean view that knowledge and probability are of “contrary and disagreeing natures.”³⁹ The conflict between fallibilism, which takes knowledge to be probabilistic, and this Humean view should be obvious. Given that fallibilism is now more or less the only option

in epistemology, apart from skepticism, we should accordingly be reluctant about accepting (EP-K).⁴⁰

Nevertheless, the principle may have support that is not ultimately grounded in the Humean view. Keith DeRose argues that (EP-K) is true because it provides the only plausible explanation for why we seem to hear a “clash” in sentences of the following types (where the possibility in question is epistemic):⁴¹

- (1) It is possible that not- p , but p .
- (2) It is possible that not- p , but I know that p .

The clash we hear in (1)-type sentences does not derive from logical inconsistency, since neither half entails the negation of the other. He argues, then, that the best explanation for why the two halves nevertheless seem to cancel each other out depends on two claims:

- (a) In flat out asserting that p , one represents it as being the case that one knows that p .
- (b) Sentences of type (2) clash because they are logically inconsistent.

And if sentences of type (2) are logically inconsistent, then (EP-K) must be true.

It is not clear, though, that there really is a clash in sentences of type (1). Suppose someone asks me what the old name for New York was. I might say, “It used to be called ‘New Amsterdam’, though I could be wrong about this—my memory isn’t what it used to be.” My acknowledgment of the possibility of error does not *cancel* my assertion that New York used to have the name ‘New Amsterdam’ but merely *qualifies* it. I may still be representing myself as knowing this proposition, but I am also providing additional information about the limitations of my SEP with respect to it.

Although good enough to clear the threshold for knowledge, my SEP still leaves room for error. A sentence of type (1) clashes, then, only on the assumption that it contains a flat out assertion that p . But this assumption is incorrect: in fact a type-(1) sentence is, as a whole, a way of asserting that p while making it clear that one is *not* making a flat out assertion. So, we do not need (b) to explain why a type-(1) sentence clashes for the simple reason that it doesn't.⁴²

Similarly, we do not need to explain the apparent clash in type-(2) sentences as resulting from logical inconsistency. Granted, there is something infelicitous about sentences of this form. But we can see this as arising from the fact that 'know' is mostly used as a threshold or non-gradable predicate. As we use the term, you either know something or you don't; we don't say that you know something better or worse than anyone else.⁴³ But this fact about the way we typically make knowledge attributions is obviously compatible with the fact that knowledge *can* come in degrees. If two beliefs both surpass the threshold distinguishing knowledge from ignorance, it is clearly possible that one might do so by a greater margin than the other.⁴⁴ When we wish to talk about a higher or lower grade of knowledge, we tend to switch from sentences of the form, "S knows that p ," to sentences of the form, "S is more (less) certain that p ."⁴⁵ There is *no* clash in sentences of the form:

(3) It is possible that not- p , but I am fairly certain that p .⁴⁶

And there is no reason to deny that sentences of type (3) can be an assertion of knowledge.⁴⁷ If so, DeRose's argument for (EP-K) fails, and (EP-K) is false. This, in turn, removes the objection to restricting the closure principle for knowledge, which is the move that underwrites the stable invariantist solution to the lottery paradox.⁴⁸ So, we can know some of the propositions that we know are entailed by propositions we know, but we can't know all of them. In particular, we cannot

know the conjunctive proposition that conjoins all of our separate bits of knowledge regarding the tickets that will lose the lottery.

2. *The Appeal of Skepticism*

The history of epistemology is, to a significant degree, shaped by the attempt to escape from skepticism. Some recent theories of knowledge, such as reliabilism, are able to do so with ease. A subject does not need to be aware that her faculties are reliable in order to enjoy the fruits of their use; all that matters is that her faculties *are* reliable.⁴⁹ If reliabilism is true, it is hard to understand why so many philosophers were so worried about skepticism. But we *do* understand why they were worried. Skepticism has an intuitive appeal that does not dissipate in the face of reliabilism. This in turn has made many philosophers suspicious of theories that are unable to explain why we do feel threatened by skepticism.⁵⁰

As we have seen, contextualism is a view devised with this theoretical burden in mind.⁵¹ If the argument in section I above is correct, though, contextualism is unable to carry this burden. The problem is that, whether or not the semantic value of 'knows' is context-sensitive, a plausible closure principle like (CP-SEP) will have skeptical scenarios be epistemically related to our ordinary knowledge claims. Assuming that our ordinary knowledge claims are held with excellent SEP, this means that we have an excellent, or at least good, SEP for the claim that we are not brains in vats.⁵² Setting aside contextualism, then, the consequence for stable invariantism is that we can know we are not brains in vats or the victims of evil demons. But if we can know this, why does skepticism still continue to worry us?

Recall that above we rejected (EP-K), the principle that connects epistemic possibility to knowledge. That principle can now be replaced with the following:

(EP-C) If S is certain that p , then it is not possible for S that not- p .⁵³

Notice that (EP-C) takes certainty to be a sufficient condition for the impossibility that not- p ; alternatively, we can read it as saying that uncertainty is a necessary condition for the possibility that p .

(EP-C) governs epistemic possibility in the widest sense. But we can limit what we regard as epistemically possible, from the wide boundaries set by certainty, by pegging it to whatever degree of knowledge suits our purposes in a given situation. In this way, we can preserve the conceptual link between knowledge and epistemic possibility. If fallibilism is correct, there is not a single cognitive relation between subjects and propositions that is knowledge. If there are many different grades of knowledge, there will be correspondingly many ways of determining the space of epistemic possibility. This space can be limited from the widest parameters, set by what one is certain about, when one treats some item of lower grade knowledge as if it were certain. For example, if I treat as certain the proposition that my guidebook is fully reliable, it will be impossible for me that the natural history museum is closed on Tuesdays. The subject's practical or theoretical interests will determine which grade of knowledge—and also which shape for the space of epistemic possibility—is appropriate in a given situation. But this is *not* to say, as the interest-relative invariantist does, that the subject's practical interests will determine whether her belief counts as knowledge at all.

Suppose, then, that I know with an average degree of certainty that the President is in Washington. When knowledge of that sort is all that matters, then I can say that it is not epistemically possible for me that the President is in New York. Given, however, that the knowledge in question is somewhat below the level of certainty, it can easily happen that this restriction on the space of epistemic possibility becomes undermined. My information on the President's whereabouts comes from *The New York Times*, which is ordinarily quite reliable.

Nevertheless, after a bit of reflection I realize that even the best newspapers occasionally make mistakes; it's possible that I will read a correction of the story in tomorrow's paper. But this realization does *not* undermine my original knowledge—it merely reminds me that I am not certain as to where the President is.

A change in one's purposes, whether motivated by practical or theoretical concerns, may render lesser degrees of knowledge unsuitable without altogether destroying them. The most effective way to do this is by considering ways in which a subject's SEP could fail. For example, in the case of a lottery, one may have an excellent SEP with respect to the proposition that a particular ticket will lose, and yet the ticket in question could win. In fact, it is because this is so obvious with lotteries that we are reluctant to ever say it is impossible for a ticket to lose.⁵⁴

Skeptical scenarios are also quite effective at undermining restrictions on the space of epistemic possibility by making it impossible to ignore situations in which our SEP fails. As a result, we never find ourselves in a situation where we can plausibly say that it is impossible for us to be brains in vats. Skeptical scenarios thus can never fully be put to rest. They show that much of our knowledge is forever open to question. The fallibilist will agree with all of this, merely adding at the end the reminder that, although open to question, it still *is* knowledge.

The “tether” of knowledge may thus be more tenuous than we would like. If fallibilism is correct, we must live with our doubts. Stable invariantism offers the simplest, most secure way of doing so.⁵⁵

REFERENCES

- Adler, Jonathan. 2002. *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Chisholm, Roderick. 1989. *Theory of Knowledge*, 3rd. ed. Englewood Cliffs, NJ: Prentice Hall.
- Cohen, Stewart. 1999. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives (Epistemology)* 13: 57-89.
- . 1998. "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery," *Australasian Journal of Philosophy* 76: 289-306.
- . 1988. "How to Be a Fallibilist," *Philosophical Perspectives* 2: 91-123.
- Cohen, Stewart, and Lehrer, Keith. 1983. "Justification, Truth, and Coherence," *Synthese* 55: 191-207.
- DeRose, Keith. 1999. "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology*, ed. by J. Greco and E. Sosa. Oxford: Blackwell.
- . 1995. "Solving the Skeptical Problem," *Philosophical Review* 104: 1-52.
- . 1991. "Epistemic Possibilities," *Philosophical Review* 100: 581-605.
- Dretske, Fred. 1981. "The Pragmatic Dimension of Knowledge," *Philosophical Studies* 40: 363-78.
- . 1970. "Epistemic Operators," *The Journal of Philosophy* 67: 1007-1023.
- Fantl, Jeremy and McGrath, Matthew. Forthcoming. "Knowledge and the Purely Epistemic: In Favor of Pragmatic Encroachment," *Philosophy and Phenomenological Research*.
- . 2002. "Evidence, Pragmatics, and Justification," *Philosophical Review* 111: 67-94.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Goldman, Alvin. 1992. *Liaisons*. Cambridge, MA: MIT Press.
- . 1979. "What Is Justified Belief?" in *Justification and Knowledge*, George Pappas (ed.). Dordrecht: Reidel. Reprinted in Goldman (1992), pp. 105-126.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hume, David. 1978. *A Treatise of Human Nature*, 2nd ed. (ed. by L.A. Selby-Bigge and P.H. Nidditch). Oxford: Clarendon Press.
- Kornblith, Hilary. 2000. "The Contextualist Evasion of Epistemology," *Philosophical Issues (Skepticism)*

10: 24-32.

Lackey, Jennifer. Forthcoming. "Norms of Assertion," *Noûs*.

Lewis, David. 1996. "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549-67.

Moore, G.E. 1959. *Philosophical Papers*, ed. by H.D. Lewis. London: George Allen & Unwin.

Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Owens, David. 2000. *Reason without Freedom*. London: Routledge.

Sosa, Ernest. 2000. "Skepticism and Contextualism," *Philosophical Issues (Skepticism)* 10: 1-18.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

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

Williams, Michael. 2001. "Contextualism, Externalism and Epistemic Standards," *Philosophical Studies* 103: 1-23.

¹ On the widespread commitment to fallibilism, see Cohen (1998, p. 91): the "acceptance of fallibilism in epistemology is virtually universal." Williams (2001, p. 5) agrees: "We are all fallibilists nowadays."

² Thus, for example, G.E. Moore refutes the skeptic merely by calling attention to his own hand; see his "Proof of an External World," in his (1959).

³ See Dretske (1970) and Nozick (1981).

⁴ For some of the most prominent defenses of contextualism, see Cohen (1988) and (1998), DeRose (1995) and (1999), and Lewis (1996). I shall ignore the differences to be found among these several versions of contextualism, as my concern here is with the fundamental thesis underlying all of them.

⁵ Here, let us suppose, Maria is an average student. The point is that she can come to know when the final will be held simply by reading my syllabus. She does not need to do anything further, such as confirm with the registrar or the Dean's office that my information about the schedule for finals is correct.

⁶ Contextualists tend to think that, not only would it be correct for me to say this, I (and most people) *would* say this when confronted with skeptical scenarios. But this seems doubtful to me. Although I am myself rather more sympathetic to skepticism than most, I do not find myself making attributions of ignorance in this way when discussing skeptical possibilities, nor does it seem to be the way my students talk. (On this point, see also Sosa (2000), p. 15.)

⁷ See my (XXXX) for an argument that the skeptic can get around this restriction by posing her challenge, not directly to the subject's beliefs, but to the attributor's prior knowledge-attributions.

⁸ *Strength of epistemic position* can be understood as a placeholder to be filled by one's preferred account of justification, warrant, etc.

⁹ This, and the other closure principles discussed below, need to be qualified so that they apply only to cases in which the subject competently infers Q from P and comes to believe Q on that basis. I omit that qualification above simply for convenience; nothing in what follows should be affected by that omission.

¹⁰ Although terms like 'excellent', 'good', etc., are sometimes used in a context-sensitive way, they also have a context-invariant use. For example, some high school textbooks have a form printed inside the cover asking the student to rate their condition—as *excellent*, *good*, *fair*, or *poor*—at the beginning of the schoolyear. The terms mark out intervals that are fixed for all conversational contexts. The closure principles formulated above should be understood, analogously, as context-invariant in this way.

¹¹ Even if SEP (i.e., evidence, reliability, etc.) does not vary from one context to another, might it not be the case that the evaluation of it—as *poor*, *good*, *excellent*, etc.—does vary? (I owe this suggestion to an anonymous referee.) If so, the contextualist could grant that our evaluation of a given subject's SEP as poor is correct while we (the attributors) are in a skeptical context, but this is compatible with our giving a different evaluation of that same SEP in another context. In particular, when we are in an ordinary context, we might correctly evaluate the subject's SEP as good or even excellent. Thus, there would be no more embarrassment in our admitting, in some contexts, that a subject's SEP is poor than there is in asserting, in some contexts, that the subject does not have knowledge.

By way of reply, let us grant that evaluative terms like ‘excellent’, ‘good’, ‘poor’, etc., are very often used in a context-sensitive way. A tennis instructor might call out, “Excellent!” when one of her beginning students makes an overhead shot, even though she would say, “Poorly done!” when an advanced student makes the same shot. But if the above contextualist objection is to succeed, it must be the case that these evaluative terms are not merely context-sensitive on *some* occasions—they must be so on *every* occasion of use. For if they do have a context-invariant use, there is nothing to prevent the skeptic from simply stipulating that these evaluative terms, when they appear in the closure principle governing SEP, are in fact being used in a context-invariant way. (The fact that the contextualist might be able to formulate a closure principle governing SEP with evaluative terms that are context-sensitive would thus be irrelevant to the skeptic.) The skeptic’s conclusion would be that, with respect to any randomly chosen belief a subject may have, the subject’s SEP for that belief is poor, where ‘poor’ is understood in a context-invariant way. The contextualist would then be forced to make the difficult choice between allowing that a subject can have knowledge even with a poor SEP and granting that (because her SEP is poor) the subject does not have knowledge after all.

The question, then, is whether there are context-invariant uses of the evaluative terms that appear in the SEP closure principle. The answer is clearly yes. Readers may be familiar with one example from their student days. Textbooks that are loaned to students year after year are often printed with a form inside the cover that asks the student, prior to receiving the book, to rate its condition as *excellent*, *good*, *fair*, or *poor*, where there are specific guidelines governing the use of these evaluative terms (e.g., the book is to be rated in poor condition if the spine is cracked and the pages are marked). Given these guidelines, there should be relatively little variation in the evaluation of a book from one context to another (though this is of course compatible with using the same evaluative term in both a context-invariant and a context-sensitive way in the same conversational context; when the speaker then shifts to a different context, she may use a different evaluative term in the context-sensitive sense, but she should continue to use the same term in its context-invariant sense). Something similar can be said, I believe, for the terms of epistemic evaluation—e.g., *probable*, *evident*, *certain*, etc.—that feature in Roderick Chisholm’s *Theory of Knowledge* (1989, see especially pp. 9-12).

Given the way he explicates the terms, it seems clear he intends them to be understood in a context-invariant way.

¹² Some philosophers may be tempted at this point to adopt some form of externalism as a way of responding to the new skeptical argument. But, if that is your answer here, it may as well have been the response you gave to the original skeptical argument; for a similar point, see Kornblith (2000). In any case, the special advantage of contextualism—its ability to honor the intuitive appeal of skepticism while still allowing us to resist it—has now been lost with respect to the new argument.

¹³ Interest-relative invariantism has been defended most rigorously by Fantl and McGrath (2002) and (forthcoming), Hawthorne (2004), and Stanley (2005). See also Owens (2000), p. 25-7, and Adler (2002), pp. 238-47. I follow Stanley in calling the basic view in question *interest-relative invariantism*. Hawthorne's version of the view is called *sensitive moderate invariantism* or *subject-sensitive invariantism*, but it includes a condition governing error salience that other interest-relative invariantists may not accept. The claim that knowledge is grounded in purely epistemic features of the subject is called *purism* by Fantl and McGrath (forthcoming) and *intellectualism* by Stanley (2005), p. 6.

¹⁴ The following cases are drawn from Fantl and McGrath (2002), pp. 67-8. See also Cohen (1999), p. 58, and Stanley (2005), pp. 3-5.

¹⁵ See Fantl and McGrath (2002), p. 72. According to Stanley (2005), p. 9, "it is immensely plausible to take knowledge to be constitutively connected to action, in the sense that *one should act only on what one knows*." Where Fantl and McGrath take the possession of knowledge to be a *sufficient* condition for rational action, Stanley apparently takes knowledge to be *necessary* for rational action. But the necessity claim is quite implausible. Suppose I urgently need to locate my brother within the next two minutes. I have some reason to believe that he has gone to the Art Institute, but my reason for believing this is not sufficient for me to know that he is there. Clearly, it would be rational for me to go search for my brother there, even though this violates the necessity claim. (To this one might reply that I at least know that searching at the Art Institute is the best thing I can do in the circumstances. But this need not be the case—I might be so pressed for time that I simply act without being in a position to know that this is the right course of action.) Moreover,

Stanley's bank cases (analogous to the train cases presented above), which are meant to underwrite his view, tend to support the sufficiency, rather than the necessity, claim. For these reasons, I shall focus on Stanley's arguments as they relate to the sufficiency claim.

¹⁶ This is a compressed version of the arguments given in Fantl and McGrath (2002), pp. 72-3, and (forthcoming), pp. 8-9.

¹⁷ Fantl and McGrath (forthcoming), p. 2. The principle they derive in their (2002), p. 73, says that *S* knows that *p* only if, for any act *A*, if *S* knows that if *p*, then *A* is the best thing she can do in light of all her goals, then *S* is rational to do *A*. See Stanley (2005), pp. 89-90, for his account of knowledge, which incorporates a practical condition (but, it is worth noting, not a condition that handles the Gettier problem).

¹⁸ That is to say, our knowledge that *p*, comes with an epistemic grounding that may fail. For that reason, we cannot treat it as an unquestioned assumption that will always be capable of guiding our deliberation. I shall say more about fallibilism below.

¹⁹ Fantl and McGrath (forthcoming), pp. 13-20.

²⁰ See Fantl and McGrath (forthcoming), pp. 19-20.

²¹ And, again, the supposed absurdity of the attempted cancellation can be explained by the fact that 'worse' is ambiguous. One option can be worse than another given *p*; alternatively, one option can be worse than another in light of the expected utilities. The attempted cancellation genuinely is absurd only if we mean 'worse' in the latter sense.

²² The person would not know, of course, in the first situation because the necessary practical condition on knowledge would not be met. But, in the second situation, the practical condition *would* be met, as would all of the purely epistemic necessary conditions. Hence, there would be no reason to deny that the person has knowledge in the second situation. This, I take it, would have to be the response of a straightforward interest-relative invariantist. But some interest-relative invariantists, like Hawthorne, take the *salience* of error possibilities to determine in part whether a subject's strength of epistemic position is sufficiently good to allow her to possess knowledge. Practical considerations can become relevant by making a broader range of error possibilities salient. In punishment/reward case 2, then, the subject doesn't have knowledge in one

scenario because the stakes are high, thereby making a broad range of error possibilities salient. If they are salient when considering the high stakes scenario, they should be salient when the subject is simultaneously considering the low stakes scenario. Hence, the subject would not have knowledge with respect to either scenario. Although this would avoid the need for relativizing knowledge to practical context, it would do so only by taking on the highly counterintuitive result that the subject fails to have knowledge in the low stakes scenario. This is implausible on its face, and it also leaves the interest-relative invariantist no way of explaining why the subject feels perfectly warranted in acting as if she has knowledge in that situation.

²³ On this point, see Hawthorne (2004), p. 180.

²⁴ Hawthorne (2004), p. 176.

²⁵ If practical concerns *do* affect the quality of the subject's evidence or the performance of her cognitive faculties, the subject's ignorance can be explained in a purely epistemic way. The subject does not know that *p* simply because her SEP is not good enough. The causal explanation as to *why* her SEP is not good enough is irrelevant (from a strictly epistemic—but not practical—point of view).

²⁶ This is the same principle to which I earlier referred as a sufficiency claim: knowing that *p* is sufficient for rational action as if *p*. This is of course equivalent to the necessity claim: rational action as if *p* is necessary for knowing that *p*.

²⁷ Notice, by contrast, that there is no problem here for purely epistemic theories. Given an account of knowledge that does not include a necessary practical condition, the claim that one is irrational to act as if *p* only if one doesn't know that *p* can be taken as an informative claim—its informativeness derives from the independent conception of knowledge. (Nevertheless, because the claim in question is false, as I have argued above, I would not advise any purists to endorse it. The point is merely that they *could* do so.)

²⁸ I shall say more about this below.

²⁹ As should be clear from the arguments above, though, I do not think there is such a relation of entailment, even in the case of 'knowledge' claims.

³⁰ Fallibilism is usually taken to be the view that the subject's SEP does not entail the truth of the belief (even when it is sufficient for knowledge). But, on that understanding of fallibilism, it is impossible to have fallible

knowledge of necessary truths, as anything—including the subject's SEP—entails the truth of a necessary proposition. According to a fully general fallibilism, then, what matters is not whether an entailment relation holds between the subject's SEP and the truth of the belief but rather whether there is a probabilistic relation between the SEP and the belief. See my (XXXX) for a fuller statement of fallibilism.

³¹ It is important to note that to claim fallibilistic knowledge is probabilistic is *not* to claim that it is knowledge *of* probabilities. One could, obviously, have probabilistic knowledge of non-probabilistic propositions and infallible knowledge of probabilistic propositions.

³² Some philosophers resist the claim that Millie can know her lottery ticket will lose on the grounds that merely statistical evidence—e.g., knowing the odds in a lottery—can never put one in a position to know that an improbable fact does not actually obtain. This seems doubtful to me. I find it intuitively plausible to say that I know a gorilla is not, at this very moment, climbing to the top of the antenna on the Sears Tower, even though I have no relevant perceptual or testimonial evidence. I know it simply because it is so incredibly unlikely to be happening. But, in any case, this issue can be sidestepped. If necessary, we can assume that Millie has seen a lot of lottery tickets lose but has never or rarely seen one win. The situation is quite similar to my knowing that my friend will receive the letter I mailed yesterday. I have observed many letters being correctly delivered and have never or rarely seen one lost. If inductive evidence of this sort is perfectly good with respect to the postal system, I see no reason to think it isn't good enough in the case of lotteries.

³³ I should note that there is no plausibility to the suggestion that her true belief, that ticket 1000 will be the winner, is justified and fails to be knowledge only because it doesn't satisfy an anti-Gettier condition.

³⁴ See Vogel (1990) and Hawthorne (2004), ch. 1.

³⁵ Here is one way in which a skeptic might create an argument that parallels the lottery situation. Suppose that, for an otherwise normally functioning human, there is a one in a million chance she will have a neurological glitch that causes her to misperceive something in her environment and thereby form a false belief. (I don't know what the probability actually is; it may be higher or it may be lower. But, given that our cognitive equipment is not infallible, it is clear that there is *some* probability that a normal human will have a malfunction at any given moment.) So, when one appears to see, say, an oak tree (in good light, from close

by, etc.), it is overwhelmingly likely that one's belief regarding the oak is true. But, if the high probability that one's lottery ticket will lose is not good enough for knowledge, then neither is the high probability that one's visual experience is accurate good enough for knowledge. (The fact that one does not infer the truth of the perceptual belief from its high probability is irrelevant. As I mentioned in an earlier note, the belief that one's lottery ticket will lose can be the result of an inductive inference, rather than an inference drawn merely from the odds.)

³⁶ Providing a precise account of which closure principles are acceptable and which are not lies outside the scope of this paper. But I think a promising method would begin by formulating closure principles for SEP as two or more premises are brought into inferential contact. Those closure principles for SEP that yield a SEP above the threshold for knowledge would then generate closure principles for knowledge. Hawthorne distinguishes between Single Premise Closure and Multi-Premise Closure, where these principles are not sensitive to SEP (pp. 33-4). My solution to the lottery paradox will require rejecting Multi-Premise Closure. I suspect that Single Premise Closure, in his formulation, will have to be rejected as well: when a conditional and its antecedent are each known with a SEP just sufficient to cross the threshold for knowledge, the SEP for the consequent will slip below the lower boundary of knowledge.

³⁷ Suppose that Millie believes of each losing ticket that it will lose while also believing that there will be one winning ticket. Although her belief set thereby contains no logical contradiction, it nevertheless will be probabilistically incoherent (I owe this point to Ernie Sosa). Where someone who accepts a contradiction is guaranteed to have a false belief, a subject with probabilistically incoherent beliefs is highly likely to be wrong in at least one of her beliefs. Although it is generally a good idea to avoid both accepting contradictions and having probabilistically incoherent belief sets, I think this is not a significant problem for my solution to the lottery paradox. First, it is epistemically reasonable to recognize one's fallibility. For that one reason, it is plausible to say that one ought to have one or both of the following beliefs: (i) at least one of my beliefs is false; (ii) it is highly likely that at least one of my beliefs is false. Accepting the former leads to a contradiction; accepting the latter leads one to a probabilistically incoherent set of beliefs. Arguably, either result is something that should be expected for a reflective, fallible subject. Second, accepting contradictions

or probabilistically incoherent beliefs is most obviously problematic when one does this explicitly. But this does not happen when one accepts either (i) or (ii) above, given that one is not in a position to identify which of one's beliefs is actually false. Moreover, it does not happen in the lottery case, either. Given the restrictions on closure principles, one is unable to properly draw the inferences that would allow one to make the probabilistic incoherence explicit.

³⁸ See Hawthorne (2004), pp. 47-8, for the connection between this principle and a defense of his Multi-Premise Closure principle. He formulates it as the *Epistemic Possibility Constraint*: “if the epistemic probability for S that p is not zero, then S does not know that not- p ” (p. 111).

³⁹ *Treatise*, Liv.1. See Hawthorne (2004), p. 25.

⁴⁰ Indeed, Hume uses the claim that knowledge and probability are essentially different as a premise in his argument for skepticism regarding reason.

⁴¹ DeRose (1991), pp. 600-1. He calls (EP-K) “Moore’s Principle” because G.E. Moore endorses it in “Certainty” (1959, p. 241).

⁴² I should also note that (a) is rendered irrelevant if a type-(1) sentence does not contain a flat out assertion. Moreover, (a) may be too strong. Lackey (forthcoming) argues that knowledge is not the norm of assertion—that is to say, she argues that the principle that one should assert only what one knows is incorrect. In its place, she substitutes the weaker principle that one should assert only what it is reasonable for one to believe. If this is correct, then even in cases of flat out assertion, we should not take the speaker to be representing herself as having knowledge.

⁴³ This is the element of truth in Dretske’s claim that knowledge “is like being pregnant: an all or nothing affair” (1981, p. 363).

⁴⁴ In the same way, a woman carrying twins into the ninth month is clearly more pregnant than a woman with one six-week-old embryo. This is the sense in which Dretske’s claim is false.

⁴⁵ This is one very good reason to be skeptical of epistemological views that rely too narrowly on intuitions about ‘knowledge’ *claims* rather than on intuitions about *knowledge* itself. Having the status of knowledge is a *determinable* property of beliefs. Much of the time, our purposes are served simply by attributing the

determinable property to a particular belief (or to the subject who has the belief). However, some contexts call instead for the attribution of some *determinate* value of knowledge; this can easily happen when theoretical or practical circumstances require us to pay careful attention to the specific level of SEP underwriting a particular instance of knowledge. Attributions of knowledge *simpliciter* are then too broad to be useful—and may in fact be misleading. Hence, we switch to attributions of some degree of certainty (e.g., being *sure* or *pretty sure*), or we make an admission of some degree of doubt, where this can still be small enough to be compatible with knowledge. As we saw above, this is what happens in the legal system: the jury does not attempt to determine whether it knows *simpliciter* the defendant committed the crime in question; rather, the jury tries to determine *beyond a reasonable doubt* whether the defendant is guilty.

⁴⁶ The fact that there is no clash is made clear by example: “Sure, it is possible the back door is unlocked, but I am fairly certain that I locked it.”

⁴⁷ Quibblers who hesitate over (3) will surely accept the following sentence form as a way of asserting knowledge: (4) There is the remotest possibility that not-*p*, but I am virtually certain that *p*. To reject even this is to be not a quibbler but a skeptic. It should be clear, moreover, that sentences of type (3) are not merely expressions of *subjective* (rather than *epistemic*) certainty. When there is a challenge to my utterance, “I’m sure I turned the oven off,” the proper response is not for me to talk about how confident I am in the belief. Rather, I ought to give my evidence for being so sure—e.g., I might mention that I remember seeing the indicator light had gone off.

⁴⁸ As I mentioned earlier, one of the strongest objections to the “tracking” view defended by Dretske (1970) and Nozick (1981) has been that the view allows for violations of the closure principle; see the presentation of Saul Kripke’s unpublished counterexamples in Feldman (2003), pp. 89-90, and see also DeRose (1995). Given that I have rejected outright closure principles, does my solution also face those objections? What Kripke objects to are situations like the following. Suppose I am in an area where lots of fake barns have been erected, and there is only one genuine barn. The genuine barn is red, but all of the fakes are yellow; still, I do not know this and cannot tell the genuine barn from the fakes. Given the tracking view, when I see the genuine barn, I am in a position to know that there is a red barn in front of me, but I am *not* in a position to

know that there is a barn in front of me—only the first belief tracks the truth. This is counterintuitive because the inference from the red-barn belief to the barn belief is so obvious. It is important to see, however, that the restrictions on closure I have been defending do not lead to implausible results like this. They prohibit only inferences where the probabilities of the premises are not sufficiently high to yield a conclusion that is probable enough to surpass the threshold for knowledge. Kripke's counterexample is damaging to the tracking account precisely because the inference is one the subject knows with near certainty; for that reason, the premise and the conclusion should have nearly identical epistemic properties.

⁴⁹ For the classic formulation of reliabilism, see the papers collected in Goldman's (1992), especially his (1979).

⁵⁰ See, for example, Cohen and Lehrer (1983) for the "new evil demon" problem. See also Fumerton (1995).

⁵¹ It is worth noting that contextualists are primarily concerned with skepticism grounded in skeptical scenarios. There are, however, other skeptical arguments—e.g., what the ancient skeptics called "the problem of the criterion"—that require a different treatment (I am grateful to Ernie Sosa for reminding me of this point).

⁵² Any particular version of fallibilism, including contextualism, interest-relative invariantism, and stable invariantism, is going to include the anti-skeptical assumption that many ordinary knowledge attributions are correct. My purpose here is not to defend that assumption but simply to show how stable invariantism is well-equipped to explain the appeal of skepticism (at least, the skepticism that is grounded in the skeptical scenarios). Hence, the contextualist's claim that it is a unique virtue of their view that it can acknowledge the force of this kind of skepticism is incorrect.

⁵³ Here, the relevant sort of certainty is epistemic, not psychological, in nature.

⁵⁴ Something similar can be said about the belief that my car has not been stolen from the parking lot (see Vogel (1990)). It is the very nature of the belief to remind me that my SEP for it could fail. None of this means, however, that the belief is not knowledge.

⁵⁵ Acknowledgments.