

Evidence and armchair access

Clayton Mitchell Littlejohn

Received: 14 January 2008 / Accepted: 18 November 2009
© Springer Science+Business Media B.V. 2009

Abstract In this paper, I shall discuss a problem that arises when you try to combine an attractive account of what constitutes evidence with an independently plausible account of the kind of access we have to our evidence. According to $E = K$, our evidence consists of what we know. According to the principle of armchair access, we can know from the armchair what our evidence is. Combined, these claims entail that we can have armchair knowledge of the external world. Because it seems that the principle of armchair access is supported by widely shared intuitions about epistemic rationality, it seems we ought to embrace an internalist conception of evidence. I shall argue that this response is mistaken. Because externalism about evidence can accommodate the relevant intuitions about epistemic rationality, the principle of armchair access is unmotivated. We also have independent reasons for preferring externalism about evidence to the principle of armchair access.

Keywords Evidence · Justification · Knowledge · Epistemic rationality · Reasons for belief

C. M. Littlejohn (✉)
University of Texas at San Antonio, San Antonio, USA
e-mail: cmlittlejohn@yahoo.com

1 Introduction

I want to discuss a problem that arises when we try to accommodate two independently plausible claims about evidence.¹ The first is a claim about what our evidence consists of; namely, that our evidence consists of what we know:

(E = K) S's evidence includes p iff S knows p .²

The second is a claim about the kind of access we have to our evidence:

(AA) If S's evidence includes p , it is possible for S to know from the armchair that p is part of S's evidence.³

To say that S knows p from the armchair is to say that p is known to S in a way that does not depend constitutively upon S's experiences. It is to say that if p is part of S's evidence, it should be possible for S to know that p is part of S's evidence a priori or through introspection or reflection. The principle is prima facie plausible because it often seems that when we form non-culpably mistaken beliefs about the external world, the mistakes do not derive from prior mistakes about what our evidence consisted of. Upon learning that we were mistaken, we tend to think that our evidence misled, not that we were misled about what our evidence was. This suggests that when we make mistakes about how things are in the external world, there was something we were right about (i.e., our reasons for forming the mistaken belief) and that what we were right about had to concern some subject matter other than what we were mistaken about (e.g., considerations that pertained to us and how things seemed to us rather than how things are outside of us). In turn, this suggests that the evidence we have for our beliefs must be provided by matters 'internal' to our perspective rather than facts about the external world to which we do not have unproblematic access.

To see what the problem is, let's suppose for the sake of our discussion that the skeptic is wrong and that it is possible to know contingent propositions about the external world such as the proposition that I have hands. Given these assumptions and E = K, it follows that:

¹ Silins (2005) was first to raise this objection to Williamson's (2000a) account of evidence. He thinks this problem arises for any view that implies that it is possible for two subjects in the same (non-factive) mental states to have different evidence. Neta and Pritchard (2007) defend some of McDowell's proposals about reasons for belief from a similar objection. Brewer (1999) also considers a similar objection in defending his account of perceptual reasons. The idea that we have a kind of privileged access to our evidence or reasons for belief is hardly new. See Ginet (1975) for a defense of an access requirement stronger than (AA).

² Williamson (2000a) defends this view. Brewer (1999), Hyman (1999), Hyman (2006), Maher (1996), and Unger (1975) face essentially the same difficulty that Williamson faces insofar as their views can accommodate AA only if it is possible to have armchair access to contingent matters of fact.

³ An anonymous referee asked whether armchair knowledge and a priori knowledge amounted to the same thing. This depends upon what counts as a priori knowledge. If you think that a priori knowledge does not include introspective knowledge, these notions do not come to the same thing. Armchair knowledge might include a priori knowledge (e.g., knowledge of moral truths or mathematical truths that is had by reflection or intuition), but it would also include introspective knowledge of one's own mental life. Silins (2005) defends the principle of armchair access. Audi (1993, p. 344), Boghossian (2008, p. 145), and Ginet (1975, p. 34) seem to endorse the principle as well. Just to be clear, it is not obvious that if S has armchair knowledge that p is part of her evidence she must have armchair knowledge that p is the case. Perhaps S can have armchair knowledge that the proposition *that she has hands* is part of her evidence even if she does not know that she has hands.

(E1) My evidence includes the proposition that I have hands.

It follows from (E1) and AA that:

(E2) It is possible for me to know that (E1) is true from the armchair.

If $E = K$ can be known from the armchair, it follows from this further assumption and (E2) that:

(E3) It is possible for me to know that I have hands from the armchair.

But, that seems absurd. Surely I cannot know from the armchair that I have hands.⁴

Our problem is at least superficially similar to the problem that McKinsey raised for externalism about thought content.⁵ If the contents of our thoughts depends upon contingent matters external to us (e.g., the presence of water), how can we consistently say that we have privileged access to our thoughts when we know that we cannot have privileged access to those aspects of the external world that go towards determining what the contents of our thoughts are? Similarly, if our evidence depended upon contingent matters external to us, how can we consistently say that we have a kind of privileged access to our evidence if we do not have privileged access to those aspects of the external world that go towards determining what our evidence is? Here, I shall defend an incompatibilist response. Although we cannot reconcile externalist conceptions of evidence with the principle of armchair access, we should reject the principle of armchair access and recognize that externalism about evidence can accommodate the considerations taken to motivate the principle of armchair access.

In Sect. 2, I shall explain the motivation behind the principle of armchair access and externalist approaches to evidence. In this section, we shall see why some say that the principle of armchair access is needed to make sense of a widely held intuition about epistemic rationality. In turn, we shall see why someone might think that externalism about evidence ought to be rejected on the grounds that it clashes with our intuitions about epistemic rationality. In Sect. 3, I shall explain why we ought to give up the principle of armchair access. Not only is there a good case to be made for externalism about evidence, the principle of armchair access generates implausible skeptical consequences. In Sect. 4, I shall explain how to reconcile externalism about evidence with widely shared intuitions about epistemic rationality.

⁴ Neta and Pritchard (2007) seem prepared to bite this bullet.

⁵ For discussion, see McKinsey (1991). While this is controversial, I believe that McKinsey's epistemological objections to thought content externalism rest on a mistake about the commitments of thought content externalism. If Putnam is right, we might know a priori that 'water' thoughts are wide but not that our 'water' thoughts depend upon the existence of water. (Our 'water' thoughts would have been wide in a dry world.) So, there is nothing we know a priori that allows us to knowingly deduce from the armchair that water exists if we know from the armchair that we are thinking that water is wet and know from the armchair that such thoughts have wide content. And, I doubt that knowing whether a thought is world-dependent in the way our 'water' thoughts happen to be but 'unicorn' thoughts happen not to be is necessary for knowing what we are thinking. I do not believe that critics of externalism about evidence have made any analogous mistakes about the commitments of evidential externalism, but I need not take a stand on that matter in this paper.

2 Evidential externalism and armchair access

In this section, I want to explain the motivation behind the principle of armchair access and the motivation for externalism about evidence. Before we get to that, some preliminary remarks about the notion of evidence at issue are in order.

First, we typically think of evidence as being evidence *for* some particular claim, belief, or hypothesis. It might seem strange to some to think of a subject's evidence as being coextensive with what the subject knows since evidence is always evidence *for* something or other.⁶ It might be useful to distinguish the propositions that are a subject's evidence *for* believing some particular proposition and the subject's evidence *simpliciter*. To understand this talk of evidence *simpliciter* (i.e., talk about evidence that does not mention any propositions the piece of evidence is evidence for), this analogy is helpful.⁷ We can say that someone has such and such piece of evidence without specifying what that is evidence for much in the same way that we can list the ingredients in someone's kitchen without saying what those ingredients are ingredients for. This manner of speaking should not be taken to support the dubious idea that there could be ingredients *simpliciter* that are not ingredients *for* any dish at all. We should be able to talk about what is in an individual's stock or fund of evidence without being committed to the bizarre idea that there are propositions that are evidence *simpliciter* that are not evidence *for* anything. According to $E = K$, nothing gets to be part of that fund of evidence unless the subject knows that the proposition is true and nothing beyond knowledge is necessary for that proposition to be included in that subject's evidence. It seems that someone can consistently say that these are the requirements for something to get into a fund of evidence while adding that nothing in that fund is evidence without being evidence *for* something. I cannot think of a case in which there is a proposition I know that could not serve as evidence for something I might consider believing (e.g., obvious consequences of what is known). In what follows, when I say that p is part of a subject's evidence without specifying what this is evidence *for*, this is like saying that someone is a brother or aunt without saying whose aunt or brother it is.

Second, some authors find it helpful to distinguish between objective and subjective conceptions of evidence. On an objective conception of evidence such as Achinstein's conception of veridical evidence, it cannot be that p is evidence for S 's belief that q unless p is true.⁸ I see nothing wrong with acknowledging that we do use this concept of evidence while acknowledging that we also sometimes use something along the lines of Achinstein's subjective conception of evidence on which it can be the case that p is S 's evidence for q even if p is false (unbeknownst to S). Once we distinguish objective from subjective conceptions of evidence, the question arises as to which of these concepts of evidence is at issue. One worry is that unless we fix on some single

⁶ An anonymous referee for this journal expressed the concern that $E = K$ commits us to an overly inclusive conception of evidence, but I should hope that $E = K$ is consistent with (a) any plausible view about when p can count as a subject's evidence for q and (b) the insistence that nothing can be evidence *simpliciter* unless it is evidence *for* something or other.

⁷ Hyman (2006, p. 892).

⁸ See Achinstein (1984).

notion, it is possible that both sides to this debate will talk past one another. True, someone might say, you should not combine $E = K$ with AA because the former has to do with an objective notion of evidence whereas the latter has to do with a subjective notion. Can we solve our problem simply by disambiguating?

I don't think that our problem can be solved this easily, at least not to the satisfaction of those who defend AA and $E = K$. Think about the dialectical situation. Those who defend AA take it upon themselves to argue that the incompatibility between AA and $E = K$ is reason to reject $E = K$. I doubt that either party to this dispute thinks that $E = K$ was ever intended as an account of the subjective conception of evidence. Those who defend AA either think that the very idea of an objective conception of evidence is suspect or insist that even if there is some objective conception of evidence there are still access requirements that facts or propositions must satisfy in order to be included in our stock of (objective) evidence. Those who defend AA and insist that $E = K$ is false thus seem to intend AA to be a constraint on the objective notion of evidence. If they intended AA to be a constraint only on some subjective notion and knew that $E = K$ was an account of the objective notion of evidence, their complaint about $E = K$ would be that it is not relevant to their concerns. Instead their complaint is that $E = K$ is false and their claim is that the incompatibility of $E = K$ with AA shows that it is false. Second, it seems there is some single, univocal notion of evidence at issue here that defenders of $E = K$ and AA can fight about. Parties to this debate think that the conception of evidence they have described is the notion of evidence that is involved in the justification of belief. These parties are not interested in the *activity* of giving justifications for a belief that might fail to show that the relevant belief is justified. Both sides are trying to characterize the kind of evidence that is largely responsible for ensuring that our beliefs have the *property* of being justified.⁹ If you worry about the potential ambiguity in talk of 'evidence', we can disambiguate by introducing a subjective and objective conception of evidence and ask whether the objective notion plays any role in the justification of belief. Those who defend AA insist that the only evidence that can potentially justify our beliefs is the kind of evidence we can know from the armchair to belong to our stock of evidence. Those who defend $E = K$ insist that it is a mistake to think that potential justifiers are limited to those things we can know from the armchair to belong to our evidence and insist that some propositions can serve as justifying evidence for our beliefs even if we cannot know from the armchair that these propositions are included in our evidence.

⁹ Silins (2005, p. 376) and Williamson (2000a, p. 185) are explicit on this point, insisting that the notion of evidence at issue is the notion of that which justifies belief. Of course, as an anonymous referee pointed out, we can draw a further distinction between objective and subjective notions of justification so this might only push the problem further back. In response, I will note that appealing to the distinction between subjective and objective notions of justification in the hopes of resolving our dispute to the satisfaction of our disputants would require showing that these two sides are really talking past one another. I do not know how plausible it is to suggest that the disagreement between those who think that the conditions that determine whether a belief can be said to be 'justified' depend entirely on subjective matters and those who think that it depends in part upon more objective matters stems from their failure to appreciate that the other side is trying only to characterize an objective or subjective notion of justification. I should also add that some authors (e.g., Feldman (1988)) say that there is a distinction to be drawn between the objective and subjective justification of action but insist that there is no parallel distinction to be drawn in epistemology.

2.1 Armchair access

According to the principle of armchair access, your evidence is limited to that which you can know to be part of your evidence from the armchair. Why think that your stock of evidence is limited in this way? Suppose we think of evidence as that which we have to go on if we are to settle some question.¹⁰ If we think of evidence this way, it seems natural to say that we need better access to the evidence that bears on whether q than we do to the fact that q . This does not explain why we would need *armchair* access to our evidence. Suppose that you are sprawled out in bed wondering whether it rained last night. You cannot see outside your apartment because your shades are drawn. You need to gather evidence. Whatever evidence you might acquire, it should be that your access to this evidence is less problematic than the access you now have to the facts about last night's weather. Even if this is so, we can still say that the fact that the streets are wet is part of your evidence for believing that it rained last night once you observe the wet streets having drawn the shades. The principle of AA does rule this out. Thus, the platitude that evidence is what we have to go on to settle a question does little to motivate AA.

Let's try a different tack. The evidential externalist ('externalist' hereafter) thinks that it is possible for two subjects in precisely the same non-factive mental states to have different bodies of evidence. Internalists deny this. Their disagreement concerns the status of the following supervenience thesis:

(E-INT) Necessarily, if two subjects, S and S', are in precisely the same non-factive mental states, S and S' have precisely the same evidence.¹¹

Anyone who accepts $E = K$ will deny E-INT and so anyone who accepts $E = K$ counts as an externalist as the term is being used here.¹² To see why, suppose S knows that she has hands but S' mistakenly believes she has hands because she has undergone a series of hallucinatory experiences. The proposition that S has hands will be included in the first subject's body of evidence but the proposition that S' has hands will not be included in the second subject's body of evidence. So, S will have evidence for the inferential belief in the disjunctive proposition that she has hands or cats are robots from outer space that S' does not have. It seems to Silins that AA is consistent with

¹⁰ See Kelly (2008).

¹¹ Conee and Feldman (2004) defend this thesis. Because some believe in factive mental states that have as contents propositions that represent aspects of the external world, it is important to distinguish INT from the claim that evidence strongly supervenes on a subject's mental states. $E = K$ is consistent with that thesis, but is inconsistent with INT. In epistemology, there are disputes between internalists and externalists that do not directly concern INT. For example, you might think of the internalist about evidence as someone who thinks that we have a kind of immediate access to our evidence. Maher (1996) thinks that our evidence includes propositions known to us by observation. If you think that observation gives us access to the external world, Maher's view is inconsistent with INT but would be consistent with a kind of access requirement on evidence.

¹² An anonymous referee thought it was odd to describe $E = K$ as a version of externalism about evidence. There are perfectly good ways of drawing the distinction between internalism and externalism on which $E = K$ counts as an internalist view (e.g., you might think that the internal includes everything to which you have access or everything which you know). I'm following the lead of authors such as Conee and Feldman (2004) and Silins (2005) in labeling Williamson's view 'externalist' because $E = K$ is incompatible with E-INT.

E-INT.¹³ Some might say that the same intuitions motivate AA and E-INT. As the denial of E-INT is often taken to be deeply problematic, this might suggest a natural rationale for accepting AA and for saying that the incompatibility of AA and $E = K$ is a good reason to reject externalist views about evidence.

Apart from finding E-INT intuitive, some have argued that E-INT is a consequence of a kind of internalism about epistemic rationality:

(R-INT) Necessarily, if two subjects, S and S' are in precisely the same non-factive mental states, S and S' are equally rational/reasonable in believing what they do.¹⁴

Many find R-INT quite intuitive. Suppose S is wildly successful in her epistemic endeavors. Whatever she believes, she knows to be true. Suppose S' is in precisely the same non-factive mental states but her beliefs fail to constitute knowledge for various reasons. Perhaps S' suffers the odd hallucination or is prone to forming justified but accidentally true beliefs. It seems that when S' fails to know for these reasons, she is no less reasonable than S is.¹⁵ Thus, the facts that differ between subjects in the same non-factive mental states seem not to matter to the rationality of these subjects' beliefs.

If epistemic rationality is primarily a matter of respecting the evidence, we can argue from R-INT to E-INT as follows. Suppose White and Plum are in precisely the same non-factive mental states. Both believe that Mustard is the killer. White saw Mustard kill his victim. Plum believes Mustard is the killer because of a series of hallucinatory experiences indistinguishable from White's veridical experiences. Suppose E-INT is false and $E = K$ is correct. White's evidence includes everything that Plum's evidence includes, but her evidence is not limited to Plum's evidence. Now, if Plum knew that her evidence included only the evidence someone would have in the 'bad' case (i.e., the case in which her beliefs are mistaken but she is in just the same non-factive mental states she is in now because of a hallucinatory experience), she ought to be significantly less confident in her beliefs than White is. If she knew this she ought to suspend judgment on some beliefs that White knows are true. If she knew

¹³ There is a controversial issue that I am bracketing, and that is that even if you adopt an internalist conception of evidence it is not obvious that we enjoy armchair access to our evidence. Suppose we say that a subject's evidence is either constituted by that subject's non-factive mental states or the contents of those states. If those contents are not narrowly individuated, the possibility of slow switching cases might cause trouble for AA. Before the switch, the subject might believe that water and oil do not mix. After the switch, the subject acquires the concept for water and is disposed to express beliefs involving that concept when using the word 'water'. When the switch is revealed, the subject might not be in a position to determine what her evidence is from the armchair because the subject needs empirical knowledge to determine whether the contents of her second-order judgments match the contents of the first-order beliefs she attributes to herself. It seems that the subject would also need empirical knowledge to know whether her judgments about what evidence she has would similarly depend upon empirical knowledge if her evidence is constituted (in part) by those first-order attitudes. This is a complication best discussed elsewhere.

¹⁴ In addition to [Conee and Feldman \(2004\)](#), [Audi \(1993\)](#), [Cohen \(1984\)](#), and [Wedgwood \(2002\)](#) defend this claim.

¹⁵ This sort of intuition figures prominently in defense of internalist supervenience theses. [Cohen \(1984\)](#) appeals to this intuition in saying that what justifies our beliefs must be what is common between a subject who has knowledge and her internal duplicate that is deceived by a Cartesian demon. [Turri \(2009\)](#) cites a similar intuition in arguing from R-INT to E-INT.

that her evidence included just the evidence someone has in the ‘bad’ case and did not adjust her attitudes accordingly, she would not be as reasonable or rational as White is. But, if $E = K$ is true she is not in a position to know that she has less evidence than White does. We can only say that Plum was as reasonable as White because Plum is ignorant of the fact that she has less evidence than someone in White’s position has. But this means that we have to say that she is rational or reasonable only because she does not know what her evidence really is. How can we say that she is nothing less than fully reasonable or rational and that epistemic rationality is a matter of respecting the evidence while saying that she does not know what her evidence is? This combination of claims is puzzling but this combination of claims is what externalists are committed to unless they deny that epistemic rationality is a matter of respecting the evidence. So, it seems there is a plausible line of argument from R-INT to E-INT. Given the intuitive plausibility of R-INT, we have an argument for E-INT and AA.

2.2 Externalism about evidence

The externalist thinks that it is possible for two subjects in precisely the same non-factive mental states to have different evidence. Because $E = K$ is incompatible with E-INT and AA, it seems to some that $E = K$ comes with heavy costs and we are better off without it. One reason not to dispense with externalist views like $E = K$ is that there are in fact good reasons to be an externalist about evidence. In this section, I will sketch a case for $E = K$ and for a weaker version of externalism.

There is some linguistic evidence that suggests that evidence consists of either facts or propositions. Consider:

- (1) There is good reason for her to believe that Plum is innocent; namely, that Plum’s prints were not the prints found on the murder weapon.

We ascribe reasons using that-clauses, so it seems that if reasons for belief just are bits of evidence, in ascribing someone evidence we refer to propositions or facts.¹⁶ Of course, not everyone believes that this sort of linguistic evidence is dispositive since not everyone believes that evidence consists of propositions or facts. Since the assumption that evidence consists of either facts or propositions is consistent with AA, however, let’s speak as if this view is correct ask whether evidence is subject to the sort of constraint that defenders of AA suggest that it is.

Consider two questions about the relation between evidence and knowledge:

- (Q1) Is S’s knowledge of p ’s truth *necessary* for p ’s inclusion in S’s body of evidence?
- (Q2) Is S’s knowledge of p ’s truth *sufficient* for p ’s inclusion in S’s body of evidence?

Let’s focus first on the claim that knowledge that p is true is necessary for p ’s inclusion in a subject’s body of evidence. Some have offered linguistic evidence in support of the necessity thesis, observing that it sounds contradictory to assert the following:

¹⁶ For arguments in support of the thesis that reasons are facts or propositions, see Dancy (2000), Darwall (1985), Neta (2008), Unger (1975), and Williamson (2000a).

- (2) White's reason for believing/saying Plum is innocent is that Plum was nowhere near the scene of the crime on the night of the murder but White does not know that Plum was nowhere near the scene of the crime on the night of the murder.

If this claim is contradictory, whenever p is S's reason for Φ -ing, S knows p .¹⁷

I have to confess that I do not place a great deal of confidence in this argument. In some contexts, we use 'know' loosely, as something akin to firmly held true belief.¹⁸ These are contexts in which we use 'know' in such a way that considerations having to do with accidental connections to the truth are irrelevant. Consider this exchange:

Mustard: What reason does White have to believe that Plum is innocent?

Scarlet: White's reason for believing Plum is innocent is that Plum was nowhere near the scene of the crime on the night of the murder.

In another conversational exchange, Peacock says to Green that White actually believes that Plum *was* near the scene of the crime. He then adds that Plum was at the scene of the crime on the night of the murder. As an outside observer, it seems to me that both of Peacock's assertions are inconsistent with Scarlet's response to Mustard's question. However, if Peacock starts to tell Green about Gettierish considerations (e.g., false newspaper stories that White knows nothing about that contain misleading evidence) we might be inclined to say that White does not know that Plum was nowhere near the scene of the crime for these purely Gettierish reasons. It does not seem to me, however, that such considerations thereby show that Scarlet's assertion is mistaken, false, or incorrect. And that suggests that it is a loose use of 'knows' that figures in (2).

Pointing to such Gettierish considerations is not a way of challenging the veracity of Scarlet's claim. This suggests that the fact that (2) sounds contradictory is not strong evidence for the necessity claim. I should add that there is a good reason to deny the necessity claim. Suppose that S knows that there is at least one barn visible from the road she has been driving down because she saw that there was a barn a little ways back. Suppose that S' fails to know that there is at least one barn visible from the road that she has been driving down because while she did in fact see a barn, there is a sufficient number of fakes in her vicinity for us to say that her belief failed to constitute knowledge. If we hold everything else equal (i.e., facts about their mental states, their dispositions to reason in certain ways, the accuracy of those mental states, etc. . . .), it strikes me as highly counterintuitive to say that S has evidence that S' lacks. This, however, is something that a defender of $E = K$ would have to accept.

If Gettier cases give us reason to revise $E = K$, they give us little reason to revise our account of evidence to bring it in line with AA. While I doubt that S must know that p is true for p to be included in S's evidence, some linguistic evidence suggests that p must be true if it is a piece of evidence. Consider two exchanges. Here's the first:

Scarlet: Does the prosecution have solid evidence against Mustard?

Green: Yes, they have all sorts of evidence against him: namely, that he was the last one to see the victim alive, that his alibi did not check out, that his fingerprints

¹⁷ See Unger (1975, p. 206).

¹⁸ See Goldman (2002).

were on the murder weapon, and that he had written a letter containing details the police think only the killer could have known.

Here's the second:

Plum: How good is the prosecution's evidence against Mustard?

Peacock: It seems that the evidence is pretty strong. However, Mustard's prints are not on the murder weapon, his alibi checks out, and he was not the last one seen with the victim. This is all perfectly consistent with the evidence that the prosecution does have.

It seems that Peacock's assertion flatly contradicts Green's assertion. But all that Peacock has done is assert that the falsity of certain propositions is consistent with other propositions about what the prosecution's evidence consists of. So, unless we say that claims about what someone's evidence consists of entail that those claims are true, it's hard to see how Peacock's assertion could contradict Green's assertion. Peacock's assertion speaks to the veracity of the prosecution's claims rather than speaking directly about the evidence that they have. This bit of linguistic evidence suggests that ascriptions of evidence are factive:

(ET) If S's evidence includes p , it is the case that p .

If we assert that some subject's evidence includes p , we cannot remain neutral with respect to the question as to whether p is true.¹⁹

Suppose that knowledge of p 's truth is not necessary for p 's inclusion in S's evidence. Suppose that truth is. Let's turn to our second question. Is knowledge *sufficient*? The sufficiency claim seems more plausible than the necessity claim. Suppose the logical connection between two propositions is sufficiently obvious to both the speaker and the subject the speaker ascribes knowledge to. It seems contradictory to say things like this:

(3) This is a candlestick and Plum knows it, but that is no reason for her to believe it is not a wrench.

To my ear, (3) is as bad as the abominable conjunctions that lead some to endorse closure principles for knowledge.²⁰ A natural explanation as to why (3) sounds so bad is that (3) cannot be true. But, the natural explanation as to why (3) cannot be true is that once you know that p is true you have crossed all of the hurdles that you must cross to acquire p as part of your evidence.

We think of evidence as something that provides a subject with a justification or a reason for a belief.²¹ Some have defended the view that knowledge of p 's truth is sufficient for it to be epistemically permissible to treat p as a reason for action or including p in a piece of practical reasoning.²² Various arguments have been given for the view, but rather than reviewing those, let me add one of my own. Were it not permissible to

¹⁹ Note that ET is consistent with E-INT and AA. Someone could combine these three views and say that our evidence consists of only propositions about our non-factive mental states.

²⁰ See DeRose (1995).

²¹ This is Silins' (2005, p. 376) gloss on the notion of evidence.

²² See Hawthorne and Stanley (2008).

treat p as a reason for action and include p in practical deliberation, there would be a conclusive reason to refrain from so doing. But, if there were a conclusive epistemic reason to refrain from including the belief that p is the case in practical deliberation, such a reason would constitute a conclusive reason not to believe p in the first place. Such a reason would ensure that the subject could not satisfy the justification condition on having knowledge that p . Thus, it seems that the principle that says that if you know p it is proper to treat p as a reason for the purposes of deliberation will admit of no counterexamples if justification is necessary for knowledge.

We should note that the principle says that knowledge of p 's truth ensures that it is proper for the subject to treat p as a reason for the purposes of deliberation, which is not the same thing as saying that when p is known p is a reason. Might it be possible to properly treat a non-reason as if it were a reason, say, when the subject is non-culpably ignorant of the fact that the non-reason is indeed not a reason? The distinction between what is proper to treat as a reason and what is actually a reason, it is not a distinction that those who accept AA would appeal to if they wanted to resist the argument I've just given for thesis that knowledge of p 's truth suffices for p 's inclusion in a body of evidence. According to AA, on the hypothesis that p is a piece of S's evidence, S is in a position to know that p is a piece of evidence. It is hard to see how someone could permissibly treat a non-reason as a reason when they were in a position to know that the non-reason is indeed not a reason. So, it is hard to see how someone who accepts AA could say that it is permissible to treat what you know as a reason for action while saying that it might nevertheless be the case that what is treated as a reason is no reason at all. It seems there is good reason to think that it is permissible to treat what you know as a reason for action. Thus, knowledge of p 's truth ought to suffice for p 's inclusion in your evidence if the reasons that properly figure in practical deliberation constitute evidence for certain beliefs (e.g., beliefs about what ought to be done, what there is reason to do, what would be regrettable not to do, etc. . .). The identification seems perfectly apt given the gloss on evidence as that which gives you a reason that can justify a belief.

Someone could be picky about the matter and insist on a distinction between basic and derived reasons for belief where basic reasons are acquired without inference or reasoning and derived reasons are derived from basic reasons by means of inference.²³ They might say that our evidence consists of only non-derivative reasons for belief. They might, for example, balk at the idea that deductive or inductive inference is a way of acquiring new evidence. If so, we might have to modify the proposal that asserts that knowledge of p 's truth suffices for p 's inclusion in a subject's evidence and say instead:

(NIKSE) If S knows p non-inferentially, S's evidence includes p .

²³ On Maher (1996) account of evidence, inference from things known through observation does not provide us with new evidence. Bird (2004) defends the idea that there can be inferential evidence. If Bird (2004) arguments are successful, then perhaps we do acquire new evidence via inference and the sceptical worry I raise in the next section is even worse than I suggest below. For if knowledge of p 's truth suffices for p 's inclusion in a subject's evidence regardless of whether that knowledge is inferential or non-inferential, the argument below would show that it follows from AA that knowledge of the external world is unattainable.

While individually NIKSE and ET are consistent with AA and E-INT, the conjunction of NIKSE and ET is incompatible with both AA and E-INT. Thus, anyone who embraces both ET and NIKSE has embraced a form of externalism about evidence. It is a version of externalism that is more modest than $E = K$. One of its virtues is that it does not generate the counterintuitive implications about Gettier cases that $E = K$ does. For example, if you accept $E = K$ you have to say that if one subject knows p but the second fails to know p just because she is in a Gettier case it thereby follows that the second has less evidence than the first. In the next Sect. 2 shall argue that we ought to give up AA rather than NIKSE and ET.

3 Against armchair access

Forced to choose between AA and externalism, we ought to choose externalism rather than AA because AA generates untoward skeptical consequences. I shall argue that if AA were true, we could not have non-inferential knowledge of facts about the external world. While it does not follow directly from this that we can never have knowledge of the external world, if you take a dim view of the claim that it is only via inference that we first acquire knowledge of the external world this is quite troubling.

Here is the argument:

- (1) My evidence is limited to propositions I can know from the armchair belong to my evidence [AA].
- (2) If the proposition *that I have hands* is part of my evidence, then I have hands [ET].
- (3) If I know that I have hands non-inferentially, the proposition *that I have hands* is part of my evidence [NIKSE].
- (4) If my evidence includes the proposition *that I have hands*, then I have hands and I know from armchair that I have hands [(1), (2), (3)].
- (5) It's absurd to think I could know from the armchair that I have hands [Assumption].
- (6) If my evidence cannot include the proposition *that I have hands*, either I'm handless or I cannot know non-inferentially that I have hands [NIKSE, ET].
- (7) My evidence cannot include the proposition *that I have hands* [(4), (5)].
- (8) Either I'm handless or I cannot know non-inferentially that I have hands [(6), (7)].
- (9) If I have hands, I cannot know non-inferentially that I do [(8)].
- (10) If I do not have hands, I cannot know non-inferentially that I do [Factivity of 'knows'].
- (11) Thus, whether I have hands or not, I cannot know non-inferentially that I have hands [(9), (10)].

I know non-inferentially that I have hands, so (11) is false. Which of the premises should we reject? The argument assumes that (5) is true, but this assumption is crucial to the initial objection to $E = K$. Deny it, and the initial objection to $E = K$ is undercut. I think that we ought to reject AA.

A point of clarification is in order. There is a difference between knowing that p is true and knowing that p is part of your evidence. Someone might challenge (4) on the

grounds that I have ignored this distinction.²⁴ Here is the justification for (4). Suppose that p is part of my evidence because p is known non-inferentially. If AA is true, it follows that I can know that p is part of my evidence from the armchair. Owing to the factivity of knowledge, if ET is true I can only know that p is part of my evidence if p is true. If ET is not only true but also known to be true from the armchair, I know from the armchair that p is part of my evidence only if p is true. Thus, given AA, NIKSE, and that ET is known from the armchair as assumptions, we can say that if p is part of my evidence I know not only that p is part of my evidence from the armchair, I also know from the armchair that p is true.²⁵

Someone sympathetic to AA might object that the skeptical argument against AA rests on two assumptions those who defend AA need not accept. The argument assumes that evidence consists of truths (ET) and that any truth that we know non-inferentially is included in our body of evidence (NIKSE). The quick response to this worry is to note that it seems that so far as skepticism is concerned, ET and NIKSE are innocuous. Taken in combination, they seem perfectly compatible with the sort of anti-skeptical view that most of us take for granted, a view on which we have expansive knowledge of the external world including immediate knowledge of our immediate surroundings. The same cannot be said for AA. By insisting that non-inferential knowledge that some proposition is true is not enough for that proposition to be included in your evidence, those who defend AA say that it is harder than an externalist who accepts NIKSE and AA thinks to acquire evidence. This is precisely the sort of thing that we should expect would lead to skeptical difficulties.

The careful response is this. We can show that AA faces skeptical worries without assuming ET, so the right response is to reject AA rather than ET. Suppose ET is false. Suppose that your evidence includes the proposition that you have hands. It follows from this assumption and AA that:

- (12) You should be able to know from the armchair that your evidence includes the proposition that you have hands.

Since knowledge entails belief, if you do know what AA says you can know, it follows that:

- (13) You believe that your evidence includes the proposition that you have hands.

It seems that if you are minimally rational and reflective, if you believe that your evidence includes p , you believe that p is the case. So, if you are rational and reflective:

- (14) You believe that you have hands.

It seems that if you are minimally rational and reflective, you will not both believe p and believe yourself not to know that p . So, having given the matter reflection you accede that you believe yourself to know that you have hands. So:

- (15) You believe that you have hands and that you know you have hands.

²⁴ Thanks to an anonymous referee for pressing me on this point.

²⁵ An anonymous referee raised this worry. Note that if you say that we cannot know from the armchair whether ET is true, this would undermine the initial objection to $E = K$. To show that absurd epistemological consequences follow from $E = K$, it was assumed that if $E = K$ is true it is known from the armchair. $E = K$ cannot be known from the armchair unless ET can be known from the armchair as well.

Assume also that you believe the assumption crucial to the argument against evidential externalism:

(16) You believe that you cannot know that you have hands from the armchair.

Now our question is this. How can you square (15) and (16) with the additional claim that you should be able to know just on the basis of what you know from the armchair that your evidence includes the proposition that you have hands? It seems as if you acknowledge that you are committed to the truth of a proposition in virtue of what you take yourself to know from the armchair alone. Yet, you know that the truth of that proposition is not something that can be known from the armchair alone. If we agree that you should not believe what you believe yourself not to know, it seems you should not believe that you have hands given just what you know from the armchair. Thus, you should not believe that the proposition that you have hands is included in your evidence. However, if you could know non-inferentially that you had hands, it follows from this and NIKSE that you should be able to know from the armchair that your evidence includes the proposition that you have hands. So, what are we to say? It seems that problems for AA arise given just NIKSE as an assumption. Rejecting ET will not make the problems for AA go away.

Someone might say that those who accept AA are like to either deny NIKSE or deny that a proposition that cannot be known from the armchair (e.g., that I have hands) are included in my evidence even when known non-inferentially. Let me say this in response. If you deny that your evidence can include the proposition that you have hands you have to choose between accepting the unattractive skeptical view that says that non-inferential knowledge of such a proposition is unattainable and denying NIKSE. Thus, it really does look as if acceptance of AA comes with a cost. It forces you to deny NIKSE in order to deny skepticism. This is bad. We want to avoid the skeptical conclusion that states that you cannot have non-inferential knowledge of the external world, and the suggestion that the best way to avoid this is to say that we have *less* evidence than NIKSE suggests is exceptionally odd. Moreover, to deny NIKSE is to insist that all evidence is provided by introspection or reflection and deny that perception is a basic or autonomous source of evidence that provides us with additional evidence or reasons not provided by introspection. The problem with saying this is that it seems that a basic source of knowledge is a basic source of knowledge only if it is a basic source of evidence or reasons for belief. If acceptance of AA forces us to deny that perception provides us with reasons or evidence, it is hard to see how we might then say that perception is a source of knowledge. To concede that perception is not a source of knowledge is to concede too much to the skeptic.

4 Externalism and epistemic rationality

In the previous Sect. 2 argued that we should reject AA on the grounds that it leads to a troubling skeptical conclusion; namely, that we cannot have non-inferential knowledge of the external world. In this section, I shall explain how an externalist about evidence can accommodate our intuitions about epistemic rationality and undermine the argument from R-INT to E-INT.

Suppose White and Plum are in precisely the same non-factive mental states. White believes that Mustard is the killer because he saw Mustard clobber the victim with a candlestick. Plum likewise believes that Mustard is the killer, but as the result of having undergone a series of hallucinations indistinguishable from White's experiences. It follows from $E = K$ or the conjunction of ET and NIKSE that:

- (1) White's evidence includes propositions that are not part of Plum's evidence. It seems, however, that:
 - (2) Plum is not less than fully rational for believing what White does.
 - (3) Plum and White are equally reasonable in believing Mustard to be guilty.

If epistemic rationality is primarily a matter of a subject respecting her evidence, we have a problem. If externalism is true, Plum's evidence consists of a smaller set of propositions than her evidence would if she had been in the good case. If Plum *knew* that her evidence consisted of just this smaller set of propositions but nevertheless believed what someone does in the good case, Plum would surely be less than fully reasonable. So, externalists can only say that (2) and (3) if they say that Plum does know what her evidence is. And herein lies the rub. Can externalists say that White and Plum are equally rational if epistemic rationality is a matter of 'respecting the evidence' when only White knows what his evidence is?

The internalist reasons that since White and Plum are equally reasonable in believing what they do, the subject in the bad case might be doing worse epistemically than the subject in the good case in some sense, but the subject in the bad case does just as good a job respecting her evidence as her counterpart in a good case. But, if she respects the evidence just as well as the subject in the good case, she cannot suffer from a kind of ignorance as to what her evidence when the subject in the good case will know what evidence he has. Plum would suffer from this ignorance, however, if externalism were true. Thus, it seems E-INT and AA must be true.

Williamson has said two things in response to the argument for E-INT. First, he acknowledges that we might not know what epistemic rationality requires of us when we do not know what our evidence is, but he is quick to point out that the demands of rationality are not the sort of thing we should expect to be luminous.²⁶ Second, he has suggested that when we compare our good and bad cases, we can say that the subjects in these cases have different bodies evidence while insisting that both have sufficient evidence to justifiably believe that things are in their surroundings much as they appear to be.²⁷ Let us consider these two responses separately to see if either is of any help to the externalist.

Even those who are sympathetic to Williamson's suggestion that rationality is not luminous will likely balk at the suggestion that Plum does not know what rationality requires of her simply because she is hallucinating. Her experiences were indistinguishable from White's experiences. The failure to detect the difference between hallucinatory and veridical experience is hardly a failure of reason or a failure that can be attributed to her reasoning. Internalists and externalists share the intuition that Plum is rational in her beliefs. Unless someone can explain how Plum could be rational

²⁶ Williamson (2000b, p. 624).

²⁷ Williamson (2000a, p. 197).

to believe p under the very same circumstances in which she believes what she does because she fails to know what rationality requires of her, it seems we should move on to consider the second response.

Should externalists say that White and Plum are rational or reasonable to believe what they do because both have sufficient evidence to justify their beliefs? The problem with saying that is that it is hard to see what role the ‘external’ evidence (i.e., the pieces of evidence that White has that Plum lacks) plays in explaining how White’s beliefs are justified if we insist that the ‘internal’ evidence (i.e., the evidence they share in common) is sufficient to justify both of their beliefs. I do not want to deny that overdetermination is possible. It clearly is. The problem is that it seems that the externalist who thinks that White and Plum must offer a kind of ‘disjunctive’ explanation as to why Plum and White are both rational in believing what they do if they insist that the external evidence is not epistemically idle and some insist that the very same facts should figure in the explanation as to why White and Plum are rational in believing what they do since they are rational to the same degree in believing what they do.

It seems that we regard Plum as rational in her beliefs given what the proximate cause of her belief formation was (i.e., the appearance of Mustard clubbing someone with a candlestick). It makes sense that she forms the beliefs she does given her aim of believing the truth and there is nothing available to her that we could cite to say that she should have expected that in forming the beliefs she does, she would fail to believe what she should. It also seems that the same proximate cause triggered White’s belief. We know that such a cause ensures that someone is rational in believing what these two subjects believe, and to say that the explanation as to why White is rational involves considerations beyond those that we use to explain why Plum is rational violates a kind of proportionality constraint. In general, if there is some fact that is partially constituted by some second fact (e.g., that she knows p is partially constituted by her believing p ; that she perceives that p is partially constituted by her seeming to see that p) and a certain effect would have been produced even if the second fact obtained in the absence of the first, it seems that what explains the effect is the second fact rather than the first. Applied to the case at hand, White veridically perceives something and Plum seems to see the very same things. The veridical experience and the seeming are sufficient for producing the same belief and triggering the intuition that the belief is rational. If we apply our explanatory assumption, it seems that the same aspects of the conscious mental events and the subjects’ mental states (i.e., that it seems to White that such and such and that it seems to Plum that such and such) explain the belief and how the belief attains the epistemic status that it does.²⁸

²⁸ This objection is due to Wedgwood (2002). The objection, if sound, seems to force Williamson to choose between saying (a) that both White and Plum are justified in their beliefs and that this fact is explained in virtue of the evidence they share in common or (b) that the evidence that White has and Plum lacks is necessary for having a justified belief. It seems to rule out the position Williamson seems to favor on which both White and Plum are justified in their beliefs and the explanation as to why this is need not be reduced to some common justifying factor. You can save justificatory externalism by opting for (a) and denying that the common pieces of evidence had by White and Plum are sufficient for justifying their beliefs, thus reconciling Wedgwood’s explanatory principle with Williamson’s account of justification, but that seems to

It seems that the evidential externalist must say either that only White knows what rationality requires of her or that both White and Plum know what rationality requires of them. If they opt for the former view, it seems they can consistently maintain that epistemic rationality is a matter of respecting the evidence, but they have to deny R-INT and explain away the intuitions that support it. If they opt for the latter view, it seems they can accommodate intuition, but only at the expense of the idea that epistemic rationality is a matter of respecting the evidence. If they opt for both the view that epistemic rationality is a matter of respecting the evidence, insist that both Plum and White know what their evidence is, and say that their evidence is the same, they have just given up on externalism. Since this leads us back to armchair access and E-INT, let us look elsewhere.

Of the two options available to the externalist, I think the externalist ought to embrace the second and say that subjects can know what it is epistemically rational to believe regardless of whether they happen to be in a good case or bad. The externalist does not have to deny that epistemic rationality is a matter of respecting the evidence; rather, they have to give an account of what it is to respect the evidence that allows us to say that someone like Plum might fail to know what the evidence is and what the evidence requires of her because of some non-culpable mistake due to something like hallucination, faulty memory, unreliable testimony, etc. . . . without failing to respect the evidence.

It will be useful to introduce an externalist account of justified belief to explain the response I am offering on behalf of the externalist account of evidence.²⁹ Whereas some evidential externalists use the concept of knowledge to explicate the concept of evidence, those who defend the knowledge account of justification use the concept of knowledge to explicate the concept of justified or permissibly held belief:

(K) Your belief that p is justified/permissibly held iff you know p .³⁰

On this view, if you do not know something is true, there is a conclusive reason to refrain from believing that it is true. This view might seem to be inconsistent with our intuitions regarding White and Plum. Not every failure to know is a rational failure, after all. Can we do justice to this point while insisting that there is a conclusive reason to refrain from believing any proposition you would not know to be true?

Externalists about justified belief cannot do justice to the intuitions that support R-INT if they say that R1 characterizes the relationship between reasons and rationality:

(R1) It is rational/reasonable for S to Φ iff there is overall reason for her to Φ .

Anyone who holds such a view has to think of reasons as the sorts of things that a subject could not fail to identify unless the subject is somehow less than fully rational. If reasons were the sorts of things we could fail to identify without thereby being less than fully rational, we would not have to regard the failure to conform to the demands

Footnote 28 continued

force the externalist to deny a commonly held intuition about justification. Rather than explain the intuition away, Littlejohn (2009) argues externalism is not incompatible with ordinary intuition.

²⁹ Someone could be an externalist about evidence and reject this view. Someone could accept this view of justified belief and reject externalism about evidence.

³⁰ Sutton (2005) defends such a view.

the reasons made as nevertheless reasonable or rational. If the externalists about evidence and justified belief modify this account of the relationship between reasons and rationality, this would help them resist the arguments offered by those who accept R-INT and R1 to ‘internalize’ evidence and the conditions that determine whether our beliefs are justified:

(R2) It is rational/reasonable for S to Φ iff it is not the case that S should have expected there to be undefeated reasons against her Φ -ing.³¹

If R2 does a better job representing the relationship between reasons and rationality than R1, the externalist about justified belief can say that the reason that White and Plum are rational in believing what they do is that it is that neither of them should have expected there to be undefeated reasons to refrain from believing what they do. There is no such reason for White because she saw him do it. There is no such reason for Plum, for while the knowledge account says that her hallucinating constitutes a reason for her to refrain from believing, it is not the case that she should have expected there to be a reason for her to revise her beliefs owing to her hallucination.

By replacing R1 with R2, the evidential externalist can say in response to the argument from R-INT to E-INT that given what their evidence was White and Plum’s beliefs should have differed, there were reasons for Plum to refrain from believing what White believes, but nevertheless Plum was perfectly rational in believing what White believed. This is because Plum was in no position to judge that her evidence was the kind of evidence someone has in the bad case. Thus, Plum is not guilty of any irrationality because Plum should not have expected that there was a reason for her not to believe what White believes. (We’re assuming that there is a reason not to believe what White believes if you have only the evidence someone has in the bad case.) This judgment that there is a reason for Plum to refrain from believing depends upon empirical information inaccessible to her. Externalists can block the internalist argument by distinguishing between judgments about rationality and judgments about reasons. Rationality is primarily concerned with responding to what is presented as a reason in (roughly) the sort of way someone objectively ought to if the reasons are as they are presented as being. Rationality is not primarily concerned with identifying and discovering the reasons that are out there.

The evidential externalist can thus undercut the argument for armchair access. Yes, the externalists will say, epistemic rationality is a matter of respecting the evidence. Yes, they will add, situations will arise in which a person will non-culpably fail to know what her evidence is and what the epistemic reasons demand of her. These claims are incompatible only on the assumption that respecting the reasons involves, *inter alia*, two kinds of knowledge: knowing what reasons there are and knowing how to respond to them correctly. However, we should not say that someone who (non-culpably) failed to identify what reasons there were or what they demanded of her thereby must not have shown proper respect for the reasons that applied to her when she formed her beliefs. Someone who (non-culpably) failed to identify a certain reason has not thereby

³¹ This view seems to be more along the lines of the views defended by Broome (1999), Dancy (2000), Raz (1999), and Parfit (1997).

shown that she is the sort of subject who has no respect for that sort of reason or does not know what that sort of reason demands.

Externalists about evidence, justification, reasons and the like are likely to be partial to the view that aspects of the world have normative significance and so partial to a view like R2 that allows us to distinguish the conditions that determine what is rational for someone to believe from the conditions that determine what reasons there are for them to believe or refrain from believing.³² That some belief misrepresents the situation, for example, is something of normative significance. It means that the belief should be excluded from consideration for the purposes of theoretical deliberation. Judgments about rationality are judgments about the subject's exercise of her capacity for responding to reasons. On R2, the correctness of such judgments does not depend upon whether the subject correctly identifies the reasons that make demands of her. Not every failure to conform to the demands of reasons is a failure of rationality in the strict sense. The identification of those external states of affairs that have normative significance can depend upon memory, perception, and testimony and if these fail the subject, the subject might fail to do what the reasons require without being guilty of any irrationality since she has not failed in her capacity as someone who responds to the reasons as they are presented. To make this concrete, the failure to pick up on the visual clues that someone is confabulating, for example, may well lead to a failure to conform to a certain kind of reason (e.g., the reason to refrain from believing lies). The subject who is taken in by a lie can be taken in without irrational because had the subject's perception of the situation been correct, her response would have been consistent with what the reasons required.³³

There are two reasons to think that R2 is the better view about the relationship between reasons and rationality. While neither is decisive, it should help the reader appreciate the attraction of R2 and see the weakness in the argument from R-INT to E-INT. First, think about some morally difficult situation about which reasonable people can disagree. We should be able to imagine that two subjects, White and Plum, consider the situation carefully, White judges that one ought to Φ rather than Ψ , Plum judges that one ought to Ψ rather than Φ , and *we* judge that neither is unreasonable in their judgment. It seems perfectly consistent with these judgments that we add that there is no more reason to Φ rather than Ψ or stipulate that there is more reason to Φ rather than Ψ . Either way, we sever the connection between what is reasonable or rational to judge and decide to do and what there is overall reason to judge and decide to do. If we accepted R1, whatever made it true that, say, there was more reason to Φ rather than Ψ would compel us to say that someone who judged otherwise and decided to act otherwise was less than rational or reasonable. But, that is to deny that we are considering a case of reasonable disagreement. And, to say that in such cases there

³² For a discussion of this sort of view, see Raz (1999).

³³ Raz explains this point in his remark that, "... to be rational in the first sense, that is, to be people with the ability to perceive reasons and respond to them, we need a range of capacities which do not directly contribute to our rationality. They include some perceptual ability, and the capacity to control our movements at will. An impairment of our perceptual ability does not diminish our rationality. Nor does lack of muscular control, or other neurological or physical impairments of our ability to move at will. Possession of at least some perceptual ability, and of some ability to control one's movements at will, are presupposed by capacity-rationality. But they are not themselves constituents of rationality" (Raz 1999, p. 68).

can be no more reason to Φ rather than Ψ would also force us to deny that White and Plum are rational since both thought that, morally speaking, one course of action was better than the other.

A second mark against R1 is that it prevents us from making sense of certain familiar distinctions between ways of addressing the charge that someone should be blamed for wrongdoing. Strawson helpfully distinguished between three ways of showing that someone's blaming another for wrongdoing is out of order.³⁴ We can offer a justification on behalf of the subject to show that there was reason to do what was *prima facie* wrong. To do this, we show that there was overall reason to do the action the agent in fact did. We can show that the subject in question is not the sort of subject that can be held responsible (e.g., the subject has gone temporarily insane, the subject was drugged or hypnotized, etc...). To do this, we offer an exemption. Third, there are the cases in which we try to show that subjects that can be held accountable should not be blamed for some particular wrong. In order to remove blame in these cases, we do not have to show that the agent's action was the action there was overall reason to do. We do, however, have to work from the assumption that the agent can be held responsible and had a capacity for responding to reasons. To remove blame, we have to tell a story about how a competent agent with the right sorts of concerns could nevertheless fail to do what there is overall reason to do. In these cases of excusable wrongdoing, it seems that we cannot explain why an excuse is in order unless we can show that the subject was rational in deciding what she did even though she failed to do what there was overall reason to do.

To make this somewhat concrete, think about three shootings. Somewhere in between cases of self-defense and cases in which someone shoots another because, say, they have gone temporarily insane is the case in which the subject shoots at a jogger believed to be a mugger. If we decide that they are non-culpably mistaken in their belief that the person approaching them is a dangerous mugger carrying a club rather than a jogger carrying a flashlight and acknowledge that had this belief been true, it would not have been wrongful to fire on the person approaching, it seems we have told a story in which the subject is no less reasonable than the subject whose actions are justified on grounds of self-defense. It also seems that we have told a story in which there was not overall reason for the person to act as they did. This is evidenced by the fact that certain kinds of violent intervention are justified in the case of 'imperfect' self-defense, but not in the case of proper self-defense.³⁵

³⁴ Strawson (1974).

³⁵ I owe the term 'imperfect' self-defense to Moore (1993) who says that such cases can include ones in which the subject correctly identifies what would be justified on grounds of self-defense, acts according to such a belief, but performs an action that is not justified on grounds of self-defense because the subject not-culpably, but mistakenly, believes that someone is threatening them. Now, I am well aware of the fact that some will want to say that in cases of imperfect self-defense, the subject's actions were justified on grounds of self-defense since she non-culpably, but mistakenly, believed that she was about to be attacked. As I am merely trying to show that a certain picture of the relationship between reasons and rationality is not wholly unmotivated and not obviously incoherent so as to make sense of an externalist view about evidence that respects the general idea that rationality is a matter of respecting the evidence, this is not the time to argue that my description of the case of imperfect self-defense as an excusable wrong is the correct one. Darley and Robinson (1998) show that this is how lay people describe it even if some philosophers do not. Moore (1993) and Robinson (1996) defend this description on conceptual grounds.

In order to make sense of the distinction between justifications, exemptions, and excuses, it seems we needed to make use of the sort of distinction between what there is reason to do and what rationality requires. Because this gives us some reason to reject R1 in favor of something along the lines of R2, we can see that there is something going for the distinction between reasons and rationality. If the externalist can use this distinction to make sense of the idea that someone respects the reasons and respects the evidence even if they fail to know what those reasons are or what the evidence happens to be, then the externalist can say that they do not need armchair access to accommodate these observations about epistemic rationality. What it is to be rational in believing p is (roughly) form one's beliefs in such a way that it is not to be expected that there is some undefeated reason against believing p . If one knows or should know that one has no evidence or weak evidence, one should not expect that one does not believe against an undefeated reason not to believe. If one fails to determine what the evidence truly because perception, memory, or testimony fails to indicate what the reasons truly are, this is not a failure of rationality *per se*. Not if the subject responds in the way she should have had the evidence been what it seemed to her to have been at the time.

While I do not think that the reply I have offered on behalf of the externalist is the last word on the subject, hopefully it shows that more has to be said to motivate the charge that the externalist cannot accommodate the platitude that epistemic rationality is a matter of respecting the evidence. Specifically, we need to know what would be wrong with a view on which claims about epistemic rationality depend upon how things are internal to the subject even if claims about what evidence there is and what reasons there are for belief depend (in part) upon how things stand in subject's external surroundings. Without some reason to think such a view is incoherent, the argument from armchair access has no force. Not only is the principle of armchair access something we are better off without, our intuitions about epistemic rationality support that principle since those intuitions can be accommodated by an externalist who incorporates ET and NIKSE into her account of evidence.

Acknowledgements I want to thank Ali Hasan and two anonymous referees for their extensive comments. I also want to thank Mike Almeida, Andrew Bailey, Eric Barnes, Richard Chappell, Andrew Cullison, Leo Iacono, Jonathan Ichikawa, Robert Howell, Errol Lord, Nico Silins, and John Turri for discussion.

References

- Achinstein, P. (1984). Concepts of evidence. In P. Achinstein (Ed.), *The concept of evidence* (pp. 145–175). New York: Oxford University Press.
- Audi, R. (1993). *The structure of justification*. New York: Cambridge University Press.
- Bird, A. (2004). Is evidence non-inferential? *The Philosophical Quarterly*, 54, 252–265.
- Boghossian, P. (2008). Content and self-knowledge. In his *Content and justification* (pp. 139–158). New York: Oxford University Press.
- Brewer, B. (1999). *Perception and reason*. New York: Oxford University Press.
- Broome, J. (1999). Normative requirements. *Ratio*, 12, 398–419.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46, 279–295.
- Conee, E., & Feldman, R. (2004). *Evidentialism*. New York: Oxford University Press.
- Dancy, J. (2000). *Practical reality*. New York: Oxford University Press.

- Darley, J., & Robinson, P. (1998). Testing competing theories of justification. *North Carolina Law Review*, 76, 1095–1143.
- Darwall, S. (1985). *Impartial reason*. Ithaca, NY: Cornell University Press.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104, 1–52.
- Feldman, R. (1988). Subjective and objective justification in ethics and epistemology. *The Monist*, 71, 405–419.
- Ginet, C. (1975). *Knowledge, perception, and memory*. Dordrecht: Reidel.
- Goldman, A. (2002). *Pathways to knowledge: Private and public*. New York: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105, 571–590.
- Hyman, J. (1999). How knowledge works. *The Philosophical Quarterly*, 49, 433–451.
- Hyman, J. (2006). Knowledge and evidence. *Mind*, 115, 891–916.
- Kelly, T. (2008). Evidence: Fundamental concepts and the phenomenal conception. *Philosophy Compass*, 3, 933–955.
- Littlejohn, C. (2009). The externalist's demon. *Canadian Journal of Philosophy*, 39, 399–435.
- Maher, P. (1996). Subjective and objective confirmation. *Philosophy of Science*, 63, 149–174.
- McKinsey, M. (1991). Anti-individualism and privileged access. *Analysis*, 51, 9–16.
- Moore, M. (1993). *Placing blame*. New York: Oxford University Press.
- Neta, R. (2008). What evidence do you have? *British Journal for the Philosophy of Science*, 59, 89–119.
- Neta, R., & Pritchard, D. (2007). McDowell and the new evil genius. *Philosophy and Phenomenological Research*, 74, 381–396.
- Parfit, D. (1997). Reasons and motivation. *The Aristotelian Society, Supplementary*, 77, 99–130.
- Raz, J. (1999). *Engaging reason*. New York: Oxford University Press.
- Robinson, P. (1996). Competing theories of justification: Deeds versus reasons. In A. Simester & A. Smith (Eds.), *Harm and culpability* (pp. 45–70). New York: Oxford University Press.
- Silins, N. (2005). Evidence and deception. *Philosophical Perspectives*, 19, 376–404.
- Strawson, P. (1974). Freedom and Resentment. In his *Freedom and Resentment* (pp. 1–25). Methuen: London.
- Sutton, J. (2005). Stick to what you know. *Nous*, 39, 359–396.
- Turri, J. (2009). The ontology of epistemic reasons. *Nous*, 43, 490–512.
- Unger, P. (1975). *Ignorance*. New York: Oxford University Press.
- Wedgwood, R. (2002). Internalism explained. *Philosophy and Phenomenological Research*, 65, 349–369.
- Williamson, T. (2000a). *Knowledge and its limits*. New York: Oxford University Press.
- Williamson, T. (2000b). Skepticism and evidence. *Philosophy and Phenomenological Research*, 60, 613–628.