Boise State University ScholarWorks

Philosophy Faculty Publications and Presentations

Department of Philosophy

5-18-2012

Two Ways to Put Knowledge First

Alexander Jackson Boise State University

This is an electronic version of an article published in *Australasian Journal of Philosophy*, 90(2). *Australasian Journal of Philosophy* is available online at: www.tandfonline.com. DOI: 10.1080/00048402.2011.587438

Two Ways to Put Knowledge First

Alexander Jackson

Boise State University

This paper distinguishes two ways to 'put knowledge first'. One way affirms a knowledge norm. For example, Williamson [2000] argues that one must only assert that which one knows. Hawthorne and Stanley [2008] argue that one must only treat as a reason for action that which one knows. Another way to put knowledge first affirms a determination thesis. For example, Williamson [2000] argues that what one knows determines what one is justified in believing. Hawthorne and Stanley [2008] argue that what one knows determines what it is rational for one to do. This paper argues that the defender of the knowledge norms can and should reject the determination theses. For example, the rationality of acting on a partial belief cannot be explained in terms of the subject's knowledge. That's no problem for the knowledge norm, which only governs acting on full belief. (Analogously, the knowledge norm of flat-out assertion does not govern hedged assertion.) One might worry that rejecting the determination theses undermines the importance of knowing. I reply that the knowledge norms set the standard for epistemic success. The importance of success is not undermined by loosening its ties to justification and rationality.

1. Introduction

At the end of the twentieth century, mainstream epistemological inquiry focused on justified belief, rather than knowledge. Justification seemed to be the primary dimension of epistemic evaluation. Influential recent work, with Timothy Williamson's *Knowledge and Its Limits* [2000] at its heart, reacts against that trend. This school 'puts knowledge first'. Various theses have been bundled together under that slogan. This paper cleaves two of them apart.¹ One way to put knowledge first holds that a *knowledge norm* governs assertion, belief, etc. Another way affirms a *determination thesis*: one's knowledge determines what one is justified in believing, and what it is rational for one to do. Both of these ways of putting knowledge first are meant to secure the primary importance of knowing. This paper argues that the defender of the knowledge norms can and should reject the determination theses. That is the best way to put knowledge first.

Let's quickly review some of the knowledge norms defended in the literature. Timothy Williamson argues that Knowledge is the Norm of Assertion (KNA) [2000: 243].

(KNA) One must: assert that p only if one knows that p.

Amongst other arguments, Williamson claims that KNA provides the best explanation for the following three phenomena. Firstly, it is natural to challenge an assertion by asking, 'How do you know?' KNA would explain why, as one would thereby be questioning whether the assertion meets the necessary epistemic standard. Secondly, you can't assert your ticket won't win the lottery, and intuitively that's because you don't know it won't. Thirdly, one can't assert: 'Though I don't know it, the keys are in the kitchen.' KNA would explain why, as one would thereby announce that one shouldn't assert that the keys are in the kitchen, and then assert it anyway.

¹ Other ways of putting knowledge first include holding that knowledge is unanalysable [Williamson 2000 chapters 1–3], and that all evidence is knowledge (which I endorse in section 6).

Williamson suggests that occurrent belief is the inner analogue of assertion. (As Ernest Sosa says [2010: 172, 189], occurrent belief is 'assent' to a proposition.) So he thinks it is plausible that Knowledge is the Norm of Belief (KNB) [Williamson 2000: 255–6].

(KNB) One should: believe p only if one knows p.

Another knowledge norm governs the grounds of action. John Hawthorne [2004: 30] endorses the Practical Reasoning Constraint: 'one ought only to use that which one knows as a premise in one's deliberations'. Jason Stanley [2005: 9] holds that, 'one should act only on what one knows'. Hawthorne and Stanley [2008: 577] defend the Action-Knowledge Principle (AKP).

(AKP) Treat the proposition that p as a reason for acting only if you know that p.

In many cases we criticize someone for acting on the basis of something they don't know. Hawthorne and Stanley [2008: 571–4] give the following example, among others. Suppose a sous-chef is told to take the cake out of the oven when it is done. We criticize the sous-chef if he takes the cake out without knowing that it is done. Doing so is incompetent (though not necessarily irrational). That we criticize the sous-chef for acting without knowing suggests that AKP is true.

We can further motivate AKP as follows [see Hawthorne 2004: 29–31, 147]. Suppose Bill reasons: 'My lottery ticket won't win; so it is worthless; so I should sell it even for one cent.' His conclusion is unacceptable, yet every step is valid. Intuitively, the reasoning is bad because Bill does not know his ticket won't win. If he did know it (say because he knew the lottery is fixed), the reasoning would be good. This suggests that knowing is the epistemic standard one should meet before treating something as a reason for action.

I am sympathetic to the knowledge norms (particularly KNB). I won't argue for those principles here; I think Williamson, Hawthorne, and Stanley do a good job of doing so. I will assume throughout this paper that the knowledge norms are correct.²

Now let's review the determination theses in the literature. Williamson argues that 'only knowledge justifies belief' [2000: 185]. That is, any justified belief is justified because it is supported by what one knows, or itself constitutes knowledge. Hence:

(WILLIAMSON'S DETERMINATION THESIS) What one knows determines what one is justified in believing.

Williamson explicitly endorses a determination thesis, as a way of putting knowledge first. Hawthorne and Stanley [2008] are not explicit. But I will argue that their response to Stephen Schiffer [2007] only makes sense if they endorse the following claim.

(HAWTHORNE AND STANLEY'S DETERMINATION THESIS) What one knows determines what it is rational for one to do (possibly in concert with one's desires).

² The above knowledge norms specify a necessary condition for epistemic propriety. As Jessica Brown [2008] observes, that's not the same thing as knowledge being necessary and sufficient. I endorse the latter interpretation of the knowledge norms: knowledge is the epistemic standard to meet. AKP* is a biconditional variant of AKP (compare Hawthorne and Stanley's RKP [2008: 578]).

⁽AKP*) Take p to be available as a reason for acting iff you know that p.

I suspect Brown [2008: 174–182] is right that in certain high-stakes situations, one should not treat as a reason something one knows and properly believes. If so, AKP* is not always true. Such cases put no pressure on the claim that knowledge is sufficient for proper belief, i.e. a biconditional variant of KNB.

Hawthorne and Stanley are not concerned with whether one's desires partly determine what it is rational for one to do. Their idea is that it is what one knows, *as opposed to one's other full beliefs or partial beliefs*, which determines what it is rational for one to do (possibly in concert with one's desires). I'll leave the possible role of desires implicit in future glosses on their determination thesis.

This paper argues for two main claims. Firstly, the determination theses have very implausible consequences. Secondly, if the knowledge norms are right, there is no need for a second way of securing the importance of knowing. The knowledge norms do so by themselves. In summary, the knowledge-lover should accept the knowledge norms, but deny that what one knows determines what one is justified in believing and what it is rational for one to do.

I'll start off by considering Hawthorne and Stanley's response [2008] to Stephen Schiffer's challenge to AKP [2007]. Prima facie, we often act rationally on the basis of *partial* beliefs rather than *full* beliefs. Schiffer thinks that such cases pose a deep problem for AKP, as I explain in section 2. Hawthorne and Stanley respond that what appears to be one's acting on a partial belief that p is really one's acting on knowledge that it is likely that p. I argue against that account in section 3. But no such claim is required to defend AKP. Hawthorne and Stanley give an implausible account of acting on a partial belief, not because they are defending AKP, but because they are trying to defend their determination thesis. In my view, the defender of AKP should admit that there are many cases of acting on a partial belief in which the rationality of so acting is not grounded in the subject's knowledge. AKP governs acting on full belief, and doesn't govern acting on partial belief. Section 4 discusses analogous concerns about whether hedged assertions are governed by KNA, arguing that they aren't.

Section 5 defends the resulting view from the accusation that it undermines the importance of knowing. In my view, the knowledge norms set the standard for epistemic success, not for epistemic rationality or justification. The importance of success does not need to be shored up. I note that it is not clear that AKP thus interpreted conflicts with the subjective Bayesian account of rationality, contra Schiffer and Hawthorne and Stanley.

Section 6 examines Williamson's determination thesis. I point out a counter-intuitive consequence. Williamson's main motivation for the view is to show that knowledge is important. Section 5 already explained that given the knowledge norms, there is no need for another way of securing the importance of knowing. I briefly survey some of Williamson's other arguments that only knowledge justifies belief, finding them thin. I conclude that we should reject Williamson's claim.

2. Schiffer's Objection to AKP

This section describes Schiffer's [2007] objection to AKP, and introduces two kinds of reply. Section 3 argues against the kind of reply given by Hawthorne and Stanley [2008]. Sections 4 and 5 defend the other kind of response to Schiffer.

Schiffer responds to Stanley's claim that, 'one should act only on what one knows' [Stanley 2005: 9]. Hawthorne and Stanley [2008] refine this claim into AKP. Let's see whether Schiffer's objection [2007: 189–190] hits AKP. He claims that very often, we act on *partial* beliefs. Having a partial belief is a matter of having a certain degree of belief, also called one's credence in the proposition. Suppose that Jane believes to degree 0.4 that it will rain tonight. Schiffer observes that Jane is 'completely justified in carrying an umbrella even though [she doesn't] know that it will rain' [2007: 189]. Schiffer complains [189],

Not only is this a prima facie counterexample to [the knowledge norm], insofar as it appears to be an example in which one is justified in acting as one does on the basis not of knowledge but of a justified partial belief, but there is a familiar and widely-accepted Bayesian account of why one's action is justified, even though one isn't acting on knowledge: one is performing that action which has the greatest expected desirability. There are two objections here. The first is that the case of Jane is a counterexample to the knowledge norm. The second is that AKP does not explain why it is rational³ for Jane to carry an umbrella, but the rival Bayesian account does. Both objections need unpacking.

The first objection is weak. Maybe Jane *acts on* her partial belief that it will rain. But she does not *act on the basis that* it will rain. So there is no counterexample to the principle that one should act on p only if one knows that p. That's the charitable interpretation of Stanley's claim that 'one should act only on what one knows'. The point is clear once we consider whether the case is a counterexample to AKP. Roughly, to treat p as a reason for action is to have a *full* belief that p. Jane's reason for carrying an umbrella is not that *it will rain*. It may well be that if Jane did take that proposition to be available as a reason for action, i.e. if she believed that it will rain, she would decide not to leave the house, rather than go out but take an umbrella. [Hawthorne and Stanley 2008: 582.] AKP gets things exactly right: Jane does not know that it will rain, and should not take it as a reason for action that it will rain.

Consider the analogous objection to the Knowledge Norm of Belief. It is true that Jane's attitude is epistemically proper, and it does not constitute knowledge that it will rain. But that's no counterexample to KNB, because Jane does not fully believe that it will rain; she has a merely partial belief that it will. The knowledge norm governs full belief and not partial belief. By fully believing, one takes oneself to know; by having a merely partial belief, one takes oneself *not* to know. KNB gets things exactly right: Jane should not fully believe that it will rain.

Schiffer's second worry is the interesting one. It is rational for Jane to carry an umbrella. Schiffer gestures at the subjective Bayesian explanation of that normative standing: given her degrees of belief and her desires, carrying an umbrella maximizes Jane's expected utility. What explanation can one give if one accepts AKP?

There are two kinds of response to this challenge. The first is that AKP *does* explain why it is rational for Jane to carry an umbrella. On this view, Jane knows something that is a good reason for her to act in that way. Hawthorne and Stanley take this position, claiming that Jane knows that *it is fairly likely it will rain*, and that is good reason for her to carry an umbrella. I evaluate that response in section 3, finding it unsatisfactory. The second kind of response to Schiffer's challenge says that AKP *does not* explain why it is rational for Jane to carry an umbrella. On this view, rational action is not always a matter of something you know being your reason for action. I will defend this response in section 5, arguing that it does not undermine the central importance of knowing.

3. Hawthorne and Stanley's response to Schiffer

In this section, I consider the response to Schiffer's challenge favoured by Hawthorne and Stanley [2008: 581–5]. On their view, AKP does explain why Jane's action is rational. Hence what seems to be her acting on a partial belief is really her acting on the basis of a full belief that constitutes knowledge.

Let's consider first why Hawthorne and Stanley would take Schiffer's challenge seriously, committing to what seems a very strange psychological claim. The case of Jane is not even putatively a counterexample to AKP. Hawthorne and Stanley do not seem to notice this. They simply assume that the advocate of AKP must use it to explain why it is rational for Jane to carry an umbrella [581]. What claim, stronger than AKP, are they really defending? I suggest that they are trying show that what one knows determines what it is rational for one to do (possibly in concert with one's desires). The case of Jane is a putative counterexample to that thesis. To defend it, one needs to show that Jane's knowledge determines that it is rational for her to carry an umbrella. I can't think what other thesis would commit Hawthorne and Stanley to explaining the rationality of Jane's action in terms of her knowledge. That is, their reply to Schiffer is an attempt to defend their determination thesis, not the knowledge norm AKP.

³ Schiffer says that Jane's action is justified. Bayesianism is usually presented as a theory of rationality, so I will restate Schiffer's objection in those terms.

There is some direct textual evidence that Hawthorne and Stanley intend to defend a determination thesis. They start the paper as follows [2008: 571].

According to one standard picture, \ldots rational action is a matter of maximizing expected utility \ldots . It is subjective degrees of belief that matter for rational action, not knowledge. On this picture, having knowledge that p is independent of whether it is rational to act on one's belief that p \ldots . In contrast to the picture just sketched, our ordinary folk appraisals of the behavior of others suggest that the concept of knowledge is intimately intertwined with the rationality of action.

This passage suggests they think that one's knowledge determines which actions are rational, as opposed to one's degrees of belief doing so.

Whether or not it is Hawthorne and Stanley's actual motivation, the relevant determination thesis commits one to the kind of reply to Schiffer they give. If that kind of reply is unsustainable, then one's knowledge does not determine what it is rational for one to do.

Hawthorne and Stanley have a specific proposal about what knowledge Jane is acting on. But let's get a couple of other proposals out of the way first. Jane's reason for carrying the umbrella is not *that she has degree of belief 0.4 in its raining this afternoon*. That introspective knowledge does not make it rational for Jane to carry an umbrella. For if Jane's degree of belief were the result of irrational pessimism, she would still know she has that degree of belief, but would not act rationally by taking her umbrella. Nor is Jane's reason for carrying the umbrella *that there is an objective likelihood of 0.4 that it will rain*. Suppose the local professional weather-forecasters know that there is no significant objective likelihood of rain, because the clouds will blow away in the next half hour. Jane does not have that information. In that case, Jane does not know that there is a significant objective chance of rain, because there isn't one; yet it would still be rational for her to carry her umbrella. So the proposal fails to explain her being rational.

Rather, Hawthorne and Stanley propose that Jane's reason for carrying the umbrella is that *it is epistemically fairly likely that it will rain* [591–5]. On their view, the epistemic probability of a given proposition for a subject is determined by the subject's total body of knowledge [582]. For example, if one knows that a coin is symmetrical, that determines that the epistemic probability for one of its landing heads when tossed is 0.5.

I will now argue that Hawthorne and Stanley's proposal is not a plausible general account of (what is intuitively) acting on the basis of a partial belief. Suppose Betty and Brian are on Avenue A, and they want to go to the good burger place Betty tried out last week. Is it down 5th street or 6th street? Betty isn't sure, but she has a mnemonic feeling that it is down 5th street. Betty has a rational degree of belief of 0.6 that the burger place is on 5th street. Intuitively, it is rational for Betty to head down 5th street, and irrational for her to head down 6th. I will argue that Hawthorne and Stanley's proposal cannot account for this.^{4, 5}

Hawthorne and Stanley must claim that Betty has knowledge that makes it epistemically probable to degree 0.6 that the burger place is on 5th street. But she has no such knowledge. Betty remembers going to the burger place before, but no more about its location than her mnemonic hunch supplies. The only strategy I can think of is to say that Betty knows that she has a hunch it is on 5th street, and that knowledge makes it likely (possibly in concert with other things she knows⁶) that the burger place is on 5th. That proposal does not get off the ground. Sometimes people introspect, and sometimes they don't. We can stipulate that Betty is not introspecting, so she does not have knowledge that *she has a hunch it is on 5th street*. Of course, such introspective knowledge is easy to come by; but Betty does not form it, and yet acts rationally.

⁴ If I am right, Betty acts rationally but not for a good reason (in the sense governed by AKP). I am not worried by this consequence.

⁵ In Jackson [forthcoming], I explain how apparent memories make beliefs rational without functioning as evidence or reasons for them.

⁶ I doubt the parenthetical can be spelled out attractively.

It does not help to insist that there is a disposition sense in which Betty 'has' that introspective knowledge. We need to explain why Betty acts on a rational basis, not just why she should head down 5th street. (That's analogous to explaining why a belief is doxastically justified, as opposed to propositionally justified.) Even if one says that Betty 'has' the knowledge she could easily acquire by introspecting, one must admit that it is not her actual basis for acting. So it is not the case that *her reason* for heading down 5th street is . Even if one insists that there is a dispositional sense in which Betty 'knows' she has a mnemonic hunch, it won't plausibly explain why Betty acts on a rational basis. So the introspective knowledge is not really what makes it rational for Betty to carry an umbrella.

Betty does not have knowledge that makes it epistemically likely that the burger place is on 5th street. So on the conception of epistemic probability Hawthorne and Stanley appeal to, it is *not true* that it epistemically likely for her that the burger place is on 5th. So Betty can't know it is epistemically likely. So by AKP, she should not take that proposition as a reason for action. So Hawthorne and Stanley haven't shown how AKP explains why it is rational for Betty to (intuitively speaking) act on her partial belief. The case of Betty stands as a counterexample to the thesis that one's knowledge determines what it is rational for one to do.

The only way for Hawthorne and Stanley to get round this problem is to change their conception of epistemic probability, so that it turns out true that it is epistemically likely that the burger place is on 5th street. I will argue that no such alternative conception of epistemic probability is compatible with the view that one's knowledge determines what it is rational for one to do. There is no way here to defend their determination thesis.

On the alternative to be considered, epistemic probability is not determined by the subject's knowledge. The epistemic probability for Betty that the burger place is on 5^{th} street is 0.6, but something other than her knowledge explains why. I worry that any such view will also say that something other than one's knowledge explains why certain full beliefs are rational, and hence why it is rational to act on them. That's incompatible with the view that one's knowledge determines what it is rational for one to do.

Let's work through one alternative view of epistemic probability. Suppose that epistemic probability is determined by one's evidence, but one's evidence does not consist in things one knows. (Denying that all evidence is knowledge is a significant retreat from putting knowledge first, one I reject in section 6.) Let's consider the view that one's evidence consists in appearances one has. Such a view will allow that it is epistemically likely for Betty that the burger place is on 5^{th} street. That's because Betty's weak memory appearance 5^{th} is part of her evidence. Presumably, one can believe rationally on the basis of one's evidence, which on this view is not part of one's knowledge. When such a belief does not amount to knowledge, AKP would not explain why one can act rationally on the belief. In such a case, one's knowledge would not determine what it is rational for one to do. So the current proposal about epistemic probability is incompatible with Hawthorne and Stanley's determination thesis.

In summary, Hawthorne and Stanley are defending the view that one's knowledge determines what it is rational for one to do—hence their insistence that Jane's knowledge determines that it is rational for her to carry an umbrella. They propose that acting on a putative partial belief is really acting on a full belief about epistemic likelihood. That's not plausible given their understanding of epistemic likelihood. Other conceptions of epistemic likelihood lead to other counterexamples to the thesis they are defending. I conclude that we should reject their determination thesis.

4. KNA and Hedged Assertion

This section considers a Schiffer-style worry about the knowledge norm of assertion (KNA). I argue that hedged assertions are not flat-out assertions, and KNA only applies to the latter. Some other epistemic norm governs hedged assertions. If that's right, we should expect full belief to be governed by a knowledge norm (KNB), and some other epistemic norm to govern partial belief. And if that's right, Hawthorne and Stanley are wrong to bring acting on a partial belief under the scope of the knowledge norm AKP. And if that's right, one's knowledge does not determine what it is rational for one to do.

Suppose Betty says to Brian, 'I *think* the burger place is on 5th street', or, 'I *believe* it is on 5th street.' Let's take these as our paradigms of hedged assertion. Betty does not know that the burger place is on 5th street, even if her hunch is correct. That does not make her hedged assertion epistemically inappropriate. So knowing that p is not the standard for hedged assertion that p.

There are two ways to make this compatible with KNA. On the first view, hedged assertion that p is really flat-out assertion of some other proposition q, and the speaker must know q. For example, maybe Betty's hedged assertion is a flat-out assertion that it is epistemically likely that the burger place is on 5th. That's analogous to Hawthorne and Stanley's response to Schiffer. On the second view, hedged assertions are not a species of flat-out assertion. For example, the content of Betty's hedged assertion is *that the burger place* is on 5th, but she does not flat-out assert that content. Hedged assertion is a different speech-act from flat-out assertion, and is governed by a different epistemic norm. That's analogous to the response to Schiffer that I propose.

Let's survey arguments against holding that hedged assertion is flat-out assertion of some related proposition. I'll look at three proposals about what is flat-out asserted. The first proposal is that Betty's hedged assertion is really a flat-out assertion that *she fully believes that the burger place is on* 5^{th} . Some care is needed here. It is plausible that Betty performs two speech-acts by saying 'I believe it is on 5^{th} street': she flat-out asserts that she believes it is on 5^{th} street, and she makes a hedged assertion that it is on 5^{th} street. Plausibly, the two speech acts stand in the same relation as holds between one's literally saying something and one's pragmatically implicating something else. (It is familiar that Gricean implication holds between speech-acts of different kinds. For example, one can implicate the order to shut the window by asking, 'Is it getting cold in here?' or by asserting, 'I'm getting cold.') On such a view, Betty's *literal* speech-act is an assertion that the burger place is on 5^{th} street; that *pragmatically implicates* the hedged assertion that the burger place is on 5^{th} street. Jeremy Fantl and Matthew McGrath suggest a Gricean mechanism by which that might work [2009: 145-8]. The natural view is that KNA governs what the speaker pragmatically implicates, not what she literally asserts but doesn't mean.

Our question is not whether Betty literally flat-out asserts that she fully believes that the burger place is on 5^{th} street. Our question is what explains the epistemic propriety of her (pragmatically implicated) hedged assertion that the burger place is on 5^{th} street. We are examining proposals according to which KNA explains the propriety of the hedged assertion, because the hedged assertion *turns out to be* a flat-out assertion of some other content. The first proposal is that Betty's hedged assertion is really a flat-out assertion that she fully believes that the burger place is on 5^{th} . That seems wrong: Betty's hedged assertion is sincere, but it is obvious to one and all that she does not fully believe that the burger place is on 5^{th} , she would flat-out assert that it is.

More carefully, we can argue that assuming KNA is correct, the current proposal gives the wrong conditions on epistemically proper hedged assertion. Betty doesn't know that she fully believes the burger place is on 5th (because knowledge is factive); yet her hedge assertion is epistemically proper. So knowing she fully believes p is not necessary for Betty to properly make the hedged assertion that p. Nor is such knowledge sufficient. Suppose Brian fully believes that the burger place in on 5th street, but just because he really likes the number 5. Then his hedged assertion, 'I *believe* it is on 5th street,' is epistemically out of order. That's so, even though he knows he fully believes it is on 5th. Knowing that one fully believes p is neither necessary nor sufficient for epistemically proper hedged assertion that p. So if KNA is right, hedged assertion that p is not flat-out assertion that one fully believes that p.

The second proposal is that hedged assertion that p is flat-out assertion that one *partially* believes that p. We can run a similar argument from KNA as we did against the first proposal. Betty's hedged assertion is proper, even if she has not formed the introspective knowledge that she has a partial belief that the burger place is on 5^{th} . Nor is such knowledge sufficient. Suppose Brian partially believes that the burger place is on 5^{th} , but just because he really likes the number 5. Then his hedged assertion, 'I *believe* it is on 5^{th} street,' is epistemically out of order. That's so, even though he knows he partially believes it is on 5^{th} . Knowing that one partially believes p is neither necessary nor sufficient for epistemically proper hedged assertion that p. So if KNA is right, hedged assertion that p is not flat-out assertion that one partially believes that p.

The third proposal is that a hedged assertion that p is a flat-out assertion that it is epistemically likely for one that p. On this proposal, hedged assertions satisfy KNA because one knows what one is flat-out asserting, namely that it is epistemically likely that p. The proposal is analogous to Hawthorne and Stanley's view that acting on a partial belief that p is acting on knowledge that it is epistemically likely that p. The worry I presented in the case of action arises again. Betty does not have knowledge that makes it epistemically likely that the burger place is on 5th street. So (on Hawthorne and Stanley's conception) it is not true that it is epistemically likely; so she can't know it is; so by KNA she should not assert that it is epistemically likely. But she is in an epistemic position to make the hedged assertion, 'I think it's on 5th street.' So a hedged assertion that p is not a flat-out assertion that it is epistemically likely that p. (We might be able to save the proposal by holding that epistemic probability is determined by the subject's evidence, which is taken to consist of appearances they have. It is unattractive for a knowledge-firster to deny that one's evidence consists in things one knows. Further, the aim of this section is to develop an analogy that makes Hawthorne and Stanley's response to Schiffer look unattractive. My opponent in that task must reject the suggested conception of evidence, as it is incompatible with the determination thesis they defend.)

I have argued that hedged assertions are not a species of flat-out assertion, and hence are not subject to KNA. They are subject to a different epistemic norm.⁷ Setting arguments aside, that conclusion has the ring of truth. Betty is well aware that she doesn't know whether the burger place is on 5^{th} street. She just has her hunch. It is no accident there is a kind of speech-act Betty can make to share her epistemic situation with Brian. It seems perverse to try to reduce that kind of speech-act to flat-out assertion. Analogously, we should not try to reduce partial belief to full belief. One can act rationally on a partial belief, so one's knowledge does not determine what it is rational for one to do.

5. Rationality and the Importance of Knowing

I have argued that the best defence of AKP holds it to be silent on the epistemic status of acting on a partial belief. That is, the advocate of the knowledge norm should reject the relevant determination thesis. This section explains why that doesn't undermine the primary significance of knowing. If that's right, there is no need to try to secure the importance of knowing by arguing that only knowledge rationalizes action. The refrain of this section is that knowledge is a different and more fundamental epistemic evaluation than rationality. One consequence is that AKP and the subjective Bayesian account of rationality do not directly conflict, contra Schiffer and Hawthorne and Stanley. I'll consider AKP and KNB in tandem; I suspect that the appeal of AKP derives from that of KNB.⁸ If one treats something as a reason then one occurrently believes it, and so is required by KNB to know it.

Let's be clear that AKP and KNB are not norms of rationality. KNB should not be taken to say that any false belief is irrational because it isn't knowledge. Of course there are false but rational beliefs. AKP should not be taken to say that acting on a false belief is irrational. Of course it is rational to act on a false but rational belief. The following example shows that we distinguish rationality from success, that is, from the standard we aim to meet. I propose we interpret KNB as saying that knowing is the standard for epistemically successful belief.

Suppose the aim of one's trip to the airport is to pick up the visiting speaker, Professor Jones. The norm governing the trip is: I must pick up Professor Jones. The norm governing the trip defines the conditions for its success, not the conditions for acting rationally or justifiedly. What's really important, what one primarily cares about, is whether one succeeds. It is cold comfort, upon realizing that one has driven an impostor into the city leaving the professor stranded, that one acted rationally in so doing. Whether one acted rationally is of limited interest, and certainly is not the only important thing.

⁷ What standard do hedged assertions aim to meet? Maybe the *epistemic* standard is something like justification, but they also aim to be true.

⁸ I suggest in footnote 2 that the biconditional variant of AKP is false, for reasons that assume KNB.

This is an electronic version of an article published in Australasian Journal of Philosophy, 90(2). Australasian Journal of Philosophy is available online at: www.tandfonline.com. DOI: 10.1080/00048402.2011.587438

I propose we interpret KNB and AKP as norms of epistemic success, as setting the epistemic standard we aim for our full beliefs to meet.⁹ (There might also be non-epistemic standards governing beliefs, such as truth¹⁰ and pragmatic expedience.) The standard we aim to meet is a kind of adequacy, not an ideal. For example, I aim to write 2000 words today, but more would be ideal. Similarly, one can improve one's epistemic position even if one already knows, so knowing is not an epistemically ideal state. According to KNB, the epistemic norm governing full belief is: I must believe p only if I know that p. As with the norm governing the trip to the airport, this norm defines the standard for success: knowledge is epistemically successful belief. Beliefs can be evaluated for rationality, but what we aim for, and what we principally care about, is knowing. Thus glossed, AKP is never in the business of assessing the rationality of one's basis for acting; *a fortiori* AKP isn't in the business of assessing the rationality of one's basis when it is a partial belief.

The proposed gloss on the knowledge norms shows them to be more fundamental than the norms of rationality. The point of doing what rationality recommends derives from the point of achieving success: the point of picking up the academic-looking stranger who claims to be Professor Jones is (hopefully) to thereby pick up Professor Jones. Christopher Peacocke [2004: 116] puts the point in terms the explanatory asymmetry between two kinds of rule.

The objective rule for making a chicken casserole has the form: obtain chicken and vegetables, then prepare them and cook them in a certain way. The subjective counterpart of this rule is: obtain what you believe to be chicken and vegetables, then do what seems to you to be preparing them and cooking them in the specified way. Someone who aims to follow the objective rule will also be conforming to its subjective counterpart . . . But the only rationale he would have for following a subjectively formulated counterpart rule is precisely that it would be a way he could hope to follow the objective rule . . . The objective rule is explanatorily more fundamental than its subjective counterpart.

Rationality and success might be conflated because both can be specified by saying what one 'should' do. There is a sense in which I 'should' pick up the academic-looking stranger who says she is Professor Jones. There is also a sense in which I 'should not' drive that person into the city, because she is not Professor Jones.¹¹ Similarly, there is a sense in which a justified belief is 'appropriate' even if it fails to be knowledge, and a sense in which it is not 'appropriate'. If one conflates the two senses, one will take KNB to say that only knowledge constitutes justified belief [Sutton 2005, Fantl and McGrath 2009: 124].

AKP and KNB define the nature of epistemic success, not of rationality. (That's compatible with rationality requiring you to *aim to* treat the proposition that p as a reason for acting only if you know that p.) Unfortunately, Hawthorne and Stanley say things like the following [2008: 577, 587 respectively].

[AKP is] the norm of practical rationality. The Action-Knowledge Principle is a norm of rational action.

They start the paper by setting Bayesianism and AKP in competition, presumably interpreting both as setting the standard for an action's being rational. (I quoted the passage towards the beginning of section 3.)

⁹ This gloss undermines one of Brown's objections to AKP [2008: 171–4]. She claims it is 'appropriate' to act on a justified belief that is too fortuitously true to be knowledge. I reply that we aim to avoid the risk of false belief involved in such cases.

¹⁰ Sosa [2010: 185] argues that the 'first-order' aim of belief is truth, which 'induces' the aim of getting the truth through competence rather than luck—i.e. knowing. While it endorses KNB, Sosa's position does not 'put knowledge' first in another way, as it explains knowledge in terms of true belief. As footnote 9 indicates, I am sympathetic to explaining the value of knowing in terms of the disvalue of risking false belief.

¹¹ Let's not fuss about whether there are two things one can *mean* by 'should', which therefore gets a contextualist semantics; or whether we always mean the same thing by 'should', motivating a relativist view. Nothing here turns on that debate.

That suggests they take AKP to set the standard for its being rational to treat something as a premise.¹² Such a principle is wildly implausible: it is rational to act on false but rational beliefs. It is worth noting that Williamson explicitly denies that KNA is a norm of 'reasonableness'. He says that KNA 'makes knowledge the condition for permissible assertion, not for reasonable assertion' [2000: 256].

Given that knowing is the standard for epistemically successful belief, the importance of knowing is secured. Whatever story one tells about rationality and justification won't undermine the central role of knowledge. In particular, rejecting the determination thesis won't undermine the role of knowledge. The rest of this section answers objections to this claim.

One might worry that my response to Schiffer makes knowing a lot less important, because one need not know in order to have a perfectly good basis for action. One can act properly on the basis of a partial belief, and one might allege that means knowing is not important in those cases.

This objection fails, but brings out an important point. Betty's walking to 5^{th} street is epistemically distinctly sub-optimal, because she doesn't *know* that's where the burger place is, even if her hunch turns out to be right. Acting on a partial belief is a fall-back; it is automatically epistemically inferior to having knowledge to act on. Epistemically successful partial belief is second-best to epistemically successful full belief.¹³ (That's not a consequence of AKP, but rather a natural supplement to it.) Betty wishes she knew, and is rightly upset she doesn't. Crucial to Betty's normative position is that she doesn't know where the burger place is. So one cannot use such a case to argue against the normative significance of knowledge. Analogously, a hedged assertion can be epistemically in order, but is automatically sub-optimal compared with flat-out asserting on the basis of knowledge. As we saw above, rationality is not the central way to evaluate beliefs. Evaluating partial beliefs and hedged assertions in the more central way involves seeing them as second-best to knowing. Thus the importance of knowing permeates the evaluations of acting on a partial belief, and of hedged assertions.

We can now answer the following objection. Grant for the sake of argument that subjective Bayesianism is right story about which actions are rational. According to this kind of Bayesian, the subject's rational degrees of belief and utility function (roughly, their desires) determine what action would maximize expected utility, and that's the rational thing for the subject to do.¹⁴ That story applies to all cases of action, and AKP only applies in the special case in which someone acts on a full belief. One might worry that the Bayesian norm of rationality is deeper, because it is more general. AKP does not govern all action *per se*, and the Bayesian norm does. I reply that the importance of knowledge and ignorance permeates all cases of action, even those where one acts rationally on a partial belief. For acting on a partial belief is importantly inferior to acting on knowledge (though it is not less rational).

Relatedly, one might worry about why we should care about reasons and knowledge if the Bayesian story about rational action is right. I reply that one is epistemically better off acting on knowledge rather than on mere partial belief. That's not a difference in how rational it is to act on such a basis; so the difference does not show up in an account of rationality, Bayesian or otherwise.

So it is less obvious that the knowledge norms conflict with subjective Bayesianism than Schiffer and Hawthorne and Stanley suppose. One might object to the (prima facie) compatibility of those theses are follows. The Bayesian might think that the facts about someone's full beliefs reduce to facts about their degrees of belief, plus some other stuff (what their utility function is, what practical questions they face, etc.).¹⁵ One might then allege that the norms governing full belief must derive from those governing degrees of belief. No new knowledge norm could emerge. I reply that norms governing constituted things

¹² Ram Neta [2009: 685] interprets Hawthorne and Stanley as taking AKP to be a norm of rationality.

¹³ Ernest Sosa [2010: 169] remarks that it is epistemically better to know p but have a non-ideal degree of belief in it, than to suspend judgment on q by having an ideal degree of belief in that.

¹⁴ I'm not sympathetic to subjective Bayesianism. Amongst other things, I'd like to see the account of evidence.

¹⁵ Brian Weatherson [2005] makes a serious attempt to reduce full belief to degrees of belief, though I remain unconvinced.

need not derive from norms governing the constituting things. For example, suppose that the facts about people hold in virtue of the facts about particles. There are moral norms governing the treatment of people. There are no moral norms governing the treatment of particles, or at least none that does not derive from the norms about people. The norms governing the treatment of people need not derive from those governing the treatment of particles. Analogously, the norms governing full belief need not derive from those governing degrees of belief.

6. Williamson's Determination Thesis

We have seen that Hawthorne and Stanley's determination thesis is implausible, and is not needed to secure the normative centrality of knowing. This section examines and rejects Williamson's determination thesis.

Williamson claims that 'only knowledge justifies belief' [2000: 185]. That is, any justified belief is justified because it is supported by what one knows, or itself constitutes knowledge. In other words, Williamson holds that what one knows determines what one is justified in believing. The claim is implausible.¹⁶ Suppose Mary serves on a jury, and forms a false but justified belief that the accused is guilty, on the basis of the evidence presented. Twenty years later, she has forgotten what that evidence was, but still believes the accused was guilty. Intuitively, her belief is justified when she recalls it at the later time.¹⁷ Williamson's view says it isn't justified, because it neither is knowledge nor is based on the subject's current knowledge. We can stipulate that Mary is not introspecting, and hence doesn't know that she has the relevant mnemonic appearance. The case shows that one's knowledge does not determine what one is justified in believing.

I think Williamson would reply that Mary does know she has a mnemonic appearance [see 2000: 198–9].¹⁸ Concerning a similar case, he says [199], 'Knowing is a state, not an activity. In that sense, one can know without consideration that things appear to be some way.' But it is psychologically implausible that knowing in such a dispositional sense is Mary's actual basis for believing. So it can't be what makes her belief *doxastically* justified. There is no point formulating Williamson's thesis in terms of propositional justification, for such a thesis can't be extended into a plausible theory of doxastic justification, and so can't be right.

I don't think the knowledge-lover should be impressed by Williamson's motivations for holding that justified belief either is knowledge or is based on knowledge. Williamson presents that thesis at the start of chapter 9 as a way of defending the importance of knowledge [2000: 184]. He quotes Crispin Wright's claim that the proper target of 'epistemologico-sceptical enquiry' is justified belief and not knowledge. He quotes John Earman's claim that what's relevant in the philosophy of science is what theories a scientist is justified in believing, not what they know. Williamson offers the following defence of the importance of knowledge [2000: 184–5].

Grant, for the sake of argument, that knowledge is important now only if it is somehow essential to the present justification of belief . . . Suppose that knowledge, and only knowledge, justifies belief . . . On that supposition, if justified belief is central to epistemologico-sceptical inquiry and the philosophy of science, then so too is knowledge.

Such a defence is otiose if the importance of knowledge doesn't rest on its role in the present justification of belief. If the knowledge norm of belief is true, then knowing is the epistemic standard we aim to meet in believing. Knowing is the standard for success. That secures the importance of knowing. If we accept KNB we take the importance of knowing to be established, and should not grant that knowledge is important

¹⁶ Williamson also claims that, 'Rationality requires one to conform one's beliefs to one's evidence.' [2000: 12, 179.] That's also challenged by the example in the text.

¹⁷ Jackson [forthcoming] defends a theory of justified belief and rationality that delivers the right results about Mary.

¹⁸ I also doubt that knowledge of her mnemonic appearance would account for the high degree to which Mary is justified in believing the accused was guilty.

only if it is essential to the justification of belief. The defender of the knowledge norm should be indifferent to Williamson's reason for thinking that only knowledge justifies belief.

Williamson has another argument that only knowledge justifies belief: it follows from the premises that 'what justifies belief is *evidence*' [2000: 185], and that one's evidence consists of things one knows ($E \subseteq K$). I am convinced by Williamson's arguments for the latter premise [2000: 193–203].¹⁹ For one's evidence to count in favour of a hypothesis, it must be propositional, and indeed be knowledge. That all evidence is knowledge is another thread in 'knowledge first' epistemology we can endorse. However, I am not convinced by Williamson's arguments that what makes a belief justified is the subject's current evidence. The immediate attraction to that view evaporates once we consider the case of Mary. Intuitively, she recalls a false but justified belief, but nothing she knows supports that belief. Given that all evidence is knowledge, Mary's justified belief is not supported by her current evidence.

Williamson gives two brief reasons for thinking that what justifies belief is evidence. His first idea is that no alternative retains the distinction between epistemic and pragmatic justification of belief. He says [2000: 207], 'It is far from obvious that any belief is justified in the truth-directed sense without being justified by evidence.' That bare accusation is unconvincing. There are many theories according to which perceptual and mnemonic epistemic justification is not a matter of having evidence. As the case of Mary shows, the correct theory of mnemonic justification will have that feature.

Williamson's second idea [2000: 208] is that, 'An epistemically justified belief which falls short of knowledge must be epistemically justified by something; whatever justifies it is evidence.' This argument ignores a familiar distinction. Everyone agrees that when a belief is justified, it is justified because of certain features it has. Not everyone agrees that those features 'justify' the belief, in the sense of being evidence for it. For example, one might think that a perceptual belief is justified because it caused by a reliable process, or because it is counterfactually 'safe'. The fact that the belief is caused by a reliable process, or is 'safe', is not evidence the subject has that counts in favour of the belief. Williamson has not given any reason to be suspicious of such an approach to epistemic justification.

In summary: Williamson introduces the view that only knowledge justifies belief as a way of safeguarding the importance of knowledge. But that motivation makes no sense if one accepts the knowledge norm of belief. Williamson's claim has counter-intuitive consequences, and his other arguments for it are weak. The defender of KNB should reject Williamson's determination thesis.

7. Conclusion

The knowledge norms are implausible if interpreted as norms of rationality or justification. For example, we should not take KNB to say that no false beliefs are rational or justified. As I explained in section 5, we should interpret the knowledge norms as setting the standard for epistemic success. Given the truth of knowledge norms thus interpreted, there is simply no need for the knowledge-lover to argue that one's knowledge determines what it is rational for one to do, and what one is justified in believing. Williamson's account of justified belief, and Hawthorne and Stanley's account of rationally acting on a partial belief, cannot be motivated on the grounds that they show why knowledge is important. Their accounts are also implausible. The defender of the knowledge-norms should reject the determination theses with a clear conscience. That's the best way to put knowledge first.²⁰

¹⁹ I pass over Williamson's stronger and more controversial claim that one's evidence is everything one knows (E=K).

²⁰ I am grateful to two anonymous referees, whose comments spurred substantial improvements to this paper. I would like to thank the following for helpful discussions and comments: Matt Benton, Andrew Cortens, Stephen Crowley, Angela Harper, John Hawthorne, Brian Kierland, Mary Salvaggio, Ernest Sosa, Jason Stanley, Andrew Wake, and Brian Weatherson.

REFERENCES

- Brown, Jessica 2008. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning, Nous 42/2: 167–189.
- Fantl, Jeremy and Matthew McGrath 2009. *Knowledge in an Uncertain World*, Oxford: Oxford University Press.

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

- Hawthorne, John and Jason Stanley 2008. Knowledge and Action, Journal of Philosophy 105/10: 571-590.
- Jackson, Alexander forthcoming. Appearances, Rationality, and Justified Belief, *Philosophy and Phenomenological Research*.

Neta, Ram 2009. Treating Something as a Reason for Action, Nous 43/4: 684-699.

- Peacocke, Christopher 2004. The Realm of Reason, Oxford: Oxford University Press.
- Schiffer, Stephen 2007. Interest-Relative Invariantism, *Philosophy and Phenomenological Research* 75/1: 188–195.
- Stanley, Jason 2005. Knowledge and Practical Interests, Oxford: Oxford University Press.
- Sosa, Ernest 2010. Value Matters in Epistemology, Journal of Philosophy 107/4: 167–190.
- Sutton, Jonathan 2005. Stick to What You Know, Nous 39/3: 359-396.
- Weatherson, Brian 2005. Can we Do Without Pragmatic Encroachment? *Philosophical Perspectives* 19: 417–443.
- Williamson, Timothy 2000. Knowledge and its Limits, Oxford: Oxford University Press.