

# Partial Belief, Hedged Assertion, and the Importance of Knowing

Alexander Jackson

3/09/10

Abstract: It is often rational to act on partial beliefs, and they cannot constitute knowledge. Stephen Schiffer (2007) thinks this is a problem for the view that knowing is of central normative significance. I criticize John Hawthorne & Jason Stanley's (2008) response to Schiffer, and defend a different reply. On the view I favour, acting on partial beliefs is automatically second-best to acting on knowledge. Analogously, a hedged assertion is automatically epistemically inferior to a flat-out assertion of something that's known.

## 1. Introduction

Influential recent work defends the central normative significance of knowing. For example, Timothy Williamson argues that Knowledge is the Norm of Assertion (KNA).

(KNA) One must: assert that p only if one knows that p. (Williamson 2000 p. 243)

Williamson suggests that occurrent belief is the inner analogue of assertion. (As Ernest Sosa (2010) says, occurrent belief is 'assent' to a proposition.) So he thinks it is plausible that Knowledge is the Norm of Belief (KNB).

(KNB) One should: believe p only if one knows p. (Williamson 2000 pp. 255-6).

John Hawthorne endorses KNA and the Practical Reasoning Constraint (2004 pp. 21-31).

(PRC) One ought only to use that which one knows as a premise in practical reasoning. (Hawthorne 2004 p. 30)

Jason Stanley holds,

One should act only on what one knows. (Stanley 2005 p. 9)

Hawthorne & Stanley (2008) defend the Action-Knowledge Principle (AKP).

(AKP) Treat the proposition that *p* as a reason for acting only if you know that *p*.  
(Hawthorne & Stanley 2008 p. 577)

Hawthorne & Stanley suggest a way to convert the AKP into a biconditional (called RKP, on p. 578). The following is a variation on their proposal.

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

I am sympathetic to these ‘knowledge norms’, particularly KNB. (See my MS1, MS2, and MS3.) In this paper, I defend that family of principles from a particular challenge. Prima facie, we often act on the basis of *partial* beliefs rather than *full* beliefs. For example, Betty has a degree of belief of 0.6 that the burger place she likes is on 5<sup>th</sup> street, and that makes it rational for her to look there first. Our challenge is to explain why these cases are not counterexamples to AKP, without undermining the importance of the knowledge norm. Stephen Schiffer (2007, pp. 189-90) presses this challenge against AKP, as I explain in section 2. John Hawthorne & Jason Stanley (2008) reply that Betty *is* acting on the basis of something she knows, namely that it is *likely* that the burger place is on 5<sup>th</sup> street. Section 3 argues against their way of defending AKP. Rather, the defender of AKP should admit that there are many cases of rational action about which AKP is silent. Section 4 discusses analogous concerns about whether hedged assertions are governed by KNA, arguing that they aren’t. For example, Betty is not subject to KNA if she says, “I *think* the burger place is on 5<sup>th</sup> street.” Section 5 defends the resulting view from the accusation that it undermines the normative significance of knowing. Knowledge is the norm of full belief, and acting on partial belief is automatically

epistemically second-best to acting on knowledge. It is central to Betty's practical situation that she does not *know* that the burger place is on 5<sup>th</sup> street.

## 2. Schiffer's Objection to AKP

This section describes Schiffer's objection to AKP, and introduces two kinds of response. Section 3 argues against the first kind of reply, which is that given by Hawthorne & Stanley. Sections 4 and 5 defend the second kind of response.

Schiffer (2007) responds to Stanley's claim (2005 p. 9) that, "One should act only on what one knows." Hawthorne & Stanley (2008) refine this claim into AKP. Let's see whether Schiffer's objections hit 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 p. 189). Schiffer complains:

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. (Schiffer 2007 p. 189)

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 Jane is justified in carrying 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 not counterexample to Stanley's original claim that one should act on p only if one knows that p. This is clearer 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 & Stanley 2008 p. 14.) 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. A full belief is a candidate to constitute knowledge; a merely partial belief isn't. By fully believing, one takes oneself to know; by having a merely partial belief, one takes oneself *not* to know. The Knowledge Norm should be interpreted as governing full belief. KNB gets things exactly right: Jane should not fully believe that it will rain.

Schiffer's second worry is the important one. Jane is justified in carrying 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 Jane is justified in carrying an umbrella. On this view, Jane knows something that is a good reason for her to act in that way. Hawthorne & Stanley take this position, claiming that Jane knows that *it is fairly likely it will rain*, and that knowledge is good reason for her to carry an umbrella. I will evaluate this response in section 3, finding it unsatisfactory. The second kind of response to Schiffer's challenge says that AKP does not explain why Jane is justified in carrying an umbrella. On this view, justified 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 & Stanley's response to Schiffer

In this section, I consider the response to Schiffer's challenge favoured by Hawthorne & Stanley (2008). On their view, AKP does explain why Jane's action is appropriate. Hence she is acting on the basis of something she knows.

Hawthorne & Stanley have a specific proposal about what that knowledge is. 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 Jane justified in carrying an umbrella. For if Jane's degree of belief were merely the result of a melancholy turn of mind, she would still know she has that degree of belief, but would not be justified in carrying the umbrella. Nor is Jane's reason for carrying the umbrella is *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 she would still be justified in taking her umbrella. So this proposal fails to explain her being justified.

Rather, Hawthorne & Stanley propose that Jane's reason for carrying the umbrella is that *it is epistemically fairly likely that it will rain* (p. 13). The epistemic probability of a given proposition for a subject is determined by the subject's total body of knowledge (p. 14). For example, if one knows that a coin is symmetrical, that determines that the epistemic probability for one of it landing heads when tossed is 0.5.

Schiffer objects that Jane need not have the relevant concept of likelihood to act rationally on her partial belief (2007 p. 180). More generally, children and some animals act rationally on degrees of belief, while lacking the concept of epistemic probability.

Hawthorne & Stanley make two points in reply. Firstly, they observe that one can have the relevant concept without having words for it, and without clearly distinguishing it from other notions of chance (p. 18). While that's true, I don't think that disarms Schiffer's worry. It still seems true that some children act on partial beliefs and lack the concept of epistemic chance.

Secondly, Hawthorne & Stanley claim that the epistemic probability of a proposition is determined by one's total body of knowledge. So Jane need not believe that it is

epistemically fairly likely that it will rain. Instead, she can act on the things she knows that constitute the evidence that makes it epistemically fairly likely that it will rain. (pp. 16-7.)

Whenever someone appropriately uses [a proposition about epistemic chance] as a reason for acting, there are propositions they know that are *not* about chances that they could instead have appropriately used as reasons for acting. (Hawthorne & Stanley 2008 p. 17)

That is a not a plausible general account of (what is intuitively) acting on the basis of partial belief. Suppose Betty and Brian are on Avenue A, and they want to go to the good burger place. Is it down 5<sup>th</sup> street or 6<sup>th</sup> street? Betty isn't sure, but she has a hunch that it is down 5<sup>th</sup> street. Betty says to Brian, "I don't know whether the burger place is down 5<sup>th</sup> street, or one of the other streets nearby. My hunch is that it is on 5<sup>th</sup>." So Betty and Brian head down 5<sup>th</sup> street. Betty has degree of belief 0.6 that the burger place is on 5<sup>th</sup> street. Hawthorne & Stanley claim that Betty has knowledge that makes it epistemically probable to degree 0.6 that her hunch is correct. But she has no such knowledge. Betty remembers going to the burger place before, but no more about its location than her hunch supplies. The only candidate to be the relevant knowledge is: *that she has a hunch it is on 5<sup>th</sup> street*. Presumably, that's only good reason for her to believe it is on 5<sup>th</sup> street if Betty also knows her hunches are somewhat reliable. But Betty need not have any information about how reliable her geographical hunches are. Maybe she just moved to the city from a village where she was certain of the locations of every landmark. Still, it is rational for her to look on 5<sup>th</sup> street, and given that she is hungry, it would be irrational for her to look on 6<sup>th</sup> street. That's because she has degree of belief 0.6 the burger place is on 5<sup>th</sup> street, and degree of belief 0.3 it is on 6<sup>th</sup>.

If that's right, then Hawthorne & Stanley can't defend themselves by claiming one can always act on the knowledge that determines the epistemic probability of a proposition. So the worry that children act rationally on degrees of belief, but lack the concept of epistemic chance, has force.

But the situation is a lot worse. For Betty does not have knowledge that makes it epistemically likely that the burger place is on 5<sup>th</sup> street. So on the conception of epistemic probability Hawthorne & Stanley appeal to, it is *not true* that it is epistemically likely for her that the burger place is on 5<sup>th</sup>. So Betty can't know it is epistemically likely. So by AKP, she should not take that proposition as a reason for action. I see no alternative conception of epistemic probability that gets round this problem, and is in the spirit of Hawthorne & Stanley's defence of AKP.

For example, it is not open to Hawthorne & Stanley to say that epistemic probability is determined by one's evidence, and deny that one's evidence consists in things one knows. (For example, Conee & Feldman (1985) say that one's evidence consists in the appearances one has.) Such a view will allow that it is epistemically likely for Betty that the burger place is on 5<sup>th</sup> street. (Maybe that's because Betty's weak memory appearance [that the burger place is on 5<sup>th</sup>] is part of her evidence.) But according to that view, one believes justifiedly on the basis of one's evidence, which is not part of one's knowledge. Thus AKP would not explain why our beliefs are justified. The spirit of Hawthorne & Stanley's reply to Schiffer is that AKP explains what makes *any* act justified or not (hence their insistence that AKP explains why Jane and Betty's actions are justified). Their claim that AKP also explains what makes beliefs justified (p. 8) is incompatible with the approach under consideration. Further, it is natural to suppose that one can act justifiedly on the basis of justified beliefs. Given the present conception of evidence, there will be beliefs whose justification does not derive from things one knows, and hence acts that are justified but not because AKP says they are. In summary, Hawthorne & Stanley cannot appeal to the conception of evidence under consideration in order to defend the claim that it is epistemically likely for Betty that the burger place is on 5<sup>th</sup> street. In the absence of a suitable notion of epistemic likelihood, the only sense in which it is likely for Betty that the burger place is on 5<sup>th</sup> street is that she has degree of belief 0.6 that that's the case. As we saw at the beginning of this section, knowledge of such a psychological fact does not make Betty justified in going to 5<sup>th</sup> street (e.g. if she thinks that the burger place is there only because she really likes the number 5).

#### 4. The Norm of Hedged Assertion

The previous section argued against Hawthorne & Stanley's response to Schiffer. I favour the other response, according to which AKP is silent on the epistemic standing of partial beliefs, and actions performed on such a basis. Jane is rational to carry her umbrella, and Betty is rational to look down 5<sup>th</sup> street, and the AKP does not explain why that's so. There is another norm governing what degrees of belief it is epistemically proper to act on. The implausibility of Hawthorne and Stanley's position is clearer if we consider a Schiffer-style objection to the knowledge norm of belief (KNB), rather than to AKP. If Betty has a partial belief and not a full belief, how can KNB explain why that attitude is epistemically proper? My view is that KNB can't explain why Betty's attitude is epistemically proper; I'll explain in section 5 why that doesn't undermine the central importance of knowing. The Hawthorne-&-Stanley style response is that Betty's attitude is one of full belief, which is epistemically proper because it constitutes knowledge, namely knowledge that it is epistemically likely that the burger place is on 5<sup>th</sup> street. But it is psychologically implausible that partial belief turns out to be that kind of full belief.

In this section I consider 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, and some other epistemic norm to govern partial belief.

Suppose Betty says to Brian, "I *think* the burger place is on 5<sup>th</sup> street", or, "I *believe* it is on 5<sup>th</sup> street." Let's take these as our paradigms of 'hedged assertion'. Betty does not know that the burger place is on 5<sup>th</sup> 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 knowing that p being the standard for flat-out assertion that p (i.e. KNA). On the first view, hedged assertion that p flat-out asserts 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 5<sup>th</sup>. (That's analogous to Hawthorne & 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 5<sup>th</sup>*, 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 two proposals about what is flat-out asserted. The first proposal is that in saying "I believe it is on 5<sup>th</sup> street," Betty flat-out asserts that *she partially believes that* the burger place is on 5<sup>th</sup>. (Of course Betty does not fully believe that the burger place is on 5<sup>th</sup>. Otherwise, she would flat-out assert that it is.) In other contexts, Betty could use that sentence to report her psychological state in that way. But that's not the right analysis of hedged assertion, assuming KNA is correct. For one thing, Betty need not have formed the introspective knowledge that she has a partial belief that the burger place is on 5<sup>th</sup>. Knowing she partially believes p is not necessary for Betty to properly make the hedged assertion that p. Nor is such knowledge sufficient. Suppose Brian partially believes that the burger place is on 5<sup>th</sup>, but not because of a hunch, but because his magic 8-ball answered 'yes' to that question. Then his hedged assertion, "I believe it is on 5<sup>th</sup> street," is epistemically out of order. That's so, even though he knows he partially believes it is on 5<sup>th</sup>. 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.

Now consider the view 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. This proposal is analogous to Hawthorne & Stanley's view that acting on a partial belief that p is acting on knowledge that it is epistemically likely that p. Schiffer's objection is weaker in the case of hedged assertion than in the case of partial belief. For it is plausible enough that a child who makes a hedged assertion must have the concept of epistemic likelihood. But the other worry I presented for partial belief is just as pressing. Betty does not have knowledge that makes it epistemically likely that the burger place is on 5<sup>th</sup> street. So 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 5<sup>th</sup> street.” So a hedged assertion that p is not a flat-out assertion that it is epistemically likely that p.

The foregoing is an argument 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. 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<sup>th</sup> 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.

As we saw in the case of Brian’s irrational faith in the verdict of his magic-8 ball, we can evaluate the epistemic standing of a hedged assertion, and the partial belief it expresses. What standard do hedged assertions aim to meet? While nothing in this paper hangs on the answer, here’s my two-pence worth. Betty’s hedged assertion is unsuccessful if the burger place is not on 5<sup>th</sup> street. She would not be happy with her utterance if they walk down 5<sup>th</sup> street without finding the intended restaurant, and the problem would be epistemic. Brian has an unjustified partial belief that the place is on 5<sup>th</sup> street, and even if he’s right, his hedged assertion of the location is not epistemically acceptable. Similarly, a hedged assertion based on a sheer guess is not epistemically acceptable. So the epistemic norm governing hedged assertion requires truth, and truth properly arrived at, as does the norm governing flat-out assertion. The norm for hedged assertion is merely more forgiving concerning the subject’s connection to the truth.

### 5. The Importance of Knowing

I have argued that the best defence of AKP and KNB holds them to be silent on the epistemic status of partial beliefs. This section defends that view from a series of worries that it undermines the normative significance of knowing.

On my view, there are many actions for which AKP plays no part in explaining why they are rational. Allegedly, the subjective Bayesian story is the right one to give for those actions. According to the Bayesian, the subject’s 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. The worry is that this story will crowd

out the normative significance of knowing in the other cases too. For suppose Katie Knows that the burger place is on 5<sup>th</sup> street, and that's her reason for going to that location. Katie has some high degree of belief that the burger place is on 5<sup>th</sup> street. She wants a burger quite quickly. So the subjective Bayesian can explain why it is rational for Katie to go to 5<sup>th</sup> street, without appealing to any facts about what she knows. Allegedly, that's all that's of normative significance. Thus knowing is not important, and AKP does not concern any interesting sense of the bases on which one 'should' act.

This objection fails, but brings out an important point. Grant for the sake of argument that subjective Bayesianism is the right account of what one rationally ought to do. (I'd like to see the account of evidence; and I'd like to know whether all credence functions are fit to be updated on the evidence, or whether only some partial beliefs are 'justified'.) And grant for the sake of argument that Bayesianism can explain why it's wrong to take it as a reason for not buying car insurance that your car won't get stolen, even if not buying car insurance maximizes expected utility. (Good luck with that. But see Weatherson 2005 for a serious attempt.) The objector grants that we do often act on full beliefs. Katie acts on her full belief that the burger place is on 5<sup>th</sup> street. (Presumably it's also true that she acts on her high degree of belief that that's where it is.) But the objection merely assumes that rationality is the only important normative notion. Of course a belief can be false but rationally held, and rationally acted on. But such a belief falls short of the standard we aim to meet, namely knowing.

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 really 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.

Similarly, in believing we aim to know. Beliefs can be evaluated for rationality, but what we aim for, and what we principally care about, is knowing. That's why the ordinary person talks a lot about knowledge, and not so much about justified or rational belief. The norm governing 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. Whether a belief is justified or rational is another matter. When one acts on the basis of a full belief, it is important that the belief constitutes knowledge, though that's not required for the action to be rational. Even if knowledge plays no role in explaining why any action, reasoning or assertion is *rational*, the importance of knowing is not threatened.

So AKP and KNB define the nature of epistemic success, not of rationality. (I develop this line in my MS1.) Unfortunately, Hawthorne & Stanley (2008) say things like the following.

[AKP is] the norm of practical rationality. (p. 8)

The Action-Knowledge Principle is a norm of rational action. (p. 21)

They start the paper by setting Bayesianism and AKP in competition, interpreting both as setting the standard for an action's being rational.

According to one standard picture, ... rational action is a matter of maximizing expected utility. ... 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*. ...

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. (p. 571)

It is compelling that there are false but rational beliefs, and it is rational to act on them. So it is vastly implausible that AKP sets the standard for an action's being rational. Further, that gloss on AKP commits Hawthorne & Stanley to rejecting subjective Bayesianism. That is a very controversial theoretical commitment of their view. My way of glossing AKP does not appear to have that burdensome consequence.

(Hawthorne & Stanley are sympathetic to an alternative decision theory, according to which expected utilities are a function of the subject's epistemic probabilities, not their degrees of belief (pp. 11, 13). Hawthorne & Stanley claim that AKP can do something even this decision theory can't, namely explain what is wrong with doing the right thing for the wrong reasons (p. 12). I don't see how AKP does so. AKP says when p is available to be your reason; it does not say whether p is a good reason for  $\Phi$ -ing. So it won't say whether the subject's reason for  $\Phi$ -ing was a good one. Let me suggest a different line. Hawthorne & Stanley claim that epistemic probabilities are determined by the subject's total knowledge (pp. 13-4). So according to this decision theory, expected utilities are determined by one's utility function and one's knowledge, rather than one's utility function and one's degrees of belief. Such a decision theory is grounded in the importance of knowledge, so can't undermine it.)

Let's move on to consider a different objection to AKP. I admit that many actions are performed on the basis of merely partial beliefs, and acting on such a basis does not fall under the purview of AKP. One might allege that this means that there are lots of cases in which it is not important whether one knows. Doesn't my position make AKP a lot less significant?

This objection fails, but brings out an important point. Betty's walking to 5<sup>th</sup> 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.<sup>1</sup> 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

---

<sup>1</sup> Ernest Sosa (2010) remarks that it is epistemically better to know p but have a non-ideal degree of belief in it, than to suspend judgement on q by having an ideal degree of belief in that.

seeing them as second-best to knowing. Thus the importance of knowing permeates even the cases in which one respects AKP and KNA by not believing or asserting the thing one doesn't know.

Here's a related objection to my view. Suppose we grant for the sake of argument that subjective Bayesianism tells the right story about which actions are rational. 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. As above, 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.

Here's the next objection. One might think that the facts about what full beliefs a subject has reduce to facts about what degrees of belief they have, plus some other stuff (e.g. what their utility function is, what practical questions they face, etc.). One might then allege that the norms governing full belief must derive from those governing degrees of belief. No new knowledge norm can emerge.

Let's grant for the sake of argument that full belief is a matter of high degree of belief and some other stuff. (I'm dubious.) This objection still fails: norms governing constituted things need not derive from norms governing the constituting things. Suppose we are physicalists: 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. Thus:

[The effect one has on people] is just a matter of {the effect one has on certain particles}. It does not follow that any norm governing [the effect one has on people] must derive from more general norms governing {the effect one has on particles}.

But the following is the analogous observation about norms of belief.

[Full belief] is just a matter of {one's degrees of belief}. It does not follow that any norm governing [full belief] must derive from more general norms governing {one's degrees of belief}.

It is not true that the norms governing grounded things must derive from norms governing the grounding things. So there is no objection to the emergence of the knowledge norm governing full belief.

Let's consider a final objection. On that view, full belief is a confused folk notion. Allegedly, the real psychological facts concern degrees of belief, and not outright belief. A serious scientific psychology never talks about someone's treating a proposition as a reason for acting; so from that perspective, AKP never applies. On that view, the real psychological facts, and hence the real norms, concern acting on degrees of belief.

Let's grant for the sake of argument that a scientific psychology will not attribute full beliefs to people. (It's the objector's lucky day.) I say that this objection misidentifies the project of which discovering AKP is a part. AKP is part of the philosophical project of understanding our non-scientific folk notions. We understand each other, unscientifically perhaps, by attributing full beliefs, and evaluating the reasons for which a person acted. It is to this kind of understanding that the notion of assertion belongs: we take others to be telling us something they know. AKP helps us understand how 'our conceptual scheme' works. Folk psychological notions are intimately tied up with how we evaluate each other. For example, talk of what someone's reason for acting was is tied up with talk about whether it was a good reason. That web of folk notions does not aim to line up with the psychological joints in nature, in the way scientific psychology tries to. Those folk notions try to carve pragmatically useful distinctions, not scientifically respectable ones. AKP helps us understand 'our conceptual scheme', which is a web of notions that is not impugned by its lack of scientific respectability. (Explaining why some notions are impugned by a lack of scientific respectability, and others aren't, is a task for metaphysics. I offer an account in my MS3 and MS4.)

It may be helpful to consider a related case. Joshua Knobe (2006) has discovered that our common-sense attributions of intentional action depend on our moral evaluations. In his first case, an executive is told that a profitable business plan will harm the

environment. The executive replies that she doesn't give environmental considerations any weight in business decisions, and approves the plan. Intuitively, the executive harmed the environment intentionally. In Knobe's second case, the executive is told that the business plan will help the environment. As before, she replies that she doesn't care, and she approves the plan merely because it will be profitable. Intuitively, the second executive did not help the environment intentionally. The difference in reactions to the two cases is statistically robust (Knobe 2006). So it seems that our ordinary thinking about intentional action is sensitive to our moral evaluations. (Pettit and Knobe (forthcoming) extend the result to other folk psychological notions.) One reaction to this news is that our ordinary thinking is unscientific, and must be given up. I recommend a different reaction: our ordinary thinking is not trying to be scientifically respectable, so it is no problem that it isn't. We are social creatures, and our ordinary ways of understanding each other are tied up with how we relate to each other socially, in particular how we evaluate others. What others fully believe and know plays a central role in that kind of human understanding.

## 6. Conclusion

Knowledge is the norm of belief. As Aristotle says, and Williamson quotes:

Everyone by nature desires to know. (Aristotle, *Metaphysics* A1 980a21, quoted by Williamson 2000 p. 1)

This is compatible with there being other ways of evaluating beliefs, such as for justification and rationality. When we are aware of our ignorance, we have to fall back on acting on partial beliefs. Even when the partial belief is epistemically proper, it is automatically second-best to knowing. Acting on a partial belief can be rational, but it is normatively central that one thereby acts in ignorance.

## References

- Conee, Earl, & Feldman, Richard. 1985. "Evidentialism." *Philosophical Studies* 48, pp. 15-31.
- Hawthorne, John. 2004. *Knowledge and Lotteries*, OUP.
- Hawthorne, John, & Stanley, Jason. 2008. "Knowledge and Action." *Journal of Philosophy* 105:10, pp. 571-590.
- Jackson, Alexander. MS1. "Aiming to Know (and Falling Short)."  
 ---- MS2. "Knowledge, Belief, and the Proper Grounds for Action."  
 ---- MS3. "A Case for Anti-Realism about Knowledge."  
 ---- MS4. "Understanding Vagueness."
- Knobe, Joshua. 2006. "The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology." *Philosophical Studies* 130, pp. 203-231.
- Pettit, Dean, & Knobe, Joshua. Forthcoming. "The Pervasive Impact of Moral Judgment." *Mind & Language*.
- Schiffer, Stephen. "Interest-Relative Invariantism." *Philosophy and Phenomenological Research* 75, pp. 188-195.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*, OUP.
- Sosa, Ernest. 2010. "How Value Matters in Epistemology." Carus Lecture III, delivered at the APA Central Division, 20<sup>th</sup> February 2010.
- Weatherson, Brian. 2005. "Can we Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19, pp. 417-443.
- Williamson, Timothy. 2000. *Knowledge and its Limits*, OUP.