

Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning

JESSICA BROWN

University of St. Andrews

1. Introduction

It is increasingly popular to suggest that knowledge is the norm of practical reasoning, or reasoning about what to do (e.g. Hawthorne 2004, Stanley 2005). This idea is central to the defence of a new version of invariantism—‘subject-sensitive invariantism’—on which whether the true belief that *p* is knowledge not only depends on such factors as one’s evidence, and the reliability of the belief-producing process, but also the stakes or how important it is that *p* be true (the view is also known as ‘sensitive moderate invariantism’ (Hawthorne 2004) and ‘interest relative invariantism’ (Stanley 2005)). I will argue against the idea that knowledge is the norm of practical reasoning, whether that is understood as a necessity or sufficiency claim. Instead, I will argue that the epistemic standards for practical reasoning vary contextually.

2. SSI and the Knowledge Norm

Like other versions of invariantism, subject-sensitive invariantism (SSI) denies the contextualist claim that ‘know’ is a context-sensitive term which expresses different properties in different contexts. Rather, ‘know’ expresses a univocal property. SSI is distinguished from traditional versions of invariantism by the range of factors that it allows affect the truth of knowledge ascriptions. Traditional versions of invariantism hold that whether *S*’s true belief that *p* is knowledge depends on such factors as *S*’s evidence for *p*, and whether *S*’s belief that *p* was formed reliably. SSI is distinctive in holding that

whether S knows that p depends on the stakes for S, on how important it is to S that p be true.¹

SSI has been defended in part by appealing to the idea that knowledge is the norm of practical reasoning, or reasoning about what to do (see Hawthorne 2004, Stanley 2005). The idea that knowledge is the norm of practical reasoning can be understood in several ways, as either a claim of necessity (e.g. Hawthorne 2004, Stanley 2005) and/or sufficiency (Fantl and McGrath forthcoming, Hawthorne and Stanley forthcoming).² It will be useful for us to distinguish the following claims:

NEC: if it is appropriate to rely on p in practical reasoning³ then you know that p.⁴

SUFF: if you know that p then it is appropriate⁵ to rely on p in practical reasoning.

In fact, SUFF needs restricting to what Hawthorne and Stanley call ‘p-dependent choices’.⁶ Since the cases considered in the paper involve p-dependent choices, we will continue to work with the simpler claim.

The bi-conditional, KN, is the conjunction of NEC and SUFF:

KN: it is appropriate to rely on p in practical reasoning iff you know that p

In the rest of the paper, I will consider the main arguments for KN and consider whether they support SUFF or NEC. Two main considerations have been used to support KN, certain examples of unacceptable reasoning and the use of knowledge in defending and criticising action. If KN is to function as a premise in an argument for SSI, then it is illegitimate to appeal to SSI in order to defend KN; rather we are looking for a defence of KN which does not rely on SSI. This will turn out to be important in assessing the case for KN and particularly SUFF.

3. Unacceptable Reasoning

One argument for KN arises from certain cases of intuitively unacceptable reasoning (Hawthorne 2004: 174–75). For instance, suppose that you have bought a £1 ticket in a lottery in which there are 10,000 tickets and the prize is £5000. The draw has taken place although it has not yet been announced; in fact your ticket has lost. You truly believe that your ticket has lost on probabilistic grounds. Someone offers you 1p for the lottery ticket. It seems intuitively unacceptable for you to reason as follows:

I will lose the lottery.

If I keep the ticket I will get nothing.

If I sell the ticket, I will get 1p.

So I ought to sell the ticket.

(By contrast, the reasoning would seem acceptable if you had heard the announcement of the lottery.)

Similarly, the following reasoning concerning an offer of life insurance seems unacceptable:

I will be going to Blackpool next year.

So I won't die beforehand.

So I ought to wait until next year before buying life insurance.

Hawthorne (2004) suggests that in these and other cases, the reasoning seems unacceptable because, in each case, the first premise is not known (176). If this diagnosis is correct, then it provides some support to NEC. However, the case is hardly conclusive since there may be alternative possible explanations of the unacceptability of the reasoning. Further, the cases offer no support to SUFF.

4. Criticism and Defence of Action

KN might seem to be supported by the way in which knowledge is cited in the defence and criticism of action (Stanley 2005, Stanley and Hawthorne forthcoming, Fantl and McGrath forthcoming). For instance, a mother may criticise her teenage daughter, saying 'You shouldn't have left the party so late; you didn't know that there would be a bus at that time'. Or I may defend the fact that I'm still in the office at noon when I have arranged to meet my partner for lunch at 1pm by pointing out that I know that there's an express train at 12.20pm (the express arrives in time for me to make our lunch meeting). In these and similar examples, the relevant defence and criticism focuses on action whereas KN concerns when it's appropriate to rely on a proposition in practical reasoning. Intuitions about the correctness of action do not necessarily reflect the appropriateness of the reasoning. For one may do the right action on the grounds of inappropriate reasoning. Still, it may be said that in our cases what's implicitly being assessed is the reasoning which lies behind the relevant action. For instance, it may be said that the mother is implicitly criticising her daughter for relying on a proposition which she didn't know—that there would be a bus at the relevant time—in deciding what to do. In this way, the cases may be thought to support KN. In more detail, cases such as that of the teenage daughter may seem to support NEC. If it's appropriate to rely on *p* in practical reasoning only if one knows that *p* that would explain why one's criticisable for acting on propositions one doesn't know. By contrast, other examples may seem to support SUFF. Consider

again the case in which I defend my remaining in the office until noon by pointing out that I know that there's an express train at 12.20pm. If knowing *p* is sufficient for it's being appropriate to rely on *p* in practical reasoning, that would explain why I can defend acting on the proposition that there's an express at 12.20pm by pointing out that I know that proposition.⁷ In other examples, I'm criticised for failing to act on a known proposition. Suppose this time that my partner and I have agreed to meet some friends for a drink in a pub at 6pm. We both know that the only train that will allow me to make this arrangement is the 5.20pm. My partner rings me up to discover I'm still in the office at 5.20pm. He berates me, saying 'You know the express is at 5.20pm. You should have left by now'. This kind of case may seem to support SUFF. If knowing *p* is sufficient for it's being appropriate to rely on *p* in practical reasoning, that would explain why I'm criticisable for not acting on a known proposition.

Although the use of knowledge in defending and criticising action is explicable on the assumption of KN, that use does not establish KN. For, we may cite a factor in defending and criticising action even if that factor is neither necessary nor sufficient for appropriate action. To see this, consider a different example of the justification of action. Suppose that in dividing up the stew I give Alison the biggest portion. I might defend my action by pointing out that Alison is training for a marathon. If, on the other hand, I had instead given the largest portion to Sarah, Alison might have criticised me saying 'You should have given me the biggest portion. I'm training for the marathon', or 'You shouldn't have given Sarah the biggest portion. She's not training for the marathon'. In certain circumstances, the suggested defence and criticism may seem appropriate. Even so, this doesn't show that training for a marathon is either necessary or sufficient for being justified in getting the largest portion of the stew: not necessary since other kinds of factors could justify getting the biggest portion (e.g. Sarah's pregnant), and not sufficient since other factors could trump training for a marathon (Sarah's starving). Analogously, that knowledge is cited in the defence and criticism of action does not show that knowledge is either necessary or sufficient for relying on *p* in practical reasoning.

This conclusion is reinforced when we consider the ways in which we criticise and defend actions in more detail. Although we sometimes defend and criticise actions by citing the absence or presence of knowledge, we also do so by citing conditions both weaker and stronger than knowledge. For instance, in defending why I'm still at the office when we're due to meet for lunch at 1pm I might cite facts about the train schedule ('There's an express at 12.20 pm'), or my evidence about the train schedule ('The timetable says there's a train at 12.20pm'; or, 'Peter said there's a train at 12.20pm'). That we cite factors weaker than knowledge in defending action counts against NEC. Similarly, although we sometimes criticise action by pointing out that the agent knew a relevant claim, we also do so by pointing to factors weaker

than knowledge. For instance, suppose I leave my partner a note saying we're out of potatoes. My partner returns home from work having stopped in the grocery store to buy beer, but doesn't get any potatoes. I may criticise his action saying, 'You should have got potatoes. You knew we don't have any'. His action is equally criticisable if it turns out that we do have potatoes since, by chance, a neighbour happened to bring some over from her garden as a surprise for us (although, of course, in such a situation, I would phrase my criticism differently, e.g. 'You should have got potatoes. You thought we didn't have any'). That we offer such criticisms of actions involving conditions weaker than knowledge counts against NEC. While we sometimes criticise actions by citing conditions weaker than knowledge we sometimes criticise actions by citing conditions stronger than, or orthogonal to, knowledge. For instance, a mother might criticise her teenager's late departure from a party by saying, 'You shouldn't have left so late. You weren't certain there'd be a bus home that late', or 'You shouldn't have left so late. You didn't know for sure that there'd be a bus at that time'. On the standard assumption that knowledge does not entail certainty, these kind of criticisms count against SUFF.

It seems, then, that when we look at the way we criticise and defend actions in more detail, we do so by citing factors both weaker and stronger than knowledge. It may be that our practice of citing factors other than knowledge can be shown to be consistent with KN.⁸ Nevertheless, the point remains that our ordinary practice of assessing actions is not obviously suggestive of KN. Rather, it suggests that the standard for practical reasoning varies with context: sometimes the standard is knowledge, sometimes it is less than knowledge, and sometimes it is more than knowledge. This alternative can easily explain the use of knowledge in defending and criticising action even while denying both NEC and SUFF. Knowledge is cited in defending and criticising actions since, in many contexts, knowledge is the standard for relying on propositions in practical reasoning. It does not follow that, in all contexts, knowledge is the standard. In the following sections, I support this alternative view by arguing against NEC and SUFF.

5. Knowledge isn't Necessary

According to NEC, if it is appropriate to rely on *p* in practical reasoning then you know that *p*. On this view, if one does not know that *p*, it is inappropriate to rely on *p* in one's practical reasoning. We can test this view by considering cases in which one lacks knowledge that *p*, say because although one has a justified true belief that *p* it fails to be knowledge because one's in a Gettier situation. Suppose, then, that *S* leaves the office at 12.00pm in order to meet her partner for lunch at 1pm. *S* believes truly that there is an express train at 12.20pm which would allow her to arrive in time to make lunch. Further, this belief is justified: *S* checked the train timetable on the internet just before

leaving the office. In fact, unbeknownst to S, she is in a Gettier situation: a hacker has got into the train website and for a joke has replaced all of the current timetables with last season's timetables. Luckily for S, according to both the old and new timetables, there is an express at 12.20pm. So, S's belief is a case of true justified belief but not knowledge. Let's now consider whether it's appropriate for S to rely on her belief that there's an express at 12.20pm in deciding what to do. It seems that it is appropriate. After all, S checked the timetable and she had no reason to suppose that a hacker had replaced the current timetable with last season's. In such circumstances, it seems appropriate for her to rely on the belief that there is an express at 12.20pm in deciding what to do. This is reflected in our practices of assessment. For instance, if S's partner comes to know about the hacking, it will not seem appropriate for him to criticise S's action. For instance, he would not say 'You shouldn't have left so late. You didn't know there was an express at 12.20pm'. More generally, it seems that if S and her partner discover the hacking and so come to find out that S's belief was not a case of knowledge after all, this would not change their view about the appropriateness of S's reasoning. It is just as appropriate for S to rely on her belief that there's an express at 12.20 in a Gettier case in which her belief is not knowledge as it would be in a non-Gettier case in which her belief is knowledge. It seems, then, that NEC conflicts with our intuitions about Gettier cases.

A defender of NEC may try to accommodate Gettier cases by explaining away the intuition that it is appropriate for the subject to rely on the relevant proposition in practical reasoning. For instance, she may remind us that we need to distinguish between the assessment of the action and the reasoning and she may suggest that the intuitive sense of appropriateness derives from the assessment of the action not the reasoning. In our example, since S forms a true belief that there is an express at 12.20pm, the action she bases on this belief is successful: she does make her lunch arrangement. Further, we may suppose that, even if S had not been the subject of a Gettier case and had consulted the current timetable, she would not have acted differently. It seems, then, that her action was appropriate. A defender of NEC may suggest that we confuse the appropriateness of S's action for the appropriateness of her reasoning. However this reply is unconvincing. Even once we distinguish the evaluation of the action and the reasoning, it still seems that it is appropriate for S to reason from the claim that there is an express at 12.20pm. It seems incorrect to say that S shouldn't have relied on the claim that there is an express at 12.20pm, or that she shouldn't have used this claim as a premise in her reasoning.

The defender of NEC may offer a different way of explaining away the intuition of appropriateness. She may suggest that we are confusing the evaluation of the agent and the reasoning (Hawthorne and Stanley forthcoming). In the Gettier case, the subject is clearly blameless for relying on the claim that there is an express at 12.20pm. She did not realise that a hacking

attempt had taken place and had no reason to suspect that it had. Perhaps, we confuse the blamelessness of the agent for the different claim that her reasoning was appropriate? Again, this suggestion is unconvincing. On the suggestion, although the agent is blameless in relying on the claim that there is an express at 12.20pm, it is nonetheless inappropriate for her to rely on this claim. But, it seems that this suggestion does not correctly describe the case. It just does seem appropriate for S to rely on the relevant claim in her reasoning.

One way of reinforcing this conclusion is to consider a different form of criticism of action in which one criticises an agent for failing to act on a belief which is knowledge. For instance, in our earlier example, I criticise my partner for failing to get potatoes on his way home. In the example, he calls in at the supermarket to get beer but fails to buy any potatoes, despite knowing that we were out. I criticise his action saying: 'You should have got potatoes. You knew we were out'. It seems that this criticism would not be undermined if it turned out that he didn't know we were out of potatoes but was in a Gettier situation in which he merely had the justified true belief that we were out. For instance, suppose that it's a system between us for me to leave lists of urgently needed items on the kitchen table right where he'll see them. Knowing we're out of potatoes I put 'potatoes' on the list on the table. After I leave, he finds a shopping list on the table and forms the justified belief that we are out of potatoes. However, unbeknownst to us, the kids have decided to make up mock shopping lists for fun. What my partner picked up was one of their lists. As he has no reason to suspect that the kids are putting fake shopping lists around the house, he has a justified belief that we are out of potatoes. Further, the belief is true: we are out of potatoes. Discovery of the kids' prank in no way mitigates the negative assessment. For instance, it would be ridiculous for my partner to try to defend his action by pointing out that he didn't know that we lacked potatoes, but only had a justified belief that we lacked them. Further, it seems that after the kids' tricks have been discovered, I can still criticise his action by saying, 'You should have got potatoes. You thought they were on the shopping list'. The defender of NEC might attempt to deal with the example by claiming that we are confusing the assessment of the reasoning with that of the action or the agent. My partner's failure to buy potatoes is not the right action (we won't have anything to feed the kids), and he's surely blameworthy for getting beer but not potatoes. However, distinguishing these dimensions of assessment does not undermine the intuition that his reasoning is at fault: surely, given his evidence, he should have treated the proposition that we are out of potatoes as a premise in reasoning about what to do. He should have taken this proposition for granted in his practical reasoning.

It seems, then, that NEC is false. According to NEC, if it is appropriate for S to rely on p in her practical reasoning then she knows that p. But we have seen that it may be appropriate for a subject to rely on a proposition

in her practical reasoning even if she is in a Gettier case in which she has a justified true belief that *p* which does not amount to knowledge.

6. Knowledge is not Sufficient

I have suggested that the role of knowledge in defending and criticising action can be accommodated within a view on which the condition for appropriately relying on *p* in practical reasoning varies with context. In the last section, I argued that in some cases it is appropriate to rely on *p* in practical reasoning even though one does not know that *p*. Someone might accept that the condition for appropriate practical reasoning varies contextually in such a way that it may sometimes be less than knowledge, but still question how that condition could sometimes require more than knowledge. It might be said, if an agent knows that *p* what more could be required in order for her to rely on *p* in her practical reasoning? I now turn to defend the idea that knowing *p* is not always sufficient for relying on *p* in practical reasoning. I start by defending the idea that a subject who knows that *p* can strengthen her epistemic position further. This idea makes room for the thought that even if a subject knows that *p* this need not be sufficient for her to rely on *p* in practical reasoning; practical reasoning may require more than knowledge.

It seems part of our everyday concept of knowledge that it allows that two people can both know something, yet one be in a better epistemic position than the other. For example, two subjects may both know that *p* but one base her belief on a more reliable method of belief formation. Such cases commonly arise when both an expert and a non-expert know some claim but the expert's basis for the knowledge claim is superior. For instance, a doctor and a patient may both know that the patient has arthritis but the doctor's knowledge is based on a wider range of data, including various test results. Examples can also be constructed in which two non-experts both know a claim although one has a better ground for her judgement than the other. For instance, A and B may both know that the film starts at 7pm, although A bases her belief on the printed listings in the weekly paper, whereas B bases her belief on the listings and telephone confirmation with the venue.⁹

It is important to see that a subject-sensitive invariantist should have no objection to the idea that two subjects may know yet one be in a better epistemic position than the other. SSI is distinctive in allowing the stakes to affect knowledge. The stakes are a non-truth conducive factor; that a lot turns for the subject on whether *p* is the case does not make *p* more or less likely to be true. That SSI allows that the stakes affect whether one knows, and in some versions, whether one has evidence and justification,¹⁰ is compatible with acknowledging that knowers may differ with respect to their epistemic position. For we may construct such cases even for subjects for whom the stakes do not differ. For instance, our cinema goers may both be in the same low stakes situation in which nothing much turns on whether

the film starts at 7pm. A truly believes that the film starts at this time on the basis of the newspaper listing alone, while B truly believes it on the basis of the listing and telephone confirmation with the venue. Given the low stakes, a subject-sensitive invariantist would hold that one can know on a relatively slender basis. Thus, A knows that the films starts at 7pm. B faces the same low stakes. The only difference is that she believes that the film starts at 7pm on the basis of the listings and the telephone confirmation. It is hard to see how, if A and B face the same stakes, B's having telephone confirmation in addition to the newspaper listing could have the result that she does not know whereas A does. Thus, even a subject-sensitive invariantist should accept that, although A and B both know, B's epistemic position is stronger than A's. In describing B's epistemic position as being better than A's, we exploit one sense of 'epistemic' which Stanley distinguishes, that in which it denotes truth conducive factors (Stanley 2005: 2).¹¹ We can fill out the idea that B's epistemic position is better than A's in a variety of ways depending on what the subject-sensitive invariantist holds are the truth-conducive factors affecting knowledge: that B can rule out certain possibilities that A cannot, e.g. that there has been a change to the program since the paper was printed; that B's belief matches the facts in a wider range of worlds, e.g. the world in which there has been a schedule change since the paper was printed; that B's belief was based on a more reliable method of belief formation than A's, etc.

Once we see that it's part of our intuitive concept of knowledge that two subjects can both know that p yet one of them be in a stronger epistemic position than the other, this puts pressure on the idea that knowing p is always sufficient for relying on p in practical reasoning. If two subjects, A and B, both know that p, but B's epistemic position is better than A's, why should we think that mere knowledge is always sufficient for relying on p in practical reasoning? Although A and B share an important epistemic commonality—they both know that p—they also differ epistemically: B is in a better epistemic position than A. Perhaps it's this difference which is crucial for practical reasoning? On this view, knowledge is not always sufficient for practical reasoning. In order to defend this view, it is useful to examine a range of cases in which it seems intuitive that a subject knows that p but it does not seem appropriate for her to rely on p in her practical reasoning.

7. Counterexamples to Sufficiency

One way of generating cases in which it seems intuitive that a subject knows that p but it does not seem appropriate for her to rely on p in her practical reasoning is to focus on bets with extreme stakes or with heavily weighted odds. Take any claim which one would ordinarily regard oneself as knowing, say where one was born or one's name. One can construct a bet with such odds or stakes that it seems irrational for one to take the bet. Nonetheless it

seems that one knows the target claim. These cases seem to provide examples in which, intuitively, one knows that *p* but it would be inappropriate to rely on *p* in one's practical reasoning (Hawthorne 2004; Hawthorne and Stanley forthcoming).

For instance, suppose that Liz knows that she was born in England. Nonetheless, it would seem irrational for her to accept a bet with the following pay-offs:

Liz was born in England: Liz gains £1

Liz was not born in England: Liz loses her home

While this first example involves very high stakes, we can generate cases without high stakes, but instead involving heavily weighted odds. For instance, although Liz knows that she was born in England, she would be irrational to accept a bet with the following pay-offs:

Liz was born in England: Liz gains a million millionths of a penny

Liz was not born in England: Liz loses £3.

Such cases are not isolated. It seems plausible that, for any claim one knows, one can set up a bet with such stakes and/or odds that it is irrational for one to take the bet.

Other putative counterexamples to SUFF do not involve bets. Here are a range of examples.

SURGEON

A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

AFFAIR

A husband is berating his friend for not telling him that his wife has been having an affair even though the friend has known of the affair for weeks.

Husband: Why didn't you say she was having an affair? You've known for weeks.

Friend: Ok, I admit I knew, but it wouldn't have been right for me to say anything before I was absolutely sure. I knew the damage it would cause to your marriage.

Here the friend admits knowing but claims that it would have been inappropriate for him to act on that knowledge by telling the husband. Of course, there are a variety of non-epistemic reasons why a friend in this situation might not reveal the affair: perhaps the husband has been under such severe pressure recently that information about the affair might tip him over the edge into suicide or a breakdown; perhaps revealing the affair would involve breaking a promise to a third party; perhaps the friend is simply squeamish. We will stipulate that none of these factors apply to the case in hand. Nonetheless, the friend's statement seems perfectly intelligible and plausible: he is saying that although he knew of the affair his epistemic position wasn't strong enough to act on that knowledge. He needed to be absolutely certain before proceeding to inform the husband.

RESULT

Two lecturers have been in an exams meeting all morning at which the students' exam results are determined. After the meeting lecturer B bumps into one of the students and, on the basis of her recollection of the meeting, informs her that she has passed. On discovering this, lecturer A criticises her colleague for giving out information about the results without checking the exam pass list.

A: You shouldn't have told her she'd passed without having the list in front of you.

B: What's the problem? I knew she'd passed—I was in the examination meeting this morning.

A: That's not the point. Think of the damage it would cause if you gave the wrong result. We can't afford to take that risk.

Note that A doesn't reply by denying that B knows, but by saying 'that's not the point'.

Intuitively these cases put pressure on SUFF: in each one, a subject claims that either she, or a third party, knows something but that it would be inappropriate for her to act on that knowledge. Although the relevant evaluations explicitly concern action, it seems that they reflect claims about the underlying reasoning. For instance, the relevant intuition in SURGEON is that the surgeon should not rely on the premise that it is the left kidney which is affected in practical reasoning. For instance, she should not reason that since it is the left kidney which is affected, she should remove the left kidney straight-away rather than check first before operating. Similarly, the relevant intuition

in RESULT is that the lecturer should not rely on the claim that the relevant student has passed in her practical reasoning. For instance, she should not reason that since the student has passed she should tell her straightaway rather than first check the pass list.

One potential response to the alleged counterexamples to SUFF attempts to argue that any proposed account of the conditions for relying on *p* in practical reasoning will face analogous counterexamples. It seems plausible that, for any condition proposed as sufficient for relying on *p* in practical reasoning, if the stakes are sufficiently high, one could meet that condition and yet it be inappropriate for one to rely on that claim in practical reasoning. For instance, suppose that it's proposed that if one knows that one knows that *p*, then it is appropriate for one to rely on *p* in one's practical reasoning. Although this condition is stronger than knowledge, it seems open to an analogous objection. For instance, one can construct some bet with such stakes or odds that it would be irrational to accept the bet even if one knows that one knows. However, this reply is effective only against an opponent who claims that there is a condition sufficient in all circumstances for relying on *p* in practical reasoning. It is utterly ineffective against the kind of view outlined here according to which there is no such condition, but rather the epistemic position required for relying on *p* in practical reasoning varies with the context and practical reasoning situation.

Setting aside this first reply, it seems that there are two ways for a defender of SUFF to respond to putative counterexamples in which the subject knows that *p* but it is not appropriate for her to rely on *p* in practical reasoning. First, she might argue that, in fact, in the examples, the subject does not know that *p*. Second, she may accept that the subject does know that *p* and attempt to explain away the intuition that it is inappropriate for the subject to rely on *p* in practical reasoning. I will discuss each in turn.

The defender of SUFF might argue that in all the supposed counterexamples, the subject does not know the relevant proposition. She could do so by arguing that some of the conditions for knowledge are not met. For instance, she may argue that in the counterexamples, the stakes undermine the subject's belief and so the subject fails to know the relevant proposition. While high stakes may sometime undermine a subject's belief, it is not plausible that they always do so. So, it seems legitimate for it to be stipulated that the counterexamples involve a subject who truly believes the relevant proposition. In any case, the idea that the stakes undermine belief is an unpromising response for a defender of SSI. Subject-sensitive invariantists support their view by contextualist cases in which, they argue, the subject knows in Low but not High, although Low and High differ only in the stakes. If high stakes undermine the agent's belief, then one could explain why the subject knows in Low but not High by appeal to the difference in belief and without supposing that knowledge partly depends on the stakes.

While accepting that the putative counterexamples involve a subject who truly believes the target proposition, the defender of SUFF could argue that this true belief does not constitute knowledge. She might appeal to SSI to argue that, given the stakes, the relevant true belief does not amount to knowledge (Hawthorne and Stanley forthcoming: 22). For instance, in the bet example, she may argue that in virtue of being offered a bet with high stakes, Liz loses her knowledge that she was born in England. Notice that this reply is of no use in dealing with examples of bets which do not involve high stakes but rather heavily weighted odds (Hawthorne and Stanley forthcoming: 22). More importantly, such a reply is illegitimate given our dialectic. The aim was to see whether the knowledge norm, and in particular SUFF, could be used as a premise in an argument for SSI. Within this project, we need a defence of SUFF which does not itself rest on SSI.

Setting aside any worry about the stakes, it seems plausible that in each case the agent knows the relevant proposition: for instance, it seems plausible that an experienced surgeon can know which kidney is diseased by suitable examination of a patient together with test results, or that a lecturer can know that some student has passed an exam on the basis of her recollection of an examination board. Of course, the agent may not be able to distinguish her situation from a possible situation (the BAD situation) in which she makes a mistake. For instance, the lecturer cannot distinguish her situation from a possible situation in which she seems to recall that the student passed although she did not. However, the suggestion that knowledge requires that the agent can distinguish her actual situation from any possible BAD situation contains an obvious sceptical threat. As a result, such a response is not promising for SSI which is a non-sceptical position.

It seems, then, that the first of the two possible strategies for responding to the putative counterexamples to SUFF fails. Let us then turn to consider the second strategy. This strategy accepts that, in each of the putative counterexamples, the agent does know the relevant proposition and attempts to explain away the intuition that it is inappropriate for the agent to rely on her knowledge in practical reasoning. The suggestion is that we confuse intuitions about whether it's appropriate for the subject to rely on what she knows in practical reasoning with intuitions about some other aspect of the situation. I will consider several ways of filling out this strategy. On the first, we confuse intuitions about whether it's appropriate for the subject to rely on what she knows in practical reasoning with intuitions about whether she knows that she meets the sufficient condition for so relying. According to the second, we confuse intuitions about whether it's appropriate for the subject to rely on what she knows in practical reasoning with intuitions about the character traits she exhibits in so doing.

The first way of filling out the strategy exploits a failure of luminosity. Williamson (2000) argues that no condition is luminous, i.e. for any condition, C, there are cases in which C obtains but one is not in a position to know

that it obtains. According to SUFF, if the agent knows that *p*, then it is appropriate for her to rely on *p* in practical reasoning. By the failure of luminosity, there are cases in which the agent knows that *p* but does not know that she does. A defender of SUFF may argue that the combination of SUFF and the failure of luminosity predicts that in cases in which the agent knows but does not know that she does, our intuitions about the acceptability of the reasoning will be mixed or unclear: on the one hand the agent meets the condition for relying on *p* in practical reasoning but, on the other, she does not know that she does. Given this, we may have no simple intuitions about the acceptability of reasoning which relies on *p*. As Hawthorne and Stanley (forthcoming) put it, ‘intuitions go a little hazy in any situation that some candidate normative theory says is sufficient to make it that one ought to *F* but where, in the case described, one does not know that situation obtains’ (18). In particular, although the agent does meet the condition for relying on *p* in practical reasoning, it may seem that she does not because we confuse intuitions about the acceptability of the reasoning with intuitions about the agent and whether she was in a position to know that she meets the condition for relying on *p* in practical reasoning. For instance, consider again the case of extreme bets. The strategy now being considered accepts that Liz knows that she was born in England but attempts to explain away the intuition that it is inappropriate for her to rely on that claim in accepting the bet. On this view, since Liz does know that she was born in England, it is appropriate for her to rely on that claim in accepting the bet. However, it seems incorrectly that it is not acceptable for her to rely on that claim in reasoning to accept the bet because she does not know that she knows that she was born in England and so does not know that she meets the condition for appropriately relying on that claim in practical reasoning.

This reply is persuasive only if there is independent reason to suppose that the relevant counterexamples are cases in which the subject knows the target proposition but does not know that she does. Williamson’s anti-luminosity argument provides reason to suppose that *C* fails to be luminous for borderline cases of *C*, those which are close to the boundary between *C* and not-*C*. Suppose that in such a borderline case of *C* one correctly judges that *C* obtains. Intuitively, this belief is not knowledge for it’s not safe: it is too near borderline cases in which *C* does not obtain but in which one would still judge that it obtains. So, Williamson’s anti-luminosity argument provides reason to suppose that the putative counterexamples to SUFF instantiate a failure of luminosity if they concern borderline cases of knowledge. But, the problem for the defender of SUFF is that intuitively the counterexamples do not concern borderline cases of knowledge. For instance, the target proposition in the case of extreme bets—that one knows one’s name or where one is born—is not intuitively a borderline case of knowledge. Indeed, part of the interest of extreme bets is that they seem applicable even to the most secure or paradigm cases of knowledge. Similarly, when a surgeon bases her belief that

it is the left kidney which is diseased on the basis of a careful examination of the patient and test results, this does not appear to be a borderline case of knowledge. Given this, there is no reason to suppose that in all the putative counterexamples, the subject knows the target proposition without knowing that she does. So, the anti-luminosity defence of SUFF fails.¹²

According to the second way of explaining away the intuition that the relevant practical reasoning is inappropriate, we confuse intuitions about the appropriateness of the practical reasoning with intuitions about the epistemic character displayed by the reasoning. On this view, one does know the target proposition and it is appropriate to rely on that proposition in practical reasoning but we mistakenly think that it's inappropriate since, in relying on that proposition, one exhibits a character trait that is an epistemic flaw. For instance, in the case of extreme bets, the suggestion would be that one does know, e.g., where one was born, and it is appropriate to rely on that proposition in reasoning to accept the bet. However, it's alleged, it seems inappropriate to rely on that proposition in accepting the bet since, in doing so, one exhibits an inappropriate epistemic character:

[S]ometimes one does what is right but in doing so manifests traits of character that will foreseeably get one into trouble in none too distant scenarios. For any good case in which one knows that *p*, there will be a range of subjectively similar possible bad cases where one doesn't know that *p* but for all one knows one does know that *p*. Suppose that the bad cases are ones that could fairly easily come about in the future, and are ones where accepting a certain *p*-dependent contract leads to bad consequences. Someone who accepts the contract in the good case will obviously be so constituted as to accept the contract in the bad case. This fact about his constitution—this trait of epistemic character—may be worrisome, especially to the extent that the consequences of contract acceptance are disastrous, and especially to the extent that bad cases are expected to arise. (Hawthorne and Stanley forthcoming: 22–23).

While this response to the putative counterexamples to SUFF may seem initially plausible, we should consider it carefully. The response combines two ideas, namely that one does know the target proposition and that it seems incorrect that it is inappropriate to rely on that proposition in practical reasoning since, in so doing, one exhibits an epistemic flaw. In particular, one exhibits an epistemic trait that will lead to disaster in cases in which one thinks one knows but does not. I suggest that there is a tension between the two elements of the reply. For the character trait to be an epistemic flaw requires that, as Hawthorne and Stanley put it, the bad case be 'not too distant', or could 'fairly easily come about'. That a trait leads to disaster in a bad case is hardly a flaw if bad cases are remote from reality, for instance, if the trait is disastrous in a Cartesian demon world. But to the extent that bad cases are 'not too distant', this undermines the supposition that one does know the target proposition. On a widespread view, a true belief that *p* is

knowledge only if safe, i.e. there is no nearby world in which one believes that *p* falsely. Thus, allowing the bad case to be nearby in order to see the character trait as an epistemic flaw undermines the supposition that one knows.

There are a variety of ways of filling out the notion of ‘nearness’. To illustrate, suppose that the measure of nearness is probabilistic. On this view, knowing the target proposition requires that bad cases not be too likely. For instance, to the extent that one’s frequently wrong about such ordinary matters as one’s name or one’s place of birth, one does not know one’s name or place of birth. So, to be consistent with the supposition that one does know the target proposition, cases in which one thinks one knows but does not must be rare. This makes it hard to fill out the idea that in relying on what one knows in taking the bet one exhibits an epistemic vice. For now we see that the relevant trait leads to difficulty only rarely. This is especially problematic for the case of bets with heavily weighted odds but without high stakes, such as the bet that pays a million millionth of a penny if one was born in England but £3 if not. In the rare cases in which one loses such a bet, one loses very little—just £3. It’s hard to see why an epistemic trait that rarely leads one to lose such small amounts as £3 could be much of an epistemic vice. In the case of bets with high stakes, the tendency to take such bets rarely leads in trouble but, when it does, the costs are severe, e.g. one may lose one’s house. Even here, it’s far from clear that a trait which occasionally leads to disaster is an epistemic vice. After all, in acting one has to take some risks, even risks of disastrous things occurring, and it’s often rational to do so. Every time one crosses a road one exhibits a tendency which leads to disaster in a few rare cases, namely the tendency to cross the road when it seems to be clear of traffic. That there are a few cases in which one thinks the road’s clear but it’s not and one ends up seriously injured or dead does not show that there is an epistemic vice in crossing roads.

I have examined two strategies for a subject sensitive invariantist to respond to putative counterexamples to sufficiency—the claim that the subject does not know the target proposition, or the attempt to explain away the intuitive sense that the subject should not rely on the target proposition in practical reasoning. Both fail, at least in the dialectical context in which SUFF functions as a premise in an argument for SSI.

8. Practical Reasoning and the Case for SSI

I have been examining whether KN can be used as a premise in an argument for SSI. In this dialectical context, it is illegitimate to defend KN by appeal to SSI. I have argued that, without such an illegitimate appeal, SUFF is vulnerable to counterexamples. Further, I have argued that Gettier cases undermine NEC. I have suggested an alternative view on which the epistemic standards for practical reasoning vary with context. How does this discussion of KN affect the case for SSI?

Stanley (2005) defends SSI by arguing that SSI can avoid certain problems facing contextualism while explaining much of the original contextualist data (2–3). For instance, since SSI is an invariantist view, it can avoid the objections to the idea that ‘know’ is a context-sensitive term (Stanley 2005, chs. 2 and 3). However, it can explain the data in at least first-person contextualist cases. For example, the bank case is constructed so that it seems appropriate for DeRose to self-ascribe knowledge that the bank is open on Saturday in Low but not High, although Low and High differ only in the stakes. Since SSI holds that whether one knows depends on the stakes, it can treat these intuitions at face value: in Low but not High, it seems appropriate for DeRose to self-ascribe knowledge since, in Low but not High, DeRose knows the relevant proposition.

However, there are well-known difficulties facing this strategy for defending SSI. There are a range of other possible invariantist explanations of the difference in the intuitive acceptability of DeRose’s knowledge claim. Like SSI, these views avoid the problems facing contextualism. But, on these alternative views, the truth value of the knowledge claim is constant across Low and High. For instance, on certain pragmatic approaches, the knowledge ascription has the same truth value in Low and High but is conversationally appropriate only in one context (Rysiew 2001, Schaffer 2004, Brown 2005). On a variety of error theories, the knowledge ascription has the same truth value in Low and High, but one of our intuitive judgements about its propriety is mistaken (Hawthorne 2004, Williamson 2005). There is an ongoing debate about the plausibility of these explanations of the data as opposed to the explanation offered by SSI (e.g., see Rysiew 2001, DeRose 2002, Brown 2005).

A defender of SSI could avoid these problems by combining contextualist cases with KN. Low and High differ not only in the intuitive acceptability of knowledge-ascriptions but also in the intuitive acceptability of practical reasoning. Recall that, in the case, DeRose truly believes that the bank is open on Saturday on the basis of his visit to the bank two weeks ago. If stakes are low, then it seems appropriate for DeRose to rely on the claim that the bank is open on Saturday in his practical reasoning, say in reasoning that since the bank is open on Saturday, he will go then rather than wait in the long Friday queue. But if stakes are high—there’ll be big trouble with the bank if the checks are not deposited before Monday—then it no longer seems appropriate for DeRose to rely on that claim in his practical reasoning. He should not reason that since the bank is open on Saturday, he should go then rather than wait in the Friday queue. Given the stakes, he should go in and check whether the bank is open on Saturday.

A defender of SSI could combine this difference in the intuitive acceptability of practical reasoning with KN to defend their interpretation of contextualist cases over rival invariantist explanations. For instance, Stanley appeals to the idea that knowledge is the norm of practical reasoning to argue against

an explanation of contextualist data in terms of psychological bias caused by framing effects:

However, the above cases reveal intuitions that are not analogous to the framing effects we see in ordinary speakers' judgments about rationality. The latter sort of judgement does not follow a discernible pattern that reflects any plausible general claim about rationality. In contrast, the intuitions we have in the above cases are just the intuitions we would expect to have, if certain antecedently plausible conceptual connections between knowledge and practical reasoning were true. (p.9)

Appeal to KN can provide a general argument against any non-truth value explanation of the contextualist data, whether pragmatic or an error theory. Consider the following argument, letting O be the proposition that the bank is open on Saturday:

- 1) KN: it is appropriate to rely on p in practical reasoning iff you know that p (by assumption).
- 2) In Low but not High, it is appropriate for DeRose to rely on O in practical reasoning.
- 3) In Low but not High, DeRose knows that O (from 1) and 2)).
- 4) Low and High differ only in the stakes (by stipulation).
- 5) Whether S knows that p partly depends on the stakes.

If knowledge is not the norm of practical reasoning, then this general strategy of arguing against any non-truth value explanation of the contextualist data is not available. This leaves it open that a defender of SSI could show on a case by case basis that any proposed non-truth value explanation of the data is problematic. For instance, there is debate about the conditions for appropriate pragmatic explanation and whether these are met in contextualist cases (see Rysiew 2001, DeRose 2002, Stanley 2005, Brown 2005). But without KN, the general argumentative strategy against non-truth value interpretations is not available.

Stanley appeals to KN at other points in his argument including, crucially, his response to third person contextualist cases. In such cases, an attributor in a high stakes scenario discusses whether a subject in a low stakes scenario has knowledge. Intuitively, it seems inappropriate for the attributor to claim that the subject knows. It is difficult for SSI to accommodate the case for, on that view, whether the subject knows depends on the stakes for the subject, not the stakes for some third party. Stanley appeals to the idea that knowledge is the norm of practical reasoning to discount intuitions in third-person contextualist cases:

it would be a mistake to allow idiosyncratic facts about my own practical situation to impinge upon my judgements about the practical rationality of someone

in a quite different situation... given the intuitive connections between knowledge and practical reasoning, one should therefore be deeply suspicious of a theory that gives undue weight to our intuitions about High Attributor-Low Subject Stakes. (98)

Stanley's idea seems to be that whether it is appropriate for a subject to rely on a claim in practical reasoning depends on the stakes for her. If knowledge is the standard of practical reasoning, then knowledge also depends on the stakes for the subject and not, as the contextualist would have it, the stakes for the attributor. If knowledge is not the norm of practical reasoning, then Stanley needs some other reason to discount intuitions in third person cases.

The problems for KN also affect Hawthorne's defence of SSI. Hawthorne (2004) presents a number of criteria for an account of knowledge, one of the most important of which is the claim that knowledge is the norm of assertion and practical reasoning (21-31). Hawthorne admits that which account of knowledge the criteria favour is a finely balanced matter, however, he suggests that SSI has the edge over contextualism and classic invariantism:

Put a gun to my head and I will opt for a treatment of the puzzles built around the materials of the 'Practical Environment' section above. But I am far from confident that this is the correct way to proceed. Then there is the further question of whether to embed those ideas within an invariantist semantical framework for 'know'. Here, though more tentatively still, I would opt for invariantism over contextualism. (188)

Appeal to KN is crucial to Hawthorne's suggestion that SSI has the edge over contextualism and classic invariantism. For instance, without appeal to the idea that knowledge is the norm for assertion and practical reasoning, non-sceptical classic invariantism and SSI perform similarly on Hawthorne's criteria (149, 185-86). Both respect the Moorean constraint and Single Premise Closure. Neither can honour both the Epistemic Possibility Constraint and the Objective Chance Principle. Whereas non-sceptical invariantism is 'in obvious tension' with Multiple Premise Closure, SSI has 'some prospect' of maintaining Multiple-Premise Closure. By contrast, Hawthorne claims, SSI is better able to accommodate the idea that knowledge is the norm for assertion and practical reasoning. We have seen that one half of KN, NEC, is undermined by Gettier cases. Further, it is hard to defend the other half of KN, SUFF, without prior appeal to SSI. As a result, it is difficult to use SUFF as an independent criterion by reference to which to determine the correct account of knowledge. The problems encountered in defending KN may well alter the fine balance Hawthorne detects in favour of SSI.

Despite the difficulties in using SUFF as a premise in an argument for SSI, SUFF and SSI may stand in a different dialectical relationship. For instance, someone could offer the combination of SUFF and SSI as the best

explanation of the data in contextualist cases rather than supposing that one is more fundamental than the other¹³ (if, as I suggest, Gettier cases are counterexamples to NEC, the combination of NEC and SSI cannot explain the data). A full examination of this suggestion is not feasible within the confines of this paper. But it may be useful to show how a rival invariantist view may explain the data. Assume a non-sceptical classic invariantist view on which the stakes do not affect whether one knows. On this view, in both contexts of the bank case, DeRose knows that O (the bank is open on Saturday). On the position defended here, the epistemic standard for relying on p in practical reasoning varies with context. This can be used to explain why, although DeRose knows that O in both Low and High, it is appropriate in Low but not High, for him to rely on O in practical reasoning. In Low, given that the stakes are low, knowledge is enough to rely on O in practical reasoning. However, in High, given that the stakes are high, DeRose needs a stronger epistemic position than knowledge in order to rely on O in practical reasoning. The notion of strength of epistemic position can be filled out in a variety of ways, e.g. the range of alternatives one can rule out, or the range of worlds across which one's belief tracks the truth. For illustration, we may use the notion of the range of alternatives one can rule out. On this view, knowledge that O does not require that one can rule out a change in hours but, given the stakes in High, DeRose should rely on O in practical reasoning only if he can rule out a change in hours. Thus, in High, although he knows that O, it is inappropriate for him to rely on O in practical reasoning.

The invariantist may offer a pragmatic explanation of the data concerning assertion and knowledge attributions. The conversation between DeRose and his wife in High concerns the practical question of whether they should wait in the queue to deposit their paycheques or wait until the next day, a Saturday. Further, in High, DeRose's wife reminds DeRose how important it is that the cheques are paid in before Monday. Given the stakes, it seems that what is relevant to the conversation in High is a strong epistemic position, one strong enough to rule out even unlikely errors, such as the possibility that the bank has changed its hours. Thus were DeRose to reply to his wife by asserting 'I know that the bank is open on Saturday', or by making the simpler claim, 'The bank is open on Saturday', he would implicate that he is in such a strong epistemic position. For if he were not in such a strong epistemic position, his assertion would not be relevant to the on-going conversation. Since, by hypothesis, DeRose is not in such a position, the simple assertion and the knowledge claim are both unwarranted even though true. (For a fuller exposition and defence of this view, see Brown 2005.)

9. Conclusion

We have been examining the idea that knowledge is the norm of practical reasoning and whether this idea can be used as a premise in an argument

for SSI. I argued that some of the data which are often cited as supporting KN—the way in which we criticise and defend actions—instead support the view that the standards for practical reasoning vary contextually. Sometimes less than knowledge is required; sometimes knowledge is enough; sometimes more than knowledge is required. We saw that Gettier cases tell against NEC (if it is appropriate to rely on *p* in practical reasoning, then one knows that *p*). Further, SUFF is undermined by an intuitive part of the concept of knowledge, the idea that of two knowers, one may be in a better epistemic position than the other. If there can be two subjects who both know that *p*, but one's epistemic position is better than the other's, why should we think that mere knowledge is always sufficient for practical reasoning? This suggestion is supported by examples in which, intuitively, an agent knows some target proposition but is not in a good enough epistemic position to rely on that proposition in practical reasoning. Last, I sketched out how the idea of context-variable standards for practical reasoning could be combined with a pragmatic approach to explain the key data in contextualist cases.¹⁴

Notes

¹ See Hawthorne 2004, Stanley 2005, Fantl and McGrath forthcoming.

² Stanley (2005) holds that one should act on *p* only if one knows that *p* (9). Fantl and McGrath (forthcoming) argue that *S* knows that *p* only if *S* is rational to act as if *p*. Hawthorne claims 'if *p* is practically relevant, it is acceptable to use the premise that *p* in one's practical reasoning if one knows that *p* and (at least in many cases) unacceptable to use the premise that *p* in one's practical reasoning if one doesn't know it' (2004: 30). Hawthorne and Stanley (forthcoming) claim that 'where's one's choice is *p*-dependent, it is appropriate to treat *p* as a reason for acting iff you know that *p*'.

³ Notice that the notion of 'relying on *p* in practical reasoning' used by Hawthorne and Stanley is rather different from the notion of acting for a reason used by Hyman. For Hawthorne and Stanley, one can rely on *p* in practical reasoning whether or not one knows that *p*; however, doing so is appropriate iff one knows that *p*. Hyman (1999) defends the thesis that the fact that *p* can be *A*'s reason for acting iff *A* knows that *p*. Thus, for Hyman, when *p* is *A*'s reason for acting, her so reasoning cannot be assessed normatively according to whether *A* knows that *p*. Rather, if she does not know that *p*, then the fact that *p* cannot be her reason for acting. It is difficult, then, to use Hyman's analysis of the logic of reason giving statements to defend the thesis of Hawthorne and Stanley.

⁴ There are a variety of formulations of the left hand side of the norm, including 'it is acceptable to use the premise that *p* in one's practical reasoning' (Hawthorne 2004: 30), 'it is appropriate to treat *p* as a reason for acting' (Hawthorne and Stanley forthcoming), 'it is appropriate to act on *p*' (Stanley 2005: 9), 'it is rational to act as if *p*' (Fantl and McGrath forthcoming). I have focussed on the notion of its being appropriate to rely on *p* in one's practical reasoning as this seems a reasonably intuitive notion.

⁵ The norm is formulated in terms of the notion of its being 'appropriate' to rely to *p* in practical reasoning, rather than the notion that one ought to so rely. As Hawthorne and Stanley (forthcoming) point out, 'it would be overly demanding to require someone to treat all of their relevant knowledge as reasons for each action undertaken'.

⁶ If *p* is completely irrelevant to the issue at hand, it seems odd to say that if one knows that *p*, it is appropriate to treat *p* as a reason for acting. Thus, Hawthorne and Stanley (forthcoming) suggest restricting the knowledge norm to 'p-dependent choices' where a choice between options

x_1 - x_n is p -dependent iff the most preferable of x_1 - x_n conditional on p is not the same as the most preferable of x_1 - x_n conditional on not- p ' (9).

⁷ This case is also explicable on the assumption of NEC (it is appropriate to rely on p in practical reasoning only if one knows that p). The use of knowledge in defending my action could be understood in the following way: I defend my acting on p by pointing out that I meet the necessary condition for this to be appropriate, namely that I know that p .

⁸ E.g., that I defend my remaining in the office by citing facts about the train schedule (there's an express at 12.20pm) is compatible with KN if knowledge is also the norm of assertion (Stanley 2005: 10). That I defend my remaining in the office by citing evidence about the train schedule (the timetable says there's an express at 12.20pm) could be explained compatibly with KN by exploiting the idea that knowledge is the norm of assertion, and holding that the reasoning in question starts from the proposition that the timetable says there's a train at 12.20pm. One might reason: the timetable says there's a train at 12.20pm, the timetable is very likely true so, given the stakes, I'll leave the office in time for the 12.20pm.

⁹ One might try to motivate the view that, of two knowers, one's epistemic position may be stronger than the other's, by appeal to the idea that, on pain of scepticism, knowledge does not require the highest epistemic standards, e.g. Cartesian certainty. Rather, knowledge requires that one meet some threshold of strength of epistemic position, less than the maximum level. However, appealing to the problem of philosophical scepticism to establish this is controversial. For some argue that one can meet available interpretations of the Cartesian standard. For instance, on a factive notion of evidence, one can rule out the BIV possibility, and distinguish it from actuality.

¹⁰ See, e.g., Stanley 2005: 88–89.

¹¹ Stanley distinguishes a second sense of 'epistemic' in which it denotes whatever the factors are which make a true belief knowledge; in this sense of epistemic, the stakes are an epistemic factor.

¹² Of course, the agent cannot distinguish her actual situation in which she in fact knows the relevant proposition from a possible BAD situation in which she falsely believes that she does. As we have already seen, on pain of scepticism, the subject-sensitive invariantist should not accept that this inability prevents the agent from knowing in the actual situation. The same constraints on knowledge should apply at both first and second order level: if a subject can know that p despite her inability to distinguish her actual situation from a possible situation in which she mistakenly believes that p , then she can know that she knows that p despite her inability to distinguish her actual situation from a possible situation in which she mistakenly believes that she knows that p .

¹³ It may be that this kind of strategy is at work in Stanley (2005) although his explicit focus is on NEC not SUFF (see, e.g., 11–12).

¹⁴ Thank you to Peter Baumann, Sandford Goldberg Martin Smith, members of the philosophy departments at St. Andrews and Cardiff, and participants at the 2007 Bled epistemology conference.

References

- Bach, Kent. 2005. 'The Emperor's New 'Knows'.' In *Contextualism in Philosophy: On Epistemology, Language and Truth*, Gerhard Preyer and Georg Peter (eds.), Oxford University Press.
- Brown, Jessica. 2005. 'Adapt or Die.' *Philosophical Quarterly Special Edition on Contextualism*.
- Cohen, Stewart. 1988. 'How to be a Fallibilist.' *Philosophical Perspectives* 2: 91–123.
- . 1999. 'Contextualism, Scepticism, and the Structure of Reasons.' *Philosophical Perspectives* 13: 57–89.
- . 2000. 'Contextualism and Scepticism.' *Philosophical Issues* 10: 94–107.

- DeRose, Keith. 1992. 'Contextualism and Knowledge Attributions.' *Philosophy and Phenomenological Research* 52:913–29.
- . 1995. 'Solving the Sceptical Problem.' *Philosophical Review* 104: 1–52. Page references from K. DeRose and T. Warfield, eds., *Scepticism*, OUP: Oxford, 183–219.
- . 2002. 'Assertion, Knowledge and Context.' *Philosophical Review*.
- . 2004. 'The Problem with Subject-Sensitive Invariantism.' *Philosophy and Phenomenological Research*, LXVIII, 346–50.
- Fantl, Jeremy and McGrath, Matthew. 2002. 'Evidence, Pragmatics and Justification', *Phil. Review* 111, 1:67–94.
- . Forthcoming. 'Knowledge and the Purely Epistemic: in Favour of Pragmatic Encroachment'.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. OUP: Oxford.
- Hawthorne, John, and Stanley, Jason. Forthcoming. 'Knowledge and Action.' *Journal of Philosophy*.
- Hyman, John. 1999. 'How Knowledge Works'. *Philosophical Quarterly* 49:433–451.
- Rysiew, Patrick. 2001. 'The Context-Sensitivity of Knowledge Attributions.' *Nous* 35:477–514.
- Schaffer, Jonathan. 2004. 'Scepticism, Contextualism and Discrimination.' *Philosophy and Phenomenological Research* 69:138–55
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. OUP: Oxford.
- . 2007. 'Précis of Knowledge and Practical Interests and Replies to Critics.' *Philosophy and Phenomenological Research* 75:168–72.
- Weatherson, Brian. 2003. 'Thoughts, Arguments and Rants: the Bank Cases.' <http://www.brown.edu/Philosophy/tar/Archives/002074.html>.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: OUP
- . 2005. 'Contextualism, Subject-Sensitivity Invariantism, and Knowledge of Knowledge,' *Philosophical Quarterly Special Edition on Contextualism*.
- . 2005a. 'Knowledge, Context and Agent's Point of View.' In *Contextualism in Philosophy: On Epistemology, Language and Truth*, Gerhard Preyer and Georg Peter (eds.), Oxford University Press.
- . Forthcoming. 'Knowledge and Scepticism.' In F. Jackson and M. Smith, *Oxford Handbook of Analytic Philosophy*.