Episteme, 9, 1 (2012) 43–62 © Cambridge University Press
doi:10.1017/epi.2011.4

PRACTICAL REASONING, DECISION THEORY AND ANTI-INTELLECTUALISM

JESSICA BROWN
jab30@st-andrews.ac.uk

ABSTRACT

In this paper, I focus on the most important form of argument for anti-intellectualism, one that exploits alleged connections between knowledge and practical reasoning. I first focus on a form of this argument which exploits a universal principle, *Sufficiency*, connecting knowledge and practical reasoning. In the face of putative counterexamples to *Sufficiency*, a number of authors have attempted to reformulate the argument with a weaker principle. However, I argue that the weaker principles suggested are also problematic. I conclude that, so far, there is no good argument for anti-intellectualism that rests on connections between knowledge and practical reasoning.

1. DEFENDING ANTI-INTELLECTUALISM

Anti-intellectualism holds that whether S's true belief that *p* is knowledge depends not only on truth-conducive factors such as S's evidence for *p*, and whether S's belief that *p* was formed reliably, but also on the stakes for S, or the odds[1] (e.g. Fantl and McGrath 2002, 2009; Hawthorne 2004; Stanley 2005). Anti-intellectualism is sometimes defended by suggesting that it offers the best explanation of our intuitions or judgements about knowledge in cases like the bank cases (e.g. Stanley 2005). This first line of argument has been challenged both in the description it gives of our intuitions/judgements[2] and its claim that anti-intellectualism is the best explanation of these data. For these reasons, it is no surprise that several anti-intellectualists are explicit that they prefer to defend their view by a second argument which appeals to alleged connections between knowledge and practical reasoning (for instance, Fantl and McGrath 2012; Weatherson forthcoming). As we will see, this second argument does not rely on the claim that, intuitively, there is a difference of knowledge across the low and high contexts of the bank cases. Rather, it argues for such a difference in knowledge by appeal to principles linking knowledge and practical reasoning. (Of course, the argument may appeal to intuitions at other points, for instance in the defence of the knowledge norm for practical reasoning.)

1 I am sympathetic to Weatherson's view that anti-intellectualism is best formulated as the view that odds, rather than stakes, affect knowledge (Weatherson forthcoming). However, in the main text I follow the more standard formulation of the view.

2 For relevant experimental literature on this issue, see May et al. 2009; Buckwalter 2010; Feltz and Zarpentine 2010; Pinillos 2012; Knobe and Schaffer forthcoming.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

Furthermore, it does not take the form of an argument that anti-intellectualism offers a better explanation than intellectualist rivals of our intuitions/judgements about knowledge.

The second, practical reasoning argument will be our focus here. I start with one main form of this argument which employs a universal principle, *Sufficiency*, according to which knowledge is sufficient for being in a good enough epistemic position to rely on a proposition in one's practical reasoning. I then turn to discuss variants of the argument that employ weaker principles, including recent variants offered by Fantl and McGrath and Weatherson.

2. ANTI-INTELLECTUALISM AND THE KNOWLEDGE NORM FOR PRACTICAL REASONING

An initial formulation of the sufficiency direction of the knowledge norm for practical reasoning is as follows:

> *Sufficiency*: If you know that *p* then you are in a good enough epistemic position to rely on *p* in practical reasoning.

Of course, even if you are in a good enough epistemic position to rely on a proposition in your practical reasoning, it does not follow that it is appropriate in other senses for you to do so. For example, perhaps one comes to know some truth in a dishonest way. Since one knows this truth, then one is in a good enough epistemic position to rely on it in one's practical reasoning. But in virtue of the dishonest way in which one acquired this knowledge, it may be morally wrong to exploit it in one's practical reasoning.

To see how *Sufficiency* may be used to motivate anti-intellectualism, let us consider a first-person version of DeRose's bank case. In this case, it is stipulated that in both the low and high contexts, DeRose truly believes that the bank is open on Saturday on the basis of the same evidence, namely his recent visit to the bank. Given the setup, it is plausible that, in the low context, DeRose knows that the bank is open on Saturday. Further, given that in the low context nothing much turns on whether the bank is open on Saturday, it seems appropriate for DeRose to rely on this claim in his practical reasoning, say his reasoning to the conclusion that instead of waiting in the long Friday queue, he will go to the bank on Saturday instead. However, in the high context, it is of great importance to DeRose whether the bank is open on Saturday: he stands to default on his mortgage payments and risk his house if he does not deposit money in his bank account before the end of Saturday. Given the stakes in the high context, it seems that he is not in a strong enough epistemic position to rely on the claim that the bank is open on Saturday in his practical reasoning. Instead, before deciding not to wait in the long Friday queue, he should check the bank's opening hours. If *Sufficiency* were true then it could be combined with the claim that, in high, DeRose is not in a strong enough epistemic position to rely on the relevant proposition in his practical reasoning to conclude that, in high, DeRose does not know that proposition. Since DeRose does plausibly know in the low context, one may conclude that knowledge is a function of the stakes. An argument of this form is explicitly given in Fantl and McGrath's recent defence of anti-intellectualism (Fantl and McGrath 2009).

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

In order for an anti-intellectualist to use *Sufficiency* to defend their position, they need to motivate *Sufficiency*. Further, if *Sufficiency* is to provide an argument for anti-intellectualism from independently motivated premises, the defence of *Sufficiency* itself should not rest on anti-intellectualism. How, then, could one support *Sufficiency* without appeal to anti-intellectualism?

One way to defend *Sufficiency* is by appeal to intuitions about cases. For instance, Hawthorne and Stanley (2008) defend a bi-conditional version of the knowledge norm for practical reasoning by appeal to intuitions about the propriety of practical reasoning and action. For instance, they point out that we criticise and defend actions by appeal to knowledge, saying such things as 'You shouldn't have gone left. You didn't know that the restaurant was to the left'; and 'Why didn't you take the 10 o'clock bus home? You knew that it was the last bus.'

This intuitive defence of the knowledge norm is problematic for several reasons. First, we offer a variety of different criticisms and defences of actions and practical reasoning. In particular, we criticise actions by citing epistemic states stronger than knowledge, e.g. 'You shouldn't have left the party so late. You didn't know for certain that there would be another bus after the 10 o'clock.' (Less problematically, we also criticise actions by citing epistemic states weaker than knowledge, for example 'You shouldn't have gone left. You had no evidence the restaurant was to the left.') Second, there are intuitive counterexamples to both directions of the bi-conditional (e.g. Hill and Schechter 2007; Brown 2008; Fantl and McGrath 2009; Littlejohn 2009; Neta 2009; Lackey 2011; Reed 2011). Some attempt to accommodate the full range of data concerning criticism and defence of action by supposing that there is a single standard for relying on a proposition in practical reasoning, where the standard is not knowledge (e.g. Fantl and McGrath 2009; Littlejohn 2009; Neta 2009). By contrast, my own preferred view is that the full range of data is best explained by supposing that the epistemic position required to rely on a proposition in practical reasoning varies with context (Brown 2008; see also Gerken 2011). One way to formulate this idea presupposes an invariantist threshold conception of knowledge on which knowledge requires a certain invariant strength of epistemic position, *s*. With this background, we may formulate the idea by saying that sometimes the strength of epistemic position required to rely on a proposition in practical reasoning is less than *s*, sometimes it is *s*, and sometimes it is more than *s*. We may flesh out this idea on a variety of different readings of strength of epistemic position, including the range of worlds across which one's belief tracks the truth, one's degree of justification, or the degree of probability of the relevant proposition on one's evidence. I will draw on this view in what follows.

Given our focus on the sufficiency direction of the knowledge norm, it will be useful to describe a putative counterexample to the sufficiency direction. For instance, in Brown's surgeon case, a student is shadowing a surgeon for the day (Brown 2008). The student observes the surgeon examining patients in the morning including one with a diseased left kidney. The decision is taken to remove it in the afternoon. Later, the student observes the surgeon in theatre where the same patient is lying anaesthetised on the operating table. The student is puzzled that the operation hasn't started and that the surgeon is consulting the patient's notes.

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?

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

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.

The surgeon example provides a case in which, intuitively, a subject knows a proposition yet is not in a good enough epistemic position to rely on this proposition in her practical reasoning but should rather check first. As a result, it makes it difficult to defend *Sufficiency* merely by appeal to intuitions about cases. Recognizing this, Fantl and McGrath accept that *Sufficiency* is 'not clearly true, prior to theorising, and that the evidence from habits of criticisms and defences [of action] . . . is not decisive' (2009: 63). Accepting that an intuitive defence of *Sufficiency* is inadequate, they attempt to provide what they describe as a 'principled' argument for *Sufficiency* from premises concerning the relation between knowledge, reasons and justification (ibid.). I will focus on their principled argument since it seems the most promising way of defending *Sufficiency*. The key question for us will be whether this principled argument avoids the problems raised by Brown's surgeon case. I will argue that it does not do so. Anyone who finds the surgeon case problematic for the intuitive defence of *Sufficiency* should find it equally problematic for the principled defence.

3. FANTL AND MCGRATH'S THEORETICAL DEFENCE OF *SUFFICIENCY*

In assessing Fantl and McGrath's principled argument for *Sufficiency*, I will use their preferred formulation of *Sufficiency*, namely (KJ):

> (KJ): If you know that p, then p is warranted enough to justify you in ϕ-ing, for any ϕ.

They explain that what it means to say that 'p is warranted enough to justify you in ϕ-ing' is that no weakness in your epistemic position with respect to p stands in the way of p justifying you in ϕ-ing (2009: 64).

Fantl and McGrath's argument for (KJ) involves a number of further principles. We will focus on two of them, the Unity Thesis and the Theoretical Knowledge-Reasons Principle. According to the Theoretical Knowledge-Reasons Principle, if you know a proposition then it is warranted enough to be a reason for belief:

> Theoretical Knowledge-Reasons Principle (Theoretical KR): If you know that p, then p[3] is warranted enough to be a reason you have to believe that q, for any q. (2009: 69)

3 Here, Fantl and McGrath assume that p itself is one's reason for belief, rather than one's belief that p (2009: 70). Although this is controversial, I won't discuss it here. For the rest of the paper, I will adopt their approach. However, those who think that the belief that p is one's reason, rather than the proposition p itself, should reformulate the discussion accordingly.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

Their argument involves combining the Theoretical Knowledge-Reasons Principle with the Unity Thesis, according to which if a proposition is warranted enough to be a reason for belief then it is warranted enough to be a reason for action:

> The Unity Thesis: If *p* is warranted enough to be a reason you have to believe that *q*, for any *q*, then *p* is warranted enough to be reason you have to ø, for any ø. (2009: 73)

Combining the Unity Thesis with (Theoretical KR) yields the claim they label (KR):

> (KR) If you know that *p*, then *p* is warranted enough to be a reason you have to ø, for any ø.

It is then a short step from this principle to their preferred formulation of *Sufficiency*, namely (KJ): if you know that *p*, then *p* is warranted enough to *justify* you in ø-ing, for any ø.

Fantl and McGrath don't do much to defend (Theoretical KR), describing it as 'relatively uncontroversial' (2009: 69). However, they do provide some defence of the Unity Thesis. They note that if it were false we would expect that we would segregate reasons by whether they are available for drawing practical or theoretical conclusions. But they argue that we do not see segregation even in high stakes cases: 'even when the stakes are high we do not find segregation; rather, when the stakes are high, we are more careful about using a consideration as a reason to believe – as careful as we are about using it as a reason to act' (2009: 73). They defend this claim by reference to an example in which one is considering whether the ice is thick enough to walk across a frozen lake. They claim that in deciding whether to walk across the lake, one tries to carefully figure out whether the ice is thick enough. Having settled this issue, one employs that conclusion in one's practical reasoning. One does not instead 'segregate', refusing to allow one's theoretical conclusions to influence one's practical reasoning (2009: ch. 3, section 3.2). Later, I will suggest an alternative explanation for the lack of segregation in this case, one which does not rest on the Unity Thesis. But, for now, let us consider the strength of their new argument for *Sufficiency*.

I will argue that Fantl and McGrath's new argument for (KR) should not impress anyone who found the surgeon case initially plausible. I will illustrate this by using a variant surgeon case in which it is the left lung, rather than the left kidney, which is diseased and so needs removing. The aim behind this slight change is to find an example in which removing an organ from the left side is a more time-consuming operation than removing the matching organ on the right side. I assume that removing a left lung is more time-consuming because of the proximity of the heart. If this is not the case, the reader should substitute an alternative such case. With this one exception, everything is as before. The student shadows the surgeon and watches her examining the sick patient in the morning where it is decided to remove the left lung in the afternoon. Later, the student is in the operating theatre and is puzzled about why the surgeon hasn't started operating yet. She asks the nurse 'Why is the surgeon checking the notes, doesn't she even know which lung it is?' The nurse replies 'Of course she knows which lung it is. But, imagine what it would be like if she removed the wrong lung. She shouldn't operate before checking the patient's records.' Here, we again have a case in which intuitively a person knows a

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

proposition but is not in a strong enough epistemic position to rely on that proposition in some practical reasoning. Despite this, it seems that the surgeon might well be in a good enough epistemic position to rely on the same proposition in other reasoning. For instance, it would seem fine for the surgeon to reason that, since it is the left lung which needs taking out, and left lungs are more time-consuming to remove than right lungs, she won't be home in time to pick up the children from school. Alternatively, she would be fine to reason that, since she removed one left lung earlier in the week, she has had two left lung procedures this week. The variant surgeon case provides a case in which a subject's epistemic position with respect to a single proposition is warranted enough to be a reason for (some) beliefs, say that she won't be home in time, but not warranted enough to be a reason to start operating immediately, without checking the records. Thus, the very case which undermined an intuitive defence of *Sufficiency* seems to undermine Fantl and McGrath's alternative principled defence of *Sufficiency*, by undermining the Unity Thesis.

Fantl and McGrath may reply that the case does not undermine the Unity Thesis by showing that a single proposition may be warranted enough to be a reason for belief but not warranted enough to be a reason for action. For they may claim that the proposition that is used to form beliefs is not the proposition that it is the left lung, but only that *it is likely* to be the left lung. For instance, they may argue that the surgeon's reasoning about when she will get home is best understood as the reasoning that, since it is likely that it is the left lung which needs taking out, and left lungs are more time-consuming than right lungs, it is likely that she won't be home in time to pick up the kids from school. If this were correct, then the case would no longer provide a case in which a single proposition is being used as a reason to form beliefs, but not as a reason to operate straightaway without checking. However, this reply fails. It seems implausible to suppose that the surgeon only uses the more qualified proposition in forming beliefs. For instance, it seems implausible that her reasoning about the number of procedures carried out during the week is best represented as follows: since I removed one left lung earlier in the week, and this is likely to be a left lung, it is likely that I have had two left lung procedures this week. This train of reasoning seems far too qualified to represent the surgeon's reasoning. Rather, I suggest, she can and would reason: since I removed one left lung earlier in the week, and here's another left lung, I have had two left lung procedures this week.

Fantl and McGrath may argue that the variant surgeon case does not in fact counterexample the Unity Thesis. The antecedent of the Unity Thesis states that, for *any proposition q*, *p* is warranted enough to be a reason to believe *q*. In the variant surgeon case, we have a proposition that is warranted enough to be reason for believing some further propositions (that one will be home late, or that one has had two left lung procedures this week). That a proposition is warranted enough to be a reason for believing some proposition does not entail that, for any proposition *q*, it is warranted enough to be a reason to believe *q*.[4]

In reply, I will argue that, even if the variant surgeon case does not clearly present a counterexample to the Unity Thesis, it still undermines Fantl and McGrath's argument

4 Of course, a similar problem affects Fantl and McGrath's own ice case to which they appeal to defend the Unity Thesis. In the ice case, there is no attempt to show that if some proposition, *p*, is warranted enough to be a reason for believing *q*, for any *q*, then it is warranted enough to be a reason for *ø*-ing, for any *ø*. Rather, at best, the case establishes that there is a proposition which is warranted enough to be reason to believe some propositions and warranted enough to be a reason for undertaking some actions.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

for (KR). We can understand the case as presenting a dilemma for Fantl and McGrath, on either horn of which one of the two premises of their argument for (KR) is false. The case makes it plausible that the surgeon knows that it is a left lung, but she isn't in a good enough epistemic position to rely on that proposition in deciding to operate without checking. The dilemma arises when we ask the question, is the proposition that it is a left lung warranted enough to be a reason for the surgeon to believe that q, for any q? Suppose that Fantl and McGrath answer 'no' to this question. On that supposition, we have a case in which the surgeon knows a proposition (that it is the left lung) but that proposition isn't warranted enough to be reason to believe that q, for any q. In this case, the first premise of Fantl and McGrath's argument, namely (Theoretical KR) is false. Alternatively, Fantl and McGrath may answer 'yes' to the crucial question, accepting that it is a case in which the proposition that it is a left lung is warranted enough to be a reason to believe that q, for any q. In this case, since the surgeon isn't in a good enough epistemic position to operate without first checking, then this is a case in which Fantl and McGrath's second principle, the Unity Thesis, is false. Assuming that the case makes plausible that the surgeon knows that it is the left lung, yet is not in a good enough epistemic position to operate without first checking, then the case puts pressure on either the first or the second premise of Fantl and McGrath's argument.

An objector to Fantl and McGrath's argument need not decide whether the variant surgeon case is best understood as providing a counterexample to (Theoretical KR), or the Unity Thesis. However, my own view is that it is best understood as a counterexample to the Unity Thesis. Thus, it is useful to consider Fantl and McGrath's claim that rejecting the Unity Thesis is 'barmy' (2009: 74). Suppose, as is plausible, that after the morning clinic, the surgeon knows that it is the left lung. By (Theoretical KR), that it is the left lung is warranted enough to be a reason for belief, for any belief. Importantly, this includes beliefs about the consequences of various actions. Thus, the surgeon could combine her knowledge that it is the left lung with the following further pieces of knowledge: (1) that if it is the left lung, then removing the left without checking will be removing the correct lung; and (2) that if it is the left lung, then checking before removing the left lung is a waste of time. She may thus infer and come to know that removing the left lung without checking is the best thing to do. Fantl and McGrath claim that it would be barmy not to allow one's theoretical conclusions into one's practical reasoning. Thus, they would claim that if the surgeon knows that removing the left lung without checking is the best thing to do, she should implement this conclusion practically. But it is far from clear that refusing to do so is barmy. According to a widely shared conception of knowledge, fallibilism, knowledge is compatible with a chance of error. Thus, the surgeon's initial knowledge that it is the left lung is compatible with a chance of error. This chance of error is not removed by her subsequent reasoning. Thus, her knowledge that removing without checking is the best thing to do is also compatible with a chance of error. Removing the wrong lung would be disastrous for the patient. Thus, it isn't at all barmy not to implement the theoretical conclusion but instead operate after first checking (for a similar argument, see Reed 2011[5]). (Note that Fantl and McGrath (2009) endorse an anti-intellectualist version of fallibilism according to which knowledge is compatible

5 Reed usefully distinguishes knowing that ϕ-ing is the best thing to do given p, and knowing that ϕ-ing is the best thing to do in the sense of having the highest expected utility. In the second sense, given the possibility of error, it is not clear that the surgeon knows that operating without checking is the best

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

only with a chance of error which is not practically relevant. They might appeal to this conception to argue that, since the error is practically relevant, the surgeon does not know that it is the left lung. However, if they were to appeal to their anti-intellectualist conception of fallibilism to defend the Unity Thesis, then the Unity Thesis could not provide the kind of defence of anti-intellectualism on independent grounds for which we are looking.)

I have argued that Fantl and McGrath's theoretical argument for *Sufficiency* should not convince anyone who already found the surgeon style cases plausible. In particular, I have suggested that the surgeon case can be used to cast doubt on the Unity Thesis. However, in response to similar cases which I presented to them earlier,[6] Fantl and McGrath have suggested that their defence of *Sufficiency* could use a weaker principle than the Unity Thesis. In particular, they suggest that their argument would be successful even if the following weaker principle is true:

> Unity (weakened): When we are in the course of practical reasoning, *p* is warranted enough to be a reason for belief only if it is warranted enough to be a reason for action (2009: 75).

If this weaker principle were true, then Fantl and McGrath could combine it with (Theoretical KR) in their argument for anti-intellectualism. In particular they could use it to argue that, in the high context of the bank case, DeRose doesn't know that the bank is open on Saturday. In the case, DeRose is engaged in practical reasoning about whether to go to the bank on Saturday rather than Friday, so Unity (weakened) applies to DeRose. The proposition that the bank is open on Saturday isn't warranted enough to be a reason for him to go to the bank on Saturday rather than Friday. So, by the weakened Unity principle, that proposition is not warranted enough to be a reason for belief. Assuming (Theoretical KR), it follows that DeRose doesn't know that the bank is open on Saturday.

However, notice that even the weakened principle is threatened by the variant surgeon case. The surgeon is facing a decision about whether or not to check before operating, and so the weakened Unity Thesis applies in the variant surgeon case. As we saw earlier, it seems plausible that a single proposition, that it is the left lung, is warranted enough to be a reason for (some) beliefs, say that the surgeon has had two left lung procedures this week, but isn't warranted enough to be a reason for just going ahead and operating without checking. We can pose a dilemma to Fantl and McGrath as before, concerning whether, in the variant surgeon case, the proposition that it is the left lung is warranted enough to be a reason to believe *q*, for any *q*. If it is not then, given the plausible claim that the surgeon knows that it is the left lung, the case provides a counterexample to (Theoretical KR). If it is, then, given that the proposition that it is the left lung is not warranted enough to operate without checking, the case provides a counterexample to the weakened Unity Thesis.

Contra Fantl and McGrath, I have suggested that the Unity Thesis is false. If, as I'm suggesting, the Unity Thesis is false, then one might wonder how we can explain the data

---

thing to do. Rather, she infers that it is the best thing to do from the proposition that it is the left lung which is diseased, where her knowledge that it is the left lung is compatible with a chance of error.

6 See their discussion of my variant ice case (Fantl and McGrath 2009: 76).

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

that Fantl and McGrath use to motivate the Unity Thesis, namely that we don't find segregation in high stakes cases, such as the original ice case. Recall that, in that case, it seems plausible that, if in the course of deciding whether to cross the ice, I form the belief that it is thick enough, then I will take that as a reason for crossing the ice. I think we can explain the lack of segregation in their original ice case by appeal to the details of that case, and in particular the fact that the reasoning about the thickness of the ice is undertaken with the aim of resolving the practical decision to walk over the ice or around it. Within the attempt to solve some practical decision it would seem pointless to employ the proposition that *p* if one believes that one's epistemic position with respect to *p* is not strong enough to act on. For, then, by one's own lights, one's belief would not be strong enough to act on. For example, in the ice case, it would seem pointless to employ the proposition that the ice is thick enough to walk on within reasoning aimed at deciding whether to walk across the ice if, by one's own lights, one's epistemic position with respect to that proposition is not strong enough to cross the ice. As a result, I suggest we can explain the lack of segregation in the ice case without appealing to Unity Thesis, or even the weakened version of it. Recall that the weakened Unity Thesis states:

> Unity (weakened): When we are in the course of practical reasoning, *p* is warranted enough to be a reason for belief only if it is warranted enough to be a reason for action.

I suggest that we can explain the lack of segregation in the ice case even if this weakened Unity Thesis is false by the following alternative descriptive principle, understood as a psychological generalisation about us:

> Employ: Within[7] the attempt to solve some practical reasoning decision, if one employs a proposition *p* as a reason for belief, then it is not the case that one believes *p* not to be warranted enough to be a reason for the relevant action.

For instance if, within the attempt to decide whether to walk across the ice, one employs the proposition that the ice is thick enough as a reason for belief, then it is not the case that one believes that proposition not to be warranted enough to be a reason for walking across the ice. Thus, it is no surprise that one moves immediately from belief to action. Thus, even if the Unity Thesis is false, we can explain the data which led Fantl and McGrath to introduce it.

Notice that the explanation I've offered of a lack of segregation in the original ice case cannot be used by Fantl and McGrath in their argument for anti-intellectualism. It is

7 This phrasing is an attempt get around the problem for the weakened Unity Thesis that, as the variant surgeon case illustrates, one may simultaneously use a proposition as a reason for belief in one project (e.g. deciding whether to make alternative arrangements for picking up the kids) but not another (say, deciding whether to operate without first checking). In the intended sense, the surgeon doesn't use the proposition that it is the left lung within the attempt to decide whether to operate without checking, even if she does use it within the attempt to decide whether to make alternative arrangements for picking up the kids. Of course, Fantl and McGrath could attempt to use this phrasing to reformulate their weakened Unity Thesis so as to deal with the variant surgeon case. However, they would still need to defend their explanation of the lack of segregation in the ice case over my alternative explanation which does not lend itself to an argument for anti-intellectualism.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

useful here to compare Fantl and McGrath's original Unity Thesis with the different principle I have defended. Let us take it for granted that, in the high context of the bank case, the proposition that the bank is open on Saturday is not warranted enough for DeRose to delay his visit to the bank until Saturday. According to the Unity Thesis, if *p* is warranted enough to be a reason you have to believe that *q*, for any *q*, then *p* is warranted enough to be a reason you have to ø, for any ø. By Unity, since the relevant proposition is not warranted enough to be a reason DeRose has to go to the bank on Saturday rather than Friday, it is not warranted enough to be a reason he has to believe. Assuming (Theoretical KR), it follows that DeRose does not know that the bank is open on Saturday. Thus, the Unity Thesis smoothly combines with the claim that the proposition that the bank is open on Saturday is not warranted enough to be a reason for DeRose to delay his visit to the bank to generate the conclusion that DeRose does not know that proposition. Now consider my alternative principle, Employ, which states that within the attempt to solve some practical reasoning decision, if one employs a proposition *p* as a reason for belief, then it is not the case that one believes *p* not to be warranted enough to be a reason for the relevant action. This principle does not concern whether one's epistemic position is in fact strong enough to act on a proposition, but rather one's beliefs about this. Thus, it does not easily combine with the claim that DeRose is not in fact in a strong enough epistemic position to act on the proposition that the bank is open on Saturday. Assuming that DeRose believes that the proposition that the bank is open on Saturday is not warranted enough to be a reason to delay his bank visit, all that follows is he won't employ that proposition in his practical reasoning about whether to delay the bank visit.

In this section, I have considered Fantl and McGrath's theoretical defence of *Sufficiency*. However, I have argued that the surgeon example not only casts doubt on the intuitive defence of *Sufficiency* but also casts doubt on their theoretical defence of *Sufficiency*. Towards the end of the section, I considered but rejected the suggestion that anti-intellectualism could be defended by appeal to a principle weaker than the Unity Thesis.

4. WEATHERSON'S NEW ARGUMENT FOR ANTI-INTELLECTUALISM

Weatherson (2012) offers an argument for anti-intellectualism with a form strikingly similar to that offered by Fantl and McGrath. However, as we will see, Weatherson offers an interestingly different defence of one of the crucial premises of the argument. Weatherson's argument involves the following four steps:

1. Describe a subject such that it would ordinarily seem plausible to claim that he knows that *p*.
2. Offer the subject a bet which it is not rationally permissible for him to take.
3. Argue that if the subject still knows that *p* after the bet is offered, then it is rationally permissible for him to take the bet.
4. On pain of scepticism, conclude that being offered the bet undermines knowledge he previously had.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

In his example, *p* is the proposition that Barry's favourite musician, Beth, is playing in the nearby park. Barry believes this proposition on the basis of the music floating up from the park to his apartment. Ordinarily, we take it that people can know such propositions on this kind of basis. The bet offered to Barry pays out $10 if the musician is Beth but 1,000 years of torture otherwise; declining the bet leaves him unaffected. In overall structure, Weatherson's argument is very similar to Fantl and McGrath's argument for anti-intellectualism in which Keith is initially described in such a way that it would ordinarily seem plausible to claim that he knows that the bank is open on Saturday and is then offered a bet which he should not take (to bet on the bank being open on Saturday at the risk of losing his house). However, whereas Fantl and McGrath support an analogue of 3 by appeal to *Sufficiency*, Weatherson supports 3 without committing himself to *Sufficiency*. He explicitly claims that his argument relies on weaker premises than those standardly used in supporting anti-intellectualism and, in particular, requires neither the necessity nor the sufficiency direction of the knowledge norm for practical reasoning (1.2). In effect, Weatherson supports 3 by appeal to the more limited principle that, in cases like Barry's, agents can rely in their practical reasoning on what they know. In addition, Weatherson supports this more limited principle by appeal to what he claims is the best philosophical interpretation of decision theory. I will separately consider these two moves.

Let's consider whether it improves the argument to rely not on the fully general claim that knowledge is always sufficient for relying on a proposition in one's practical reasoning, but only on the more limited claim that, *in cases like the one under consideration*, it is so sufficient. He characterises the relevant set of cases as ones where one has to make a relatively simple bet which affects only the person making the bet. Some intellectualists who deny the general principle that knowledge is always sufficient for one to rely on a proposition in one's practical reasoning also deny the existential claim that sometimes knowledge is so sufficient. For instance, some hold that knowledge is never sufficient and that some stronger epistemic position is always required, such as knowledge of knowledge. However, intellectualists who deny *Sufficiency* need not hold that knowledge is never enough to rely on a proposition in practical reasoning. Instead, an intellectualist may hold that there is no fixed epistemic standard for relying on a proposition in one's practical reasoning. Rather, the strength of epistemic position required varies from case to case according to the stakes: sometimes it is less than the invariant threshold of strength of epistemic position required for knowledge, and sometimes more (Brown 2008; Gerken 2011). So, there need be no disagreement between Weatherson and the latter kind of intellectualist over the truth of the following existential claim: in some practical reasoning situations, if a subject knows that *p*, then she can rely on *p* in her practical reasoning. Of course, when it comes to the case of Barry, a non-sceptical intellectualist should say that Barry knows the proposition but isn't in a good enough epistemic position to rely on it in his practical reasoning. How, then, can Weatherson support his view about Barry (that since it's not rationally permissible for Barry to take the bet he doesn't know)? He won't appeal, as Fantl and McGrath do, to any argument for the general truth of *Sufficiency*. Instead, Weatherson attempts to appeal to the best philosophical interpretation of decision theory to support the view. Let's see how successful this is.

Weatherson's idea seems to be that he can defend the crucial premise 3 by appeal to the following principles concerning the best philosophical interpretation of decision theory:

A. It is rationally permissible, in cases like Barry's, to take dominating options.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

B. It is always right to model decision problems by including what the agent knows in the framework. That is, our decision table should include what the agent knows about the payoffs in different states, and leave off any state the agent knows not to obtain.

Notice that the first principle, unlike *Sufficiency*, is not a universal principle. Rather, the first principle is explicitly restricted to cases in which one has to make a relatively simple bet which affects only the person making the bet. I will consider later how Weatherson deals with surgeon style cases, in which one person's action (the surgeon's action) affects another. But, for now, let us see how assumptions A and B enable Weatherson to defend step 3 of his argument: if Barry still knows that *p* after the bet is offered, it is rationally permissible for him to take the bet. Consider the decision table for the bet which Barry faces.

	Beth	**Not-Beth**
Take bet	$10	1000 years of torture
Decline bet	Status quo	Status quo

Granting assumption B and that Barry knows that Beth is playing in the park, then Barry can ignore the right-hand column of the table and just concentrate on the left-hand column. Ignoring the right-hand column of the table, taking the bet dominates declining the bet. By assumption A, it is rationally permissible for Barry to take the dominating option in this case. So, on the assumption of A and B, it is rationally permissible for Barry to take the bet.

Notice that Weatherson's defence of step 3 requires both A and B. Later, I will cast doubt on assumption B. If B is false, then even if Barry knows that the musician is Beth, it doesn't follow that he can cross off the right-hand column of the decision table. As a result, it doesn't follow that taking the bet dominates declining the bet. So even with assumption A in place, namely that it is rationally permissible in cases like Barry's to take dominating options, it wouldn't follow that it is rationally permissible for Barry to take the bet.

Weatherson hopes to defend assumptions A and B by appeal to decision theory. We might wonder how appeal to decision theory could help. Given assumption A, assumption B amounts to the claim that, in cases like Barry's, if one knows that *p*, one can treat *p* as a certainty in one's practical reasoning. Of course, the latter assumption is precisely what is in dispute in the philosophical debate over *Sufficiency*. Fallibilist traditional invariantists[8] hold that knowledge is compatible with the possibility of error so that when the stakes get high enough, one shouldn't rely on what one knows in one's practical reasoning. Weatherson hopes to determine this important philosophical issue by appeal to decision theory. But, it's not clear that decision theory by itself will determine the philosophical issue about when one can treat a proposition as a certainty in one's practical reasoning. If one wants to settle this issue, one needs to discuss a range of different potential suggestions including that there is no constant standard, as well as various possible suggestions

8 I am assuming an intellectualist version of fallibilist invariantism rather than an anti-intellectualist version such as that defended by Fantl and McGrath (2009).

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

for what the constant standard is (justified true belief, justified belief, knowledge, knowledge of knowledge, etc.), as well as arguments about whether any possible suggested standard is necessary, sufficient, or both necessary and sufficient. But decision theory is not focused on these issues. Indeed, some suggest that the correct moral of decision theory is to reject traditional philosophical talk of belief and so knowledge, and instead only use the notion of degrees of belief. Weatherson may grant that decision theory itself has not been concerned with these issues but he may well still claim that the best philosophical interpretation of decision theory involves assumptions A and B. Let's consider his argument.

Weatherson gives two intuitive examples to try to defend B (S1.1). These involve students in an introductory decision theory course who seem to misunderstand decision tables. Suppose that the instructor gives the students a decision table.

	State 1	State 2
Choice 1	−$200	$1,000
Choice 2	−$100	$1,500

The instructor wants the students to see that choice 2 dominates choice 1, and so it is rational to take choice 2. However, two students fail to see this. One student, Stewart, thinks that the agent might not be able to collect $1,500 if she makes the second choice and the state of the world is state 2. For he thinks that people aren't always as reliable about paying out on bets when the amount of money is over $1,000. Another student, Dominic, suggests that the table should include a third state of the world in which there is a vengeful God who will take everyone who chooses choice 1 to heaven and everyone who chooses choice 2 to hell, so that choice 2 does not dominate choice 1. In different ways, these two students seem to have misunderstood the decision table. Weatherson's proffered explanation is that they have misunderstood the rules of decision tables and, in particular, that they haven't latched onto assumption B, namely that it is always right to model decision problems by including what the agent knows in the framework. Since we know there is no vengeful God, we don't need to bother about the state of the world in which there is a vengeful God. In addition, Weatherson claims, it is implicit that any value put into a cell of the table is such that we know that the combination of the relevant choice and state will yield that value. So, when the instructor asks the students to consider this table, she is asking them to assume that the agent knows that the combination of each state and choice yields the value stated in the table. For this reason, Stewart's objection is a misunderstanding. In this way, then, Weatherson claims that we can explain how the students misunderstand the table by appealing to the idea that the best interpretation of decision theory assumes B, that it always right to model decision problems by including what the agent knows in the framework. If B is true, then choice 2 dominates choice 1. If we make the further assumption A, that it is rationally permissible in cases like this to take dominating options, then it is rationally permissible for the agent to take option 2.

We have seen, then, how Weatherson defends his assumption B by saying that it can help explain the misunderstandings of the students, Stewart and Dominic. However, Weatherson's case for this assumption would be undermined if there were an alternative

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

explanation of the misunderstanding which does not rest on B. This is what I will now argue. I will suggest that we can explain the misunderstanding by the alternative assumption that it is implicit that the instructor is stipulating a problem for the students to solve. Thus it is implicit that the students should attempt to solve the task set by the instructor (what should the agent do?) on the basis of the information in the table. More fully, they should take it as given that the table correctly states the payoffs for the relevant combinations of choices and states of the world, and should assume that the table states all the relevant information. Adding more columns or rows to the table, or suggesting different values for the cells in the existing table would change the problem set. If the students are allowed to add extra columns or change the values in the cells of the table as they like, then they simply change the problem they are addressing. Further, if each student makes different changes to the decision table, there will be no single answer to the different problems they are addressing. But presumably the instructor wanted the class as a whole to concentrate on a single problem, namely the problem set, rather than that each student consider some different problem of their own making. In effect, the students have been given a conditional task: assuming the payoffs stated in the table and that the table states all the relevant information, what should the agent do? Each of the students in question misunderstands this task and, in effect, suggests that they, or the entire class, should address some other problem, rather than the problem set.

We've seen how we can explain the misunderstanding of the students without endorsing Weatherson's assumption B. However, Weatherson has a second argument to support his assumption B, that it is always right to model decision problems by including what the agent knows in the framework. He suggests, first, that some epistemic relation guides the construction of decision tables and, second, that there are reasons to prefer knowledge for this role to stronger or weaker epistemic relations. To see the role of epistemic relations in constructing decision tables, consider a decision faced by a real world agent. Suppose that she has been offered a bet costing $1,000 and, in fact, if she takes the bet she will win $2,000 thus recouping the cost of the bet and winning an additional $1,000; if she declines the bet, her position remains unaltered. Nonetheless, it is not rational for her to take the bet if she has no idea that the bet in fact has these payoffs. For instance, it would not be rational for her to take the bet if it amounts to betting that a ball will land on number 28 on a roulette wheel, where that outcome has a probability of 1/38 on her evidence. This suggests that it is not the actual outcome of the bet which matters, but one's epistemic relation to it.

Even if we accept this, it is so far open just what is the relevant epistemic relation. Weatherson suggests that the best philosophical interpretation of decision tables assumes that the agent stands in the knowledge relation to the table. He admits that it is hard to find counterexamples to the alternative suggestion that the relevant relation is justified true belief. However, he argues against this suggestion by arguing that justified true belief is a less natural relation than knowledge, and that we should prefer more to less natural relations. Further, Weatherson plausibly argues that we should not in general require a state stronger than knowledge, such as knowledge of knowledge, given that agents often lack knowledge of knowledge, and yet we want decision theory to be applicable even so. He concludes that the relevant epistemic relation is knowledge.

Weatherson's argument assumes that there is some single epistemic relation which guides the construction of decision tables. He then argues that this single relation is not less than knowledge (since anything less would be a less natural relation than knowledge)

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

and it is not more than knowledge (since agents often lack knowledge of knowledge). One might question his suggestion that any relation less than knowledge, such as justified true belief, is a less natural relation than knowledge and that we should prefer more natural relations. Further, we may question why we should assume that there is some single epistemic relation which guides the construction of decision tables. If we were to question this assumption, then the fact that we should not generally require a relation stronger than knowledge, such as knowledge of knowledge, would not show that some notion stronger than knowledge is not sometimes required. To see the issue, it is useful to draw an analogy with the knowledge norm of practical reasoning. Some assume that there is a single epistemic relation which is the norm for relying on a proposition in practical reasoning. If one makes this assumption, then one could defend the idea that the single epistemic relation is knowledge by arguing against other alternatives. In particular, one could object to the suggestion that something stronger than knowledge, say knowledge of knowledge, is always required, by arguing that it is too demanding to always require something stronger than knowledge. However, if one denies that there is a single epistemic relation which is the norm of practical reasoning, then this form of argument fails. In particular, that knowledge of knowledge is not always required does not show that it is never required. More generally, that something stronger than knowledge is not always required does not show that it is never required. Similarly, we may challenge Weatherson's argument for the claim that knowledge guides the construction of decision tables. At best, Weatherson has established a conditional claim: if there is a single relation that guides the construction of decision tables, then that relation is knowledge. But why should we assume that there is such a single relation? Perhaps there is no single epistemic relation which guides the construction of decision tables. If so, then the fact that it would be too demanding to always require more than knowledge, say knowledge of knowledge, does not show that it is too demanding to sometimes require more than knowledge. On the suggested view, the epistemic relation sufficient for writing states onto the decision table varies so that it is sometimes less than knowledge, sometimes knowledge, and sometimes more than knowledge.

Let me summarise the difficulties for Weatherson's positive arguments for his assumption B, namely it is always right to model decision problems by using what the agent knows in the framework. We saw that it is not clear that the best interpretation of decision theory is committed to this assumption. We can explain the mistakes made by the beginning students Stewart and Dominic without this assumption. Further, even if we accept that an agent's epistemic relation to the table is crucial, Weatherson's argument for the claim that the relevant epistemic relation is knowledge involves three doubtful assumptions: first, that any relation weaker than knowledge is less natural than knowledge, second, that we should prefer more to less natural relations and, third, that there is a single epistemic relation which guides the construction of decision tables. The third assumption needs defence as someone could reasonably claim that there is no such single epistemic relation. On the alternative suggested, the epistemic relation required to strike a state off a decision table may vary so that it may be more or less than knowledge.

We have seen that Weatherson's case for his crucial assumption B is problematic. As a result, he is yet to show that he can rely on this assumption in arguments for anti-intellectualism. Further, I will now show that the assumptions crucial to Weatherson's argument face intuitive objections. Thus, we will see that Weatherson's argument faces analogous problems to the argument offered by Fantl and McGrath. Both arguments

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

rely on crucial premises which face intuitive counterexamples and the alleged independent argument they offer for these assumptions is problematic. While the assumptions used in Weatherson's argument differ from the premise *Sufficiency* used in Fantl and McGrath's argument, we will see that the kinds of cases which provide intuitive counterexamples to *Sufficiency* also turn out to provide intuitive counterexamples to Weatherson's assumptions A and B.

Let us start by considering Reed's jellybean case (Reed 2011). Sally is participating in a psychological experiment which involves her being asked the question 'Was Caesar born in 100 BC?' If she gives the right answer, then she gets one jellybean. If she gives the wrong answer, then she gets a nasty electric shock. It is costless to refuse to answer. As Reed presents the case, it is intuitive that Sally knows the proposition in question and takes herself to know it, but she is not in a good enough epistemic position to rely on it in her practical reasoning. Indeed, she refuses to answer the question, thinking 'I know but it's not worth risking it'. Let us take this intuitive description at face value (i.e. take it that Sally knows the proposition but that it is sensible for her to decline the bet). The decision table for the jellybean case is as shown here.

	C	**Not-C**
Take bet	One jellybean	Nasty electric shock
Decline bet	Status quo	Status quo

On Weatherson's assumption B, if Sally knows that Caesar was born in 100 BC she can cross off the right-hand column of the table. But, then, taking the bet dominates declining it. By assumption A, it is rational for her to take the dominating option. But this seems to be the wrong result. Surely, she's rational to decline the bet. So Weatherson's assumptions when combined with the intuitive description of the case seem to give us a problematic result. Further, from the perspective of traditional fallibilism about knowledge, assumption B is implausible when applied to the case of Sally. After all, for a fallibilist, knowledge is compatible with a small chance of error, so even if Sally knows that Caesar was born in 100 BC, the right-hand column should still remain in the decision table. To cross off the right-hand column would be to treat what is known as incompatible with a chance of error. If one rejects B as I'm suggesting, then even if Sally knows that Caesar was born in 100 BC, it doesn't follow that she can cross off the right-hand column of the table. So taking the bet does not dominate declining it. As a result, even if assumption A is true, it cannot be applied to argue that Sally is rational to take the bet. (Note, that at this point Weatherson cannot just appeal to his assumptions A and B to argue that, since Sally shouldn't take the bet, she doesn't know the answer to the question. For, as we have seen, the positive arguments he uses to try and establish these assumptions, and especially B, are problematic.)

Weatherson may try to defend his assumptions A and B against the intuitive challenge presented by the jellybean case by claiming that, given the high stakes of the decision facing Sally, she lacks knowledge that Caesar was born in 100 BC. Thus, even on his assumption B, namely that it is always right to model decision problems using the agent's knowledge, he may allow that the right-hand column should remain in the decision table so that taking the bet does not dominate declining it. Thus, he may say, his account

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

of decision theory is compatible with its being rationally permissible for Sally to decline the bet. However, this would be to apply anti-intellectualism to the jellybean case to defend his interpretation of decision theory, rather than to defend his interpretation of decision theory from principles independent of anti-intellectualism. But if the aim was to use decision theory to provide independent motivation for anti-intellectualism, then it is clearly illegitimate to defend the interpretation of decision theory by appeal to anti-intellectualism. Rather what we were interested in was a defence of the relevant interpretation of decision theory which does not involve prior appeal to anti-intellectualism.

It seems, then, that Weatherson's assumptions give us an implausible result about the jellybean case. To the extent that Weatherson appeals to anti-intellectualism to avoid this problematic consequence, he seems to be employing anti-intellectualism to defend his assumptions about decision theory, rather than providing independent motivation for those principles in such a way that they can be used to motivate anti-intellectualism.

Now let us consider the surgeon case in which, intuitively, the surgeon knows that it is the left kidney but is not in a good enough epistemic position to operate without first checking. It may seem natural to model the surgeon's choice using this table.

	Left	**Right**
Remove left	1	−1
Remove right	−1	1
Check notes	1 − m	1 − m

This way of modelling the decision reflects the idea that there is a small additional cost of checking so that operating after checking has less value than operating without checking. According to Weatherson's assumption B, we should model this decision problem using the surgeon's knowledge. So assuming that the surgeon knows that it is the left kidney, we should delete the right-hand column of the decision table. However, having done this, removing the left kidney without checking dominates removing the left kidney after checking, for there is a small additional cost of checking. Despite this, Weatherson could seek to avoid the conclusion that the surgeon should remove the left kidney without checking by arguing that, in cases like this, his assumption A does not apply according to which it is rationally permissible to take dominating options. Recall that he explicitly denies that he is assuming that it is always rationally permissible to take dominating options. Rather, according to assumption A of his argument about Barry, 'it is rationally permissible, in cases like Barry's, to take dominating options'. Crucially, Barry's case, unlike that of the surgeon, involves a decision affecting only the agent and not also some third party. Why, though, might it not be rationally permissible to take the dominating option of removing the left kidney without checking? Weatherson suggests that there is a procedural rule in good hospitals requiring that in every operation one must double-check what organ or limb to operate on. Thus, even though operating without checking dominates operating after checking, the procedural rule eliminates the option of operating without first checking.

An alternative way for Weatherson to deal with the surgeon example would be by allowing the procedural rules to affect the values entered in the cells of the table. Thus,

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

we may suppose that there is a negative value associated with flouting the procedural rules, n, which is larger than the negative value of checking, m, so that removing the left kidney without checking has an overall value lower than removing it after checking. On this alternative, the decision table would be as shown.

	Left	Right
Remove left kidney	1 – n	−1 – n
Remove right kidney	−1 – n	1 – n
Check notes	1 – m	1 – m

In this case, assuming B and that the surgeon knows that it is the left kidney, removing the left kidney after checking dominates removing the left kidney without checking, given that the disvalue of flouting the procedural rules, n, is greater than the disvalue of checking, m. On this presentation of the problem, Weatherson need not claim that this is a case in which it is not rationally permissible to take a dominating option. However, we need to ask why it is that there is a negative value associated with flouting the procedures which is larger than the negative value of checking? As before, Weatherson may claim that enforcement of procedural rules in hospitals leads to better outcomes overall than if there were no such rule.

We've seen that both suggested ways of dealing with the surgeon example rely on the idea that there is something special about the case, say that there are procedural rules in force in hospitals which either overrule the generalisation that it is rationally permissible to take dominating options, or affect the values entered in the table so that removing the left kidney after checking dominates removing the left kidney without checking. However, it is not clear that talk of procedural rules really deals with the surgeon example, and other problematic cases for the anti-intellectualist. We would presumably think that the surgeon ought to double-check which kidney to take out even if there were no procedural rules established in the hospital requiring him or her to do so. Furthermore, there are a range of examples in which, intuitively, a subject knows a proposition yet shouldn't rely on it in her practical reasoning which do not involve institutional settings in which procedural rules may be established. For instance, Brown's affair case involves a subject who knows that her friend's wife is having an affair, but isn't in a strong enough epistemic position to act on that information yet. In reply to her friend's question, 'Why didn't you say something? You've known of the affair for ages', she replies, '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' (Brown 2008). Further, many of the problematic cases for anti-intellectualism involve neither harm to third parties nor institutional settings with professional norms (e.g. the jellybean case). In such cases, it is hard to defend either of Weatherson's suggestions for dealing with the surgeon example. Take the first strategy, of allowing that there are cases in which one is not rationally permitted to take a dominating option. Once we realise that the problematic cases include cases which do not involve harm to third parties or institutional norms, it seems that this first strategy risks multiplying the kinds of case in which one is not rationally permitted to take a dominating option. But this limits the interest and importance of decision theory to decision-making. Take the second strategy, of putting non-standard values into the decision table itself. This might seem acceptable to the extent that there are just a few

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

cases to be dealt with in this way, cases which we can understand to be distinct in terms of, say, harm to third parties or institutional settings. But once we see the range of examples which cause problems for the anti-intellectualist, we would need to put non-standard values into the decision tables for a great many problems. Applying the strategy of inserting non-standard values into tables seems implausible when applied across such a wide range of cases.

In conclusion, Weatherson's appeal to the philosophical interpretation of decision theory fails to get around the problems facing an argument to anti-intellectualism from *Sufficiency*. Even though Weatherson's argument does not itself appeal to *Sufficiency*, it does not surmount these problems. Weatherson's argument replaces the appeal to *Sufficiency* with appeal to his assumptions A and B. But Weatherson's own arguments for these assumptions, and especially B, are problematic. Further, some of the very cases that present intuitive counterexamples to *Sufficiency* also present intuitive counterexamples to the combination of A and B.

5. CONCLUSION

We have been examining a number of different variants of the suggestion that anti-intellectualism is supported by putative connections between knowledge and practical reasoning. One such classic argument exploits *Sufficiency*: if a subject knows a proposition, then she is in a good enough epistemic position to rely on it in her practical reasoning. However this argument faces difficulties. *Sufficiency* is open to putative counterexamples. These undermine both the attempt to defend *Sufficiency* by appeal to intuitions about the appropriate criticism and defence of actions, and Fantl and McGrath's attempted principled argument for *Sufficiency* from general principles connecting knowledge, reasons, and action. We then examined a second variant of the general strategy provided by Weatherson. Weatherson attempts to appeal to certain principles about decision theory to support anti-intellectualism without committing himself to the universal principle of *Sufficiency*. However, we saw that his interpretation of decision theory, and especially assumption B, is controversial and looks particularly implausible when applied to the putative counterexamples to *Sufficiency*. I conclude that, so far, there is no plausible argument for anti-intellectualism from considerations concerning practical reasoning.

ACKNOWLEDGEMENTS

Thanks to Mikkel Gerken, Nick Hughes, Matthew McGrath, and Brian Weatherson for discussion and useful comments on the paper.

REFERENCES

Brown, Jessica. 2008. 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.' *Nous*, 42(2): 167–89.

Buckwalter, Wesley. 2010. 'Knowledge isn't Closed on Saturdays.' *Review of Philosophy and Psychology*, 1: 395–406.

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press

Fantl, Jeremy, and McGrath, Matthew. 2002. 'Evidence, Pragmatics, and Justification.' *Philosophical Review*, 111(1): 67–94.

2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

2012. 'Shifty Epistemology.' In J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.

Feltz, Adam, and Zarpentine, Chris. 2010. 'Do you Know More When it Matters Less?', *Philosophical Psychology*, 23(5): 683–706.

Gerken, Mikkel. 2011. 'Warrant and Action.' *Synthese*, 178: 529–47.

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

Hawthorne, John, and Stanley, Jason. 2008. 'Knowledge and Action.' *Journal of Philosophy*, 105 (10): 571–90.

Hill, Christopher, and Schechter, Joshua. 2007. 'Hawthorne's Lottery Puzzle and the Nature of Belief.' *Philosophical Issues*, 17: 102–22.

Knobe, Joshua, and Schaffer, Jonathan. Forthcoming. 'Contrastive Knowledge Surveyed.' *Nous*, publ. online: 15 Dec. 2010, doi: 10.1111/j.1468-0068.2010.00795.x.

Lackey, Jennifer. 2011. 'Assertion and Isolated Secondhand Knowledge.' In J. Brown and H. Cappelen (eds), *Assertion*, pp. 251–76. Oxford: Oxford University Press.

Littlejohn, Clayton. 2009. 'Must we Act Only on What we Know?' *Journal of Philosophy*, 463–73.

May, Joshua, Sinnott-Armstrong, Walter, Hull, Jay, and Zimmerman, Aaron. 2009. 'Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study.' *Review of Philosophy and Psychology*, 1: 265–73.

Neta, Ram. 2009. 'Treating Something as a Reason for Action.' *Nous*, 43: 684–99.

Pinillos, Ángel. 2012. 'Knowledge, Experiments and Practical Interests.' In J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.

Reed, Baron. 2011. 'Defence of Stable Invariantism.' *Nous*, 44(2): 224–44.

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

Weatherson, Brian. 2012. 'Knowledge, Bets and Interests'. In J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.

Forthcoming. 'Defending Interest Relative Invariantism'. http://brian.weatherson.org/papers.shtml.

JESSICA BROWN is currently Director of, and Professor in, the Arché Philosophical Research Centre at the University of St Andrews. She works primarily in epistemology on topics including contextualism, invariantism, epistemic norms of practical reasoning and assertion, scepticism and the transmission of warrant. She is currently directing a major research project on philosophical methodology. Arché Research Centre, Department of Philosophy, University of St Andrews, Edgecliffe, The Scores, St Andrews, Fife, KY16 9AL, Scotland

https://doi.org/10.1017/epi.2011.4 Published online by Cambridge University Press