

Knowledge and Probability

1. Introduction

The notion of subjective probability or degree of belief is sometimes thought to render the notion of knowledge toothless. Jeffrey (1968: 166), for instance, holds that we should ‘deny that the notion of knowledge has the importance generally attributed to it’ and that we should ‘try to make the concept of [degree of] belief do the work that philosophers have generally assigned the grander concept’. Similarly, but in a slightly more conciliatory vein, Lewis (1996: 563) suggests that the ascription of knowledge is ‘one of the messy short-cuts—like satisficing, like having indeterminate degrees of belief—that we resort to because we are not smart enough to live up to really high, perfectly Bayesian, standards of rationality’; although ascriptions of knowledge ‘may yet be indispensable in practice’, ‘they are a handy but humble approximation’.

One might not go as far as Jeffrey and Lewis do but still think that, where the rationality of our actions is concerned, degrees of belief—in combination with utilities—do all the heavy lifting. One might grant that the notion of knowledge plays important roles in epistemology, but maintain that it plays no significant role in practical deliberation.

But while the preceding thought is a natural one, it has been challenged. For instance, Hawthorne (2004: 30) claims that ‘one ought only to use that which one knows as a premise in one’s deliberations’, whereas Stanley (2005: 9) holds that ‘one should only act on what one knows’.¹ If Hawthorne and Stanley are right, then pace Jeffrey and Lewis, knowledge has its pride of place in practical deliberation.²

¹ Similar views have also been put forward by Williamson (2000) and Fantl and McGrath (2009).

² See Weatherston (2012) for another reason as to why knowledge plays an important role in decision making. In this paper, I focus on the kind of reason put forward by Hawthorne and Stanley.

In this chapter, I will consider some objections to the kind of view espoused by Hawthorne and Stanley. I will also discuss an intriguing defence of the view that has recently been advanced by Moss (2013). Moss's defence appeals to the notion of *probabilistic knowledge*—knowledge that is constituted by one's degrees of belief. Although the defence, if successful, will give knowledge its pride of place in practical deliberation, it brings with it its own challenges to the very notion of knowledge.

2. Knowledge and Action

Stanley (2005: 10) writes:

A standard use of knowledge attributions is to justify action. . . . When my wife asks me why I turned left rather than going straight, I reply that I knew that it was the shortest direction to the restaurant. When it turns out that it was not a way to go to the restaurant at all, my wife will point out that I only *believed* that it was the shortest way to the restaurant. To say that an action is only based on a belief is to criticize that action for not living up to an expected norm.

But granting that it is standard to use knowledge attributions to justify action, one might still question the claim that if an action is based on a belief that does not amount to knowledge, that action is to be criticized. For, the thought goes, we often act on our degrees of belief, and it is often proper to do so—no norm is violated in doing so—even though such degrees of belief do not constitute knowledge. To use Schiffer's (2007: 189–90) example, suppose Jane has a 0.4 degree of belief in the proposition that it will rain. Then given certain desires and given that her degree of belief is justified, her carrying an umbrella need not violate any norm. This is so even though her degree of belief, not being truth apt, does not constitute knowledge.

In response, Hawthorne and Stanley (2008: 581–2) hold that one 'possible maneuver'

is to maintain that, in such cases, one is acting on one's knowledge of the relevant chances.³ For example, on this view, if Jane's carrying an umbrella violates no norm, it is because she is acting on her knowledge that there is a 0.4 chance that it will rain. It should be noted that Hawthorne and Stanley are not using the word 'chance' to refer to some kind of objective probability; rather, they are using the word 'to express *epistemic probabilities*, where one's epistemic probability for a proposition is determined by *the total body of one's knowledge*' (Hawthorne and Stanley's emphases). Against the worry that animals or children might rationally act on their degrees of belief even though they might not be able to grasp the concept of (epistemic) chance (Schiffer 2007: 190), Hawthorne and Stanley (2008: 585) reply that we 'should not overintellectualize the concept of chance'—they maintain that 'there are very basic concepts of chance that we employ' even if we might not be able to understand its nature or understand philosophical analyses of chance.

Let's grant that we should not overintellectualize the concept of epistemic chance. Still, there are other worries. Cresto (2010: 329) thinks that Hawthorne and Stanley owe us an account of what epistemic chances are—that they should tell us, for example, how the value of the epistemic chance of p given K is obtained, for any p and any K . Even if a subject may know what the relevant chances are without knowing much about confirmation theory or inductive logic, a theorist *qua* theorist cannot just appeal to the notion of epistemic chance without spelling it out in more detail. Cresto (2010: 327) also thinks that the appeal to our knowledge of epistemic chances 'adds an unnecessary complication, and ultimately distorts the nature of the underlying phenomenon'. If we ask someone why she is carrying an umbrella, a natural reply is that she desires not to get wet, and she knows (or believes) that it is going to rain. A less natural reply would be that she has the higher-order knowledge that she desires not to get wet and she has the higher-order knowledge that her total body of

³ Also, see Stanley (2005: 10).

evidence or knowledge supports the proposition that it will rain.⁴ Baumann (2010: 20) seems to be expressing a similar worry when he holds that it is ‘unclear why the agent has to know their probabilities and why it is not sufficient just to rely on them’.⁵

Hawthorne and Stanley (2008: 585) have a ready reply to the worries. They think that when we base an action on our knowledge of chances, such knowledge is in turn based on knowledge of some other facts about the world. To use their example, our knowledge that there is a 0.6 chance that the restaurant is on the left might be based on our knowledge that three out of five people we have asked for directions indicated that the restaurant is on the left. Further, even someone who does not have the concept of chance will be able to act on knowledge of such facts about the world. Ultimately, Hawthorne and Stanley think that the appeal to knowledge of epistemic chances is not essential. But if that is right, then it is not essential, after all, that they give us a detailed account of epistemic chance. They will also avoid the worry that an act of carrying an umbrella or an act of going to the left is based on higher-order knowledge about our own mental states rather than first-order knowledge about the world.

Baumann (2012: 20) thinks that Hawthorne and Stanley’s attempt to avoid appealing to knowledge of probabilities fails: Why is our knowledge that three out of five people indicated to us that the restaurant is on the left relevant to whether we should go right or left, he asks, ‘if not because it implies something about or simply reduces to a belief about the relevant probabilities?’

Perhaps someone sympathetic to Hawthorne and Stanley’s position might respond as follows. A creature without the concept of chance may act on its degrees of belief. But the performance of such an act is proper only if its degrees of belief are justified by the relevant

⁴ For my purposes I follow Cresto (2010) in using terms such as ‘higher-order knowledge’ or ‘second-order belief’ to refer to knowledge or beliefs about one’s attitudes, where such attitudes need not be beliefs.

⁵ Also, see Schiffer (2007: 190).

pieces of knowledge. For example, if it is proper to act on a 0.6 degree of belief in the restaurant being on the left, such a degree of belief, according to the current response, must be justified by some relevant piece of knowledge, such as the knowledge that three out of five people have indicated that the restaurant is on the left. Perhaps, in this sense, we can be said to be acting on the relevant piece of knowledge even though, in the first instance, we are acting on a 0.6 degree of belief in the restaurant being on the left.

I shall not dwell on whether the response works. For even if we answer Baumann's worry, another one lurks. There seem to be cases in which it is in perfect order for us to act on our degrees of belief even though there is no relevant knowledge to be had. Consider cases in which one obtains a certain degree of belief in a proposition through the 'deliverances of experience' (Van Fraassen 1980: 167–9). For instance, Jeffrey (1968: 172–3) holds that 'in examining a piece of cloth by candlelight, one might come to attribute [subjective] probabilities 0.6 and 0.4 to the propositions *G* that the cloth is green and *B* that it is blue', where such an assignment of degrees of belief is a 'direct result of the observation [by candlelight]'. Presumably, there are cases in which such degrees of belief are rational and we can act appropriately on them. If one holds that in these cases, we must have knowledge of the relevant epistemic chances, then one faces Cresto's and Baumann's worries about the appeal to such knowledge. But it is not clear if there is any other relevant knowledge to be had. Perhaps one might hold that we must have knowledge that an experience of such-and-such qualities has a certain tendency to indicate the presence of an object with a certain colour. But this brings us back to the worry that, on Hawthorne and Stanley's account, whether it is proper for someone like Jane to perform a particular action depends not on her first-order knowledge about the world, but on her knowledge about her own mental states.

Consider also another kind of case that involves memory beliefs. Several epistemologists hold that one's belief may be justified even if one has forgotten one's

original justification for it and has not gained any new relevant evidence. For example, Goldman (1999: 280) holds that many ‘justified beliefs are ones for which an agent once had adequate evidence that she subsequently forgot’. And Bernecker writes (2008: 113-14):

Suppose at t_1 , you come to justifiably believe (and know) that p on the basis of a trustworthy friend’s having told you so. At t_2 , as is normal, you have forgotten a great deal of what you once knew. Among those things that you have forgotten is that it was your trustworthy friend who told you that p . For all you know you could have acquired the belief that p by reading an article in the *National Inquirer* or from some other unreliable source. You no longer remember the original justification for believing p and you have not gained new evidence concerning the proposition in question. But since you remember p most of us would feel comfortable saying that you are still justified in believing p .

Although Goldman and Bernecker seem to have in mind a binary notion of belief, one might think that a similar point holds with respect to degrees of belief. Suppose that we justifiably come to have a degree of belief of 0.7 in p on the basis of a somewhat reliable friend telling us that p . Some time later, we retain this degree of belief in p . Now, we might no longer remember our original evidence for having such a degree of belief, and we might not have gained any new evidence concerning p . But insofar as our memory allows us to remain justified in having a binary belief that p even if we have forgotten our original evidence for it, it should also allow us to remain justified in having a certain degree of belief in p even if we have forgotten our original evidence for it.

How does the discussion concerning memory beliefs bear on the discussion of knowledge and rational action? Well, suppose that one may have a justified degree of belief in a proposition even if one has forgotten one’s original evidence (and has not gained any new relevant evidence) for it. If such a degree of belief is justified, then an action that is

based on it—as well as on the relevant desires—may be justified as well. But there is no relevant knowledge about external facts to be had. After all, the case is precisely one in which one has forgotten the relevant evidence. Now, given the discussion above, a binary belief may be justified and may amount to knowledge even if one has forgotten one’s original evidence for it. But presumably, since it is a categorical error to hold that a degree of belief of 0.7 in p is true or false, and since knowledge entails true belief, such a degree of belief cannot amount to knowledge.⁶ Perhaps one might appeal to knowledge of the reliability of one’s memory or knowledge of the vividness of one’s memory impressions concerning p . But this move is unlikely to appeal to Hawthorne and Stanley given that, ultimately, they want to tie rational action to first-order knowledge about the world.

3. Probabilistic Knowledge to the Rescue?

Moss (2013) has recently put forward an account of *probabilistic knowledge* that promises to deal with the worries raised against Hawthorne and Stanley. According to her, even though a degree of belief in p , unlike a binary belief that p , may not be something that is true or false, it can nonetheless constitute knowledge (2013: 1).⁷ For instance, if Moss is right, a degree of belief of 0.6 in a piece of cloth being green can constitute knowledge concerning the cloth’s being green, and a degree belief of 0.7 in p can constitute knowledge even if one has forgotten one’s original evidence for p . In such a case, one might maintain a strong connection between rational action and knowledge. For example, in the case of Jane, one might claim that her act of carrying an umbrella is proper only if her degree of belief of 0.4 in the proposition that it will rain constitutes knowledge. If Moss is right, there is no need to appeal to knowledge of epistemic chances or knowledge of other relevant facts.

⁶ We’ll see below that Moss (2013) would challenge such a claim.

⁷ Moss is not claiming that knowledge is a gradable notion. But for such a view, see Hetherington (2011), in particular, chapters 2 and 5.

Moss argues for the view that degrees of belief can constitute knowledge by arguing that it ‘yields simple solutions to some difficult problems’. She focuses on two main problems. The first problem concerns how to understand the language of subjective uncertainty. The second problem concerns cases that involve probabilistic analogues of Gettiered beliefs. I will go through each problem and then discuss Moss’s solution.

3.1 The First Problem

As Moss notes, we often utter sentences such as the following:

- (1) It’s probably not raining.
- (2) Joe’s more likely to be in the library than to be in his office.

Drawing on the work of Yalcin (2007), Moss (2013: 5) maintains that such sentences are not truth apt; for her purposes, she proposes to understand the sentences via ‘a very simple expressivist semantics’ according to which, in uttering the first sentence, one expresses the advice to give a high degree of belief to the proposition that it is not raining, and in uttering the second sentence, one expresses the advice to give a higher degree of belief to Joe’s being in the library than to Joe’s being in his office.

Why think that (1) and (2) are not truth apt? Consider the following sentences:

- (3) Suppose that it’s raining, and it’s probably not raining.
- (4) If it’s raining and it’s probably not raining, then...

Both (3) and (4) are infelicitous. But they should not be so if (1) were truth apt. If (1) were truth apt, then when I utter (1), I would be asserting, for instance, the proposition that the epistemic probability that it is raining given my total body of evidence is low. But there is no infelicity in uttering the following:

- (5) Suppose that it is raining and the epistemic probability that it is raining given my total body of evidence is low.

(6) If it's raining and the epistemic probability that it is raining given my total body of evidence is low, then...

This suggests that in uttering a sentence such as (1), we are not trying to report a probabilistic fact about the world.

Given that the above is right, how should we understand a sentence such as the following?

(7) Sue believes that it's probably not raining.

On expressivist theories, (7) does not ascribe a relation between Sue and a proposition concerning probability. In particular, (7) does not say that Sue has a binary belief concerning the probability of rain. Instead it is really a claim about Sue's degrees of belief. It says that Sue's degrees of belief conform to a certain constraint—that Sue assigns a high degree of belief to the proposition that it is not raining.

But now, there is a problem concerning a sentence such as the following:

(8) Sue knows that it's probably not raining.

Traditionally, knowledge that p is understood to entail that p is true. Traditionally, when we ascribe knowledge to a subject, we ascribe a relation between that subject and a true proposition. However, given that 'It's probably not raining' is not truth apt, (8) does not ascribe an attitude towards truths.

3.2 The Second Problem

The second problem is that there are probabilistic analogues of Gettiered beliefs, and there is a question of how we should account for them. Consider the following example from Moss (2013: 9–10):

Several assistants at a reputable entomology supply company select specimens from cultures in response to mail order requests for cereal flies. There are two cultures, one

consisting mainly of *G. hackmani* specimens and one consisting mainly of *G. balachowskyi* specimens. A diligent lab assistant sends out specimens from the former culture along with a letter to the recipient saying that their specimens are more likely *G. hackmani* than *G. balachowskyi*. Meanwhile, several disgruntled lab assistants have started sending out the very same sort of letter along with specimens from the latter culture. On receiving your letter, which happens to be from the diligent lab assistant, you assign more credence to the proposition that your specimens are *G. hackmani* than to the proposition that they are *G. balachowskyi*.

Moss observes that in such a case, our believing that our specimens are more likely *G. hackmani* than *G. balachowskyi*—our assigning a higher degree of belief to the proposition that our specimens are *G. hackmani* than to the proposition that they are *G. balachowskyi*—seems justified on the basis of the letter we receive. Yet, Moss (2013: 10) claims, there is nonetheless something ‘epistemically incorrect’ about our doxastic state. Although our specimens are indeed more likely to be *G. hackmani* than to be *G. balachowskyi*, this seems to be due to luck—in the way that in Carl Ginet’s fake barn case (Goldman 1976), it is due to luck that we are looking at a real barn. Further, we are not inclined to say that we *know* that our specimens are more likely to be *G. hackmani* than *G. balachowskyi*. Suppose that, just as there are Gettiered beliefs, there are Gettiered degrees of belief. It would be nice if we have a theory that picks out certain common properties that non-Gettiered beliefs and non-Gettiered degrees of belief share, and that tells us why such properties are absent in the example above (Moss 2013: 11).

3.3 Moss’s solution and a worry concerning factivity

Moss (2013: 12-13) thinks that both problems can be solved by positing probabilistic knowledge—by holding that one’s degree of belief can constitute knowledge. Probabilistic

knowledge helps solve the first problem. If degrees of belief can constitute knowledge, then Sue's believing that it's probably not raining can constitute knowledge, and Sue can know that it's probably not raining even though 'It's probably not raining' is not truth apt. Probabilistic knowledge also helps solve the second problem. We assign a higher degree of belief to the proposition that our specimens are *G. hackmani* than to the proposition that they are *G. balachowskyi* just in case we believe it is more likely that our specimens are *G. hackmani* than that they are *G. balachowskyi*. Such a belief is epistemically incorrect in just the way that a standard Gettiered belief is epistemically incorrect—they both fail to constitute knowledge.

One might wonder: How exactly is the positing of probabilistic knowledge supposed to help solve the first problem? After all, one might think, the worry is precisely that knowledge is factive, but so-called probabilistic knowledge is not factive. This, as Moss is well aware, is the biggest worry for her. To mitigate it, she distinguishes two senses in which an attitude might be factive, and argues that although knowledge is not factive in the first sense, it is factive in the second sense. In the first sense, for knowledge to be factive is for 'S knows that *p*' to be true only if '*p*' is true (Moss 2013: 14). If Moss is right, knowledge is not factive in this sense—on her view, it might be true that Sue knows that it's probably not raining even though the sentence 'It's probably not raining' is not truth apt. But knowledge may still be factive in the sense that the inference from 'S knows that *p*' to '*p*' is valid, where validity is not defined in terms of truth preservation (Moss 2013: 14).⁸ According to Moss (2013: 15–16), 'any expressivist must develop notions of consequence, validity, and inconsistency within the context of an expressivist semantic theory'; instead of defining validity in terms of truth, an expressivist might hold that an inference from 'S knows that *p*'

⁸ Moss (2013: 18-20) also considers versions of the safety condition and the sensitivity condition for probabilistic knowledge.

to '*p*' is valid in the following sense: 'if you follow the advice that "S knows that *p*" expresses, you thereby follow the advice that "*p*" expresses'.⁹

Moss considers the objection that the first sense of factive is the important one, 'since *factive* verbs should relate their subjects to *facts*' (2013: 17; Moss's emphases). In response, she holds that she is not interested in settling a verbal dispute as to what 'factive' means; she thinks that the second sense of 'factive' does the work that we want a notion of factivity to do.¹⁰ Nonetheless, one might worry that something important has been left out. It is often said that belief aims at truth.¹¹ One might think that there is something epistemically good or correct about a true belief, whether or not the belief is justified. Now, if knowledge entails true belief, then knowledge that *p* brings with it whatever a true belief that *p* brings (and more). There is something good about knowing that *p* in part because there is something good about having a true belief. Can we account for this given Moss's favoured sense of 'factive'?

It is not immediately clear how we might do so. Suppose we follow the advice that 'Sue knows that it's probably not raining' expresses and thereby follow the advice that 'It's probably not raining' expresses. On Moss's view, 'It's probably not raining' expresses the advice to give a high degree of belief to the proposition that it is not raining. But what is epistemically good about following such advice—what is good about the mere adoption of such a high degree of belief? Moss does not say.

⁹ What about the worry that, traditionally, knowledge that *p* entails having a justified binary belief that *p*, but on Moss's view, it doesn't? Moss does not consider this worry explicitly. But it's plausible that she would say that what is important about the concept of knowledge is that 'S knows that *p*' entails 'S justifiably believes that *p*', and on her view, such an entailment is preserved. For instance, on her view 'Sue knows that it's probably not raining' entails 'Sue justifiably believes that it's probably not raining' even though the truth of the latter requires, not that Sue has a binary belief about the probability of rain, but that Sue has a high degree of belief in the proposition that it's not raining.

¹⁰ See Moss (2013: 15–17) for some examples illustrating how the second sense of 'factive' does such work.

¹¹ For discussion of this claim, see, for example, Wedgwood (2002) and Engel (2004).

In light of the above, one might suggest that we need an analogue of truth for degrees of belief—that we need to say what makes for good advice. Hájek (ms) has argued that truth is to belief as objective chance is to degree of belief. Just as there is something epistemically good about a true belief—whether it is justified or whether it constitutes knowledge—there is something epistemically good about a degree of belief that matches the relevant objective chance—whether it is justified or whether it constitutes knowledge. Given this, one might suggest that a degree of belief of x in p constitutes knowledge only if the objective chance that p is true is x . For example, on the current suggestion, the advice to give low credence to the proposition that it is raining counts as good advice only if the objective chance that it is raining is low.

Here is a worry with the suggestion. Given that the objective chance of a past event is either 1 or 0, depending on whether the event occurred, we cannot have probabilistic knowledge of past events if the relevant degrees of belief are neither 1 nor 0.¹² But presumably, proponents of probabilistic knowledge would want to maintain that we can have such knowledge—presumably, on their view, Sue’s degree of belief of 0.9 in the proposition that it did not rain yesterday can constitute knowledge even if the objective chance of the proposition being true is not 0.9.

To get around the above worry, one might suggest that instead of appealing to objective chance, we appeal to epistemic probability. Now, Moss (2013: 23) does consider the suggestion that ‘your credence in a proposition constitutes knowledge just in case it equals your epistemic probability for that proposition’. But she points out that the suggestion does not work. Recall the lab specimens example. Suppose that the diligent lab assistant states in the letter that the specimens are very likely to be *G. hackmani*, and the disgruntled lab assistants do the same in their letters. Suppose that as a result of reading the letter sent by

¹² Lewis (1980) and Schaffer (2007) think that past events are not chancy. But for a dissenting view, see Loewer (2001) and Hoefer (2007).

the diligent lab assistant, we come to have a very high degree of belief in our specimens being *G. hackmani*. It seems that the epistemic probability of the specimens being *G. hackmani* given the relevant evidence is similarly high. But intuitively, Moss claims, our degree of belief does not constitute knowledge.

Now, one might wish to retreat to the weaker claim that one's degree of belief in a proposition constitutes knowledge only if it equals the epistemic probability of that proposition given one's evidence or background knowledge. However, a larger problem remains: The epistemic probability of a proposition given one's evidence or background knowledge does not seem to serve as a good analogue of truth. As Moss (2013: 23) points out, we might understand epistemic probability in terms of the degree of belief it is rational for us to have in a proposition or in terms of the justification or evidence that we have for believing a proposition. So the weaker claim is akin to saying that one's degree of belief constitutes knowledge only if it is rational or justified by one's evidence.

Perhaps the appeal to some kind of epistemic probability will still work if we follow Yalcin (2011) in distinguishing between *rational* speech acts and *advisable* speech acts. A speech act is 'rational in the sense that someone equipped with the evidence of the speaker would be responding appropriately to the evidence by accepting the content of the speech act', whereas it is 'advisable in the sense that a person equipped with full information about the relevant situation would be responding appropriately to that information by accepting the content of the speech act' (Yalcin 2011: 311). Now, consider the slightly modified lab specimens example above. Suppose we are one of those who get a letter from a disgruntled lab assistant, and we claim, 'It's very likely that our specimens are *G. hackmani*'. Is our claim rational? Is it advisable?

There is a way to understand rationality and advisability in terms of different kinds of epistemic probability. Since the epistemic probability of the specimens being *G. hackmani*

given *our evidence* is high, we would be responding appropriately to our evidence by accepting that it's very likely that our specimens are *G. hackmani*. So our claim is rational. But it is not advisable. Since the epistemic probability of the specimens being *G. hackmani* given *full information about the relevant situation* is low, someone equipped with such full information will *not* be responding appropriately to that information by accepting that it's very likely that our specimens are *G. hackmani*. It easy to see that there can also be advisability without rationality. Suppose we get a letter from the diligent lab assistant, but without reading it, we assign a high degree of belief to our specimens being *G. hackmani* because of wishful thinking (we really needed to get our hands on *G. hackmani*). Then it is not rational for us to claim, 'It's very likely that our specimens are *G. hackmani*'. But the claim is nonetheless advisable (in the sense spelt out by Yalcin). For the epistemic probability of the specimens being *G. hackmani* given full information about the relevant situation is high.

Given the distinction between rational and advisable speech acts, one might suggest that advisability serves as the analogue of truth for degrees of belief. More precisely, one might suggest that whereas non-probabilistic knowledge that p requires that p be true, probabilistic knowledge that it's likely that p requires that the epistemic probability that p given full information about the relevant situation be high.

The suggestion is promising. As mentioned, knowledge is epistemically good in part because it entails true belief, and there is something epistemically good about having a true belief, whether such a belief is justified or amounts to knowledge. If the suggestion works, it will enable us to hold that, similarly, probabilistic knowledge is epistemically good in part because it entails believing a claim that is advisable, whether or not the belief is justified or amounts to knowledge.

However, more has to be done to spell out what it is to have full information about a relevant situation. Presumably, in many situations, due to the limited information available to us, we can at most believe that a proposition is very likely to be the case—we shouldn't be absolutely certain of the proposition. But someone with *full* information about the relevant situation might well have information that the proposition is in fact true. So the epistemic probability of the proposition being true given such full information is 1. But then, if a degree of belief constitutes knowledge only if it equals the relevant epistemic probability, it would seem that most intermediate degrees of belief will not constitute knowledge.

3.4 *Two other worries*

Setting the issue of factivity aside, one might also worry about how to understand a sentence such as the following:

- (9) Sue knows that either it's very likely to rain or it's very unlikely to rain (but does not know which).

Suppose Sue hears a weather forecaster whom she knows to be very reliable make an announcement about rain. Due to loud background noise, Sue isn't sure whether the forecaster said 'It's very likely to rain' or 'It's very *unlikely* to rain', although she knows that it was one of the two. Given this, it seems that we should accept (9). But if Sue's knowledge is probabilistic knowledge, what are the relevant degrees of belief?¹³

It won't do to hold that either Sue has a very high degree of belief or she has a very low degree of belief in the proposition that it will rain. In the first case, if her degree of belief were to constitute knowledge, it would constitute knowledge that it's very likely to rain. In the second case, if her degree of belief were to constitute knowledge, it would constitute

¹³ This is a version of the Frege-Geach problem for expressivism with respect to probability talk.

knowledge that it is very unlikely to rain. But Sue neither knows that it's very likely to rain nor knows that it's very unlikely to rain.

One might think that since Sue does not know whether it's more likely to rain or more likely not to rain, she might end up assigning a 0.5 degree of belief to the proposition that it will rain. As we'll see in a bit, I think such a scenario is possible. Nonetheless, if such a degree of belief were to constitute knowledge, it would constitute knowledge that it's 50% likely to rain, and not knowledge that either it's very likely to rain or it's very unlikely to rain.

Rothschild (2012: 109-12) considers a similar problem and suggests that the solution lies in *imprecise* degrees of belief. In particular, suppose we model a person's degrees of belief by a set of probability functions rather than a single probability function. Then we can hold that Sue's doxastic state is modelled by a set of probability functions which are such that some of them assign a very high value to the proposition that it will rain and the rest assign a very low value to the same proposition. Such a doxastic state, one might hold, constitutes knowledge that it's very likely to rain or it's very unlikely to rain.

Those who think that degrees of belief should be precise will not be satisfied with this suggestion.¹⁴ For they will hold that Sue is irrational if she has imprecise degrees of belief, but knowing that it's very likely to rain or it's very unlikely to rain (without knowing which) does not make her irrational. Of course, if one denies that degrees of belief can even be imprecise to begin with, one will also be dissatisfied with the suggestion.

Perhaps one might counter by holding that it is clear that Sue can know that either it's very likely or it's very unlikely to rain. One might then hold that since the best way to account for this is by holding that such knowledge is constituted by rational imprecise

¹⁴ See Elga (2010) and White (2010) for arguments that rational degrees of belief should be precise.

degrees of belief, there is good reason to hold that the latter exist. (Thus we have a new argument for imprecise degrees of belief.)

But here's a worry: It seems possible for Sue to know what she knows and yet lack the relevant imprecise degrees of belief. Consider the following exchange:

Joe: How confident are you that it'll rain?

Sue: **I'm 50% confident.**

Joe: Is there any reason why?

Sue: Well, **I know that either it's very likely to rain or it's very unlikely to rain.**

But I don't know which, and I've no more reason to think that it'll rain than I've reason to think that it won't.

Sue's claims in bold seem consistent. One might object that, given her evidence, she is irrational in being 50% confident that it will rain; instead her degree of belief in rain should be imprecise. However, even if this is right, the point remains that it is possible for Sue to know that either it's very likely to rain or it's very unlikely to rain (without knowing which), and then assign a precise degree of belief to the proposition that it will rain—whether or not she is rational to do so. But then, Sue's knowledge is not constituted by imprecise degrees of belief.

Yet another worry has to do with one's confidence in one's claims. (But this might be more of a request for clarification rather than a worry per se.) Consider the following exchange:

Sue: It's going to rain.

Joe: How confident are you?

Sue: Very confident.

The exchange above is intelligible. In replying to Joe, Sue is saying that she's very confident that it's going to rain. Now, consider a slightly different exchange:

Sue: It's probably raining.

Joe: How confident are you?

Sue: Very confident. (Or: Somewhat confident)

This exchange seems intelligible too, and no less intelligible than the previous one. In replying to Joe, Sue is saying that she's very confident that it's probably raining. But how are we to understand the latter? Supposing that 'Sue believes that it's probably raining' tells us that Sue has a high degree of belief in the proposition that it's raining, what does 'Sue is confident that it's probably raining' or 'Sue is somewhat confident that it's probably raining' tell us about Sue's degrees of belief? Presumably, since each sentence is significantly different from the other two, they shouldn't all be telling us exactly the same thing—that Sue has a high degree of belief in rain. The worry then is that to distinguish between the sentences, we have to understand latter two as reporting Sue's degree of belief in the proposition that it's probably raining. But such a view is presumably antithetical to the expressivist's position. If the latter two sentences are to be understood in such a way, it is at least puzzling why the first sentence should not be understood as reporting Sue's belief in the same probabilistic proposition.

4. Conclusion

One might think that far from being rendered toothless, the notion of knowledge continues to bare its fangs. In particular, one might think that knowledge plays an important role in deliberation. I have focused on the view that one should act only on what one knows, and discussed some objections to it as well as Hawthorne and Stanley's responses to those objections. I have also considered a defence of the view offered by Moss. Moss's defence avoids several of the worries faced by Hawthorne and Stanley. On her view, there is no conflict between acting on one's degrees of belief and acting on what one knows. For

according to her, one's degrees of belief may constitute knowledge. But interestingly, her view introduces some challenges to the very notion of knowledge itself. Although she thinks that the main challenge—which has to do with the factivity of knowledge—can be answered, I have argued that her response might be a bit too quick.

References

- Baumann, P. (2012), 'Knowledge, Practical Reasoning and Action', *Logos & Episteme*, 3: 7–26.
- Bernecker, S. (2008), *The Metaphysics of Memory*, Dordrecht: Springer.
- Cresto, E. (2010), 'On Reasons and Epistemic Rationality', *Journal of Philosophy*, 107: 326–30.
- Elga, A. (2010), 'Subjective Probabilities Should be Sharp', *Philosophers' Imprint*, 10: 1–11.
- Engel, P. (2004), 'Truth and the Aim of Belief', in D. Gillies (ed), *Laws and Models in Science*, London: King's College Publications.
- Fantl, J. and McGrath, M. (2009), *Knowledge in an Uncertain World*, Oxford: Oxford University Press.
- Goldman, A. I. (1976), 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, 73: 771–91.
- Goldman, A. I. (1999), 'Internalism Exposed', *Journal of Philosophy*, 96: 271–93.
- Hájek, A. (ms), 'A Puzzle about Degree of Belief'.
- Hawthorne, J. (2004), *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008), 'Knowledge and Action', *Journal of Philosophy*, 105: 571–90.
- Hetherington, S. (2011), *How to Know: A Practicalist Conception of Knowledge*, Wiley-Blackwell.
- Hofer, C. (2007), 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance', *Mind*, 116: 549–96.
- Jeffrey, R. C. (1968), 'Probable Knowledge', in I. Lakatos (ed), *The Problem of Inductive Logic*, Amsterdam: North-Holland Publishing Company.

- Lewis, D. (1980), 'A Subjectivist's Guide to Objective Chance', in R. C. Jeffrey (ed), *Studies in Inductive Logic and Probability*, 2, Los Angeles: University of California Press.
- Lewis, D. (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549–67.
- Loewer, B. (2001), 'Determinism and Chance', *Studies in History and Philosophy of Modern Physics*, 32: 609–20.
- Moss, S. (2013), 'Epistemology Formalized', *Philosophical Review*, 122: 1–43.
- Rothschild, D. (2012), 'Expressing Credences', *Proceedings of the Aristotelian Society*, 112: 99–114.
- Schaffer, J. (2007), 'Deterministic Chance?', *British Journal for the Philosophy of Science*, 58: 113–40.
- Schiffer, S. (2007), 'Interest-Relative Invariantism', *Philosophy and Phenomenological Research*, 75: 188–95.
- Stanley, J. (2005), *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- van Fraassen, B. C. (1980), 'Rational Belief and Probability Kinematics', *Philosophy of Science*, 47: 165–87.
- Weatherson, B. (2012), 'Knowledge, Bets, and Interests', in J. Brown and M. Gerken (eds), *Knowledge Ascriptions*, Oxford: Oxford University Press.
- Wedgwood, R. (2002), 'The Aim of Belief', *Philosophical Perspectives*, 16: 267–97.
- White, R. (2009), 'Evidential Symmetry and Mushy Credence', in T. S. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*, Oxford: Oxford University Press.
- Williamson, T. (2000), *Knowledge and its Limits*, Oxford: Oxford University Press.
- Yalcin, S. (2007), 'Epistemic Modals', *Mind*, 116: 983–1026.
- Yalcin, S. (2011), 'Nonfactualism about Epistemic Modality', in A. Egan and B. Weatherson (eds), *Epistemic Modality*, Oxford: Oxford University Press.