

- Lewis, D. 1986. Probabilities of conditionals and conditional probabilities II. *Philosophical Review* XCV: 581–9.
- Ramsey, F.P. 1926. Truth and probability. In *The Foundations of Mathematics*, ed. F. P. Ramsey, 158–98. London: Routledge and Kegan Paul
- Ramsey, F.P. 1929. General propositions and causality. In *The Foundations of Mathematics*, ed. F. P. Ramsey, 237–55. London: Routledge and Kegan Paul.
- Stalnaker, R. 1981. A defense of conditional excluded middle. In *Ifs*, eds. W. L. Harper, R. Stalnaker, and G. Pearce, 87–104. Dordrecht: Reidel.

Probabilistic Knowledge in Action

BY CARLOTTA PAVESE

1. Introduction

According to a standard assumption in epistemology, if one only *partially* believes that p , then one cannot thereby have knowledge that p . For example, if one only partially believes that it is raining outside, one cannot know that it is raining outside; and if one only partially believes that it is likely that it will rain outside, one cannot know that it is likely that it will rain outside. Many epistemologists will agree that epistemic agents are capable of partial beliefs in addition to full beliefs and that partial beliefs can be epistemically assessed at least along some dimensions. However, it has been generally assumed that such doxastic attitudes cannot possibly amount to knowledge.¹

In *Probabilistic Knowledge*, Moss challenges this standard assumption and provides a formidable defence of the claim that *probabilistic beliefs* – a class of doxastic attitudes including credences, degrees of beliefs and partial beliefs – can amount to knowledge too. Throughout the book, Moss goes to great lengths to show that probabilistic knowledge can be fruitfully applied to a variety of debates in epistemology and beyond.

My goal in this essay is to explore a further application for probabilistic knowledge. I want to look at the role of probabilistic knowledge within a ‘knowledge-centred’ psychology – a kind of psychology that assigns knowledge a central stage in explanations of intentional behaviour. My suggestion is that Moss’s notion of probabilistic knowledge considerably helps further both a knowledge-centred psychology and a broadly intellectualist picture of action and know-how that naturally goes along with a knowledge-centred psychology. At the same time, though, it raises interesting issues about the notion of explanation afforded by the resulting psychology.

1 I’d like to thank Sarah Moss for comments on an earlier draft. Unless otherwise noted, all references are to her *Probabilistic Knowledge*.

My plan is as follows. In §2, I summarize certain crucial aspects of Moss's notion of probabilistic knowledge. In §3, I motivate a knowledge-centred psychology and I show that this sort of psychology naturally goes with a broadly intellectualist picture of know-how. In §4, I apply Moss's notion of probabilistic knowledge to overcome some outstanding objections to a knowledge-centred psychology and to intellectualism. The final section (§5) raises two issues that distinctively arise for a *probabilistic knowledge*-centred psychology.

2. *Probabilistic knowledge introduced*

One might believe that it will rain. In this case, one has a full, or outright, belief that it will rain. But one might also believe that it *might* rain, or that it is *likely* to rain, or that it is likely that it will rain *if* the clouds are gathering, or that it is *more likely* that it will rain *than* that it will not rain. Moss argues that these sorts of ascriptions can be used to ascribe *probabilistic beliefs* (Chs. 3 and 4).

Moss defends two distinct claims about probabilistic beliefs: the first claim is the *probabilistic knowledge claim* – that is, that probabilistic beliefs can amount to knowledge. The second is the *probabilistic content claim* – that is, the claim that probabilistic beliefs, and hence probabilistic knowledge, have a distinctively *probabilistic* content.

The alternative to the *probabilistic content claim* is the view that probabilistic beliefs are to be understood as *complex probabilistic attitudes* towards simple non-probabilistic contents. On this view, believing, for example, that it is likely that it will rain is a matter of *believing to a high degree* that it will rain; and believing that it is more likely that it will rain than that it will not rain is a matter of *believing to a higher degree* that it will rain than that it will not rain. Moss provides a battery of arguments to show that this view of probabilistic beliefs ought to be replaced by a view on which the content of probabilistic beliefs – not their attitudes – is probabilistic. Accordingly, the content of the probabilistic belief that it will likely rain is not to be modelled as the proposition that it will rain, and not even as a proposition about the objective chances that it will rain, about its evidential probability or about its subjective probability. Rather, it is to be modelled as sets of *probabilistic spaces*, where a probabilistic space is a mathematical entity made out of a domain of possibilities, an algebra of propositions and an assignment of probabilities over that algebra.

As Moss (85) notes, the *probabilistic knowledge claim* and the *probabilistic content claim* are independent. One might endorse the former while clinging to a complex probabilistic attitude view. And one might endorse Moss's *probabilistic content claim*, while rejecting the *probabilistic knowledge claim*. Since my focus here will be on the latter claim, my discussion in the following will abstract from the *probabilistic content claim*.

What does it mean to say that probabilistic beliefs can be knowledge? Consider one's 10% credence that it will rain. This is an example of a probabilistic belief. Now, for that credence to amount to knowledge, it ought to be true, for knowledge is factive. What does it mean for such a credence to be true? Moreover, on many views of knowledge, for a doxastic attitude to amount to knowledge, it ought to be *safe* – for example, it could not easily be false (e.g., Sosa 1999, Williamson 2000, Pritchard 2005, Aarnio 2010). What does it mean for a credence to be easily false?

In response to these issues, Moss (120–1; 154–5) argues that we can understand the factivity and the safety of probabilistic beliefs in *deflationistic* terms starting from what she argues are their standard ascriptions. For example, consider (1a–1d):

- (1a) S believes that it is likely that it will rain.
- (1b) S believes that it is more likely that it will rain than that it will not rain.
- (1c) S believes that it might rain.
- (1d) S believes that it will rain if the clouds are gathering.

According to Moss, there is at least a reading of each of (1a–1d) that ascribes to S a probabilistic belief: on this reading, (1a) ascribes to S a high credence that it will rain; (1b) ascribes to S a higher credence that it will rain than that it will not rain; (1c) ascribes to S the probabilistic belief that has as its content the set of probabilistic spaces whose domains of possibilities encompass some possibilities where it rains; and finally, (1d) ascribes to S the conditional credence that it will rain given that the clouds are gathering.

So far so good. Now, under what circumstances are *these beliefs* true? Here, Moss's answer is deflationistic: there is nothing more to its being true that it is likely that it will rain than its being likely that it will rain; and there is nothing more to its being true that it is more likely that it will rain than that it will not rain than its being more likely that it will rain than that it will not rain. The same goes for safety (107). One's safely believing that it is likely that it will rain is just a matter of not easily falsely believing that it is likely that it will rain; and that amounts to nothing more than one's not easily believing that it is likely that it will rain while it is not likely that it will rain. The important lesson of this discussion is that the proponent of probabilistic knowledge does not need to give up the safety or the factivity of knowledge provided that one understands them in suitably deflationistic terms.

With this in the background, in Chapter 5, Moss argues that probabilistic beliefs are the sort of doxastic attitudes that can amount to knowledge by showing, among other things, that it makes sense for them to be justified or unjustified and that they can be Gettiered. And in Chapter 9, Moss shows that probabilistic knowledge can be fruitfully employed in a defence of a

knowledge norm of action, according to which it is permissible to act only on what one knows, of the sort defended by Hawthorne and Stanley (2008), Weisberg (2013) and Williamson (2017), among others.

Although I think that Moss has made a great case for this latter application of probabilistic knowledge, I think that probabilistic knowledge has a bigger role to play in vindicating a picture on which knowledge enters central stage in the practical domain. As some scholars have argued, the relation between knowledge and action is, arguably, not merely normative but *explanatory*: knowledge enters central stage in *psychological explanations* of behaviour (Williamson 2000, Gibbons 2001, Nagel 2013, Hyman 2015).² In particular, it enters into psychological explanations of *intentional actions*. The normative claim that knowledge is the norm of action and the psychological claim that knowledge enters into psychological explanations of behaviour are quite different claims. The former is compatible with there being intentional actions that are not knowledge-guided and hence, in some sense, epistemically defective but still intentional. According to the psychological claim, by contrast, intentional action can only be there *if knowledge is there*, for intentional action is explained, at least in part, by knowledge.³

The next section motivates a particular kind of knowledge-centred psychology. Then, the following section shows that probabilistic knowledge helps fend off several objections to this sort of knowledge-centred psychology.

3. *Towards a knowledge-centred psychology*

A long tradition assigns beliefs a central role to play in folk psychological explanations of intentional behaviour (e.g., Stich 1978, Fodor 1987, Lewis 1974, Stalnaker 1984, Dretske 1988). But, more or less explicitly, this tradition confines psychological explanations to an explanation of *attempts*. As an illustration of this, consider the usual example of a psychological explanation, where one's belief that there is water in the fridge and one's desire to drink it together are supposed to explain one's attempt to grab a bottle of water from the fridge. Success happens when the world co-operates – when there is indeed water in the fridge. If one's belief is true, then one will succeed

2 Williamson (2000) argues that knowledge enters into psychological explanations of behaviour, with the famous example of the thief who keeps looking for the diamonds *because* he knows that they are there; Gibbons (2001) explicitly argues for the role of knowledge in explaining intentional actions by looking at a variety of examples which I will discuss later. Also, Nagel (2013) defends the thesis that knowledge is a mental state and the role of knowledge in psychological explanations. For another sort of defence of a knowledge-centred account of intentional action, cf. Hyman 2015.

3 Moss (5–6: n. 4) discusses the role of beliefs in psychological explanation and the role of knowledge as a norm for action but does not discuss the possibility of a knowledge-centred psychology.

at finding a bottle; if one's belief is false, one will not succeed at finding water. The dominant thought behind a belief-centred psychology – at least as I will understand it here – is the idea that, as far as the *psychological* explanation of behaviour goes, whether the world complies (e.g., whether there is water in the fridge) is irrelevant: what we want to explain is the fact that *one attempted to get water from the fridge*, whether or not one has succeeded. And one's belief that there is water in the fridge, together with one's desire to drink it, suffices to explain one's attempt, whether or not one's belief is true.⁴

This assumption that psychological explanation should be confined to explaining attempts, rather than successes, relies on two questionable ideas. The first is the idea that actions are *decomposable* – into attempts, on one hand, and into bodily movements, causal chains or successes on the other. The second questionable idea is that attempts are the only 'mental' components of actions – that is, that whatever needs to be added to attempts to get an intentional action is non-mental.

If both of these ideas were true, it would be legitimate to confine psychological explanations to an explanation of attempts rather than successes. After all, if attempts were the only mental components of actions, they would properly fall only under the domain of *psychological* explanations. On the other hand, if actions were not so decomposable into mental and non-mental components, and if actions could be intentional – and hence mental – beyond the attempts' contribution, then this residual 'mentality' of actions too would call for a psychological explanation.⁵

Indeed, there are good reasons to think that actions are not decomposable into mental and non-mental components and that even if attempts were in some sense components of actions, the mentality of actions would not be exhausted by the mentality of attempts. Here is an argument for this conclusion. If attempts exhausted the mentality of actions, then provided that one attempted to perform an action ϕ , one's eventual success at ϕ -ing would have to be intentional. For on this picture, the intentionality, and hence the mentality, of actions would be exhausted by their attempt. However, there are a variety of cases in which one attempts at ϕ -ing, succeeds and yet their action fails to be intentional. That suggests that the intentionality of actions cannot be reduced to the intentionality of attempts; and intentionality being one mark of the mental, that suggests that the mentality of actions cannot be reduced to the mentality of attempts.

4 That a belief-centred psychology aims to explain attempts rather than successes is made explicit in Stich 1978.

5 Cf. McDowell 1995. Cf. Levy 2013 and Williamson 2017 for the point that standard action theory assumes the decomposability of actions into mental and non-mental components. Cf. Gibbons 2001 for a similar argument that psychological explanations should explain successes and not merely attempts.

For considerations of space, I will here consider only two illustrative examples of this phenomenon, both from Gibbons (2001). The first example is *Bobby and the Bomb*. Bobby intends to kill his uncle by a bomb in his house and then, after moving a safe distance away, pressing the large red button on the remote control device. He does not know much about how these things work. He thinks that pressing the button will cause the bomb to detonate but has no idea about the details of this process. His belief is true and, we can suppose, justified. But here is what happens. A satellite, launched by the National Security Agency and designed to prevent bombings of just this kind, intercepts Bobby's transmission; this causes the satellite to send a warning to the intended victim; but, because of an unfortunate choice of frequency, this causes the bomb to detonate. Bobby killed his uncle and caused the bomb to detonate and did intend both things. But he did not do either of these things intentionally. In this example, Bobby attempts to perform an action (killing his uncle) and succeeds but intuitively his action is not intentional. Hence, the intentionality, and hence mentality, of this action is not reducible to the intentionality and mentality of the attempt.

Now, *Bobby and the Bomb* involves a deviant causal chain. But there are other examples of intended and attempted actions that are not intentional which do not involve deviant causal chains. Consider a second example, slightly revised from one given in Gibbons (2001): *Cindy and the Lottery*. In it, Cindy mistakenly believes that someone rigged a lottery in her favour and that she will be handed the winning ticket at the ticket store. On this basis, she believes of a particular ticket that is being handed to her, that if she buys it, she will win. She buys the ticket and wins. So her belief that she will win the lottery by buying that ticket is true. It is even justified. Buying a winning ticket is a perfectly reliable way of winning a lottery. But intuitively, she did not intentionally win the lottery.⁶

Here again, Cindy intends to win the lottery and attempts to do it by buying a ticket that she believes truly and justifiably to be a winning ticket. No deviant causal chain plays a role here in making the ticket the winning ticket. And yet, Cindy's victory is not intentional: fair lotteries cannot be intentionally won. That suggests, once again, that the intentionality of actions cannot be reduced to the intentionality of attempts. But if so, we need a sort of psychology that differs from a belief-centred psychology in that it is not only tailored to explaining attempts.

A kind of psychology where knowledge, rather than belief, enters central stage is promising in this respect. To begin with, note that both cases above can be accounted for by a knowledge-centred psychology. Fred does not know that he can provoke the explosion through his plan. That is why his success is too coincidental to count as intentional. Lottie does not know that

6 Cath (2015) argues against Gibbons' example being a case of intentional action without knowledge. But see Pavese (2018, forthcoming) for a reply.

she can win the lottery by purchasing that particular ticket. That is why her victory is too coincidental to count as intentional. More generally, if one possesses knowledge, then one's belief cannot be lucky (Sosa 1999, Williamson 2000).⁷ Hence, a knowledge-centred psychology can explain why luck can undermine the intentionality of Bobby's and Cindy's successes: by undermining their knowledge. And it is independently plausible that an action based on knowledge is sufficiently under one's control to count as intentional.

Hence, a knowledge-centred psychology can be motivated on the basis of intuitions about cases. But there is also a more theoretical argument for a knowledge-centred psychology, one that hinges on the features that a psychological explanation ought to satisfy in order to count as a satisfactory explanation of its intended *explanandum*.⁸ Good explanations ought to be *counterfactually robust* (Railton 1981, Woodward 2000, Strevens 2008). In other words, an explanation of an event that deserves the name ought to hold in a variety of sufficiently similar situations. For example, a satisfactory explanation of a car accident ought to show that the car accident would also have happened in a variety of sufficiently similar situations.

Now, psychological explanations of actions can be modally robust only provided that the beliefs that they invoke are themselves counterfactually robust. For example, if my reaching for the bottle in the fridge has been caused by my belief that there is the bottle in the fridge, and if my belief is not robust, an explanation of my success invoking that belief will not be counterfactually robust, for it will fail in some sufficiently similar circumstances. Hence, it will not be a satisfactory explanation of my success at satisfying my thirst. That means that for psychological explanations of actions to be good explanations, they ought to invoke modally robust beliefs. This provides some theoretical motivation for thinking that knowledge is the best candidate to feature in satisfactory explanations of actions, for even justified true beliefs give no guarantee of being sufficiently modally robust, if they fall short of knowledge.

These are some of the motivations for a knowledge-centred psychology (for other motivations, see Nagel 2013). My final point in this section is that this sort of knowledge-centred psychology marries well with a broadly intellectualist view of know-how and action. First note that the kind of knowledge that, on a knowledge-centred psychology, explains intentional action is exactly the same kind of knowledge that, on a broadly intellectualist picture, is required by know-how. To see this, consider the kind of knowledge that

7 Some object to a modal requirement on knowledge. See Beddor and Pavese 2018 for a recent defence.

8 For arguments for a knowledge-centred psychology that highlight the modal profile of explanations, see Greco 2016 and Pavese 2018. For more theoretical considerations in favour of a knowledge-centred psychology, see Nagel 2013.

would be needed for explaining intentional action. Start with Goldman's (1970) action theory, according to which one intentionally φ s when one has a plan to φ , where a plan to φ is a belief that specifies the means to φ (cf. also e.g., Audi 1986, Bratman 1987, Ginet 1990, Harman 1976, Velleman 1989/2007, Mele and Moser 1994):

(Intentionality/Belief): If s intentionally φ s, then there are some means m_1, \dots, m_n to φ such that s truly believes that m_1, \dots, m_n are means for oneself to φ .

(Intentionality/Knowledge): can be formulated along the same lines:

(Intentionality/Knowledge): If s intentionally φ s, then there are some means m_1, \dots, m_n to φ such that s knows that m_1, \dots, m_n are means for oneself to φ .

This is, incidentally, the sort of knowledge that Bobby and Cindy lack in the previous examples.

Now, according to standard formulations of intellectualism (Stanley and Williamson 2001, Stanley 2011, Pavese 2015, 2017), one knows how to φ only if, for some means ψ to φ , one knows that ψ is a means for one to φ :

(Intellectualism about Know-How): s knows how to φ is at least in part of a matter of knowing,⁹ for some means ψ to φ , s knows that ψ is a means for oneself to φ .

Hence, the knowledge that *(Intentionality/Knowledge)* requires for intentional action is the same that intellectualists require for know-how.

In fact, a knowledge-centred psychology and intellectualism about know-how are strictly connected views, supporting each other. Start from *(Know-How/Intentionality)*, endorsed by many scholars, intellectualists and anti-intellectualists alike (Ryle 1949, Stanley and Williamson 2001, Stanley 2011, Hawley 2003, Hornsby 2004, 2011, Setiya 2012, Pavese 2015, 2017, 2018):

(Know-how/Intentionality): If s intentionally φ s, s knows how to φ .

Furthermore, suppose that *(Intentionality/Knowledge)* is true so that the intentionality of an action is to be explained at least in part in terms of propositional knowledge. Then by *(Know-how/Intentionality)* and *(Intentionality/Knowledge)*, we get that if one intentionally φ s, one both knows how to φ and one has propositional knowledge of some means to φ :

9 I am stating intellectualism as the view that know-how *requires* knowledge of the means for action. Stating intellectualism as a fully reductive claim would require talking about practical modes of presentation, which I cannot discuss here.

(Know-how, Intentionality, Knowledge): If s intentionally φ s, s both knows how to φ and for some means m_1, \dots, m_n , s knows that means m_1, \dots, m_n are means for oneself to φ .

The intellectualist picture provides the best explanation for why *(Know-How, Intentionality, Knowledge)* should hold. According to this explanation, *(Know-How, Intentionality, Knowledge)* is true not just because of a coincidental aligning of propositional knowledge and know-how in intentional action. Rather, its truth is grounded on the very nature of know-how.

But there is also an argument that goes in the other direction – that is, from intellectualism to a knowledge-centred psychology. Suppose that both *(Intellectualism about Know-How)* and *(Know-How/Intentionality)* are true. Then, the intentionality of an action by an agent requires knowledge of the means to perform it. Hence, intentional action is guided by knowledge – that is, *(Intentionality/Knowledge)* follows. A knowledge-centred psychology not only invites an intellectualist view of know-how and action but is also motivated by it.

4. *Probabilistic knowledge, knowledge-centred psychology and intellectualism about know-how*

So far, I outlined a picture on which a knowledge-centred psychology should replace a belief-centred psychology and on which a knowledge-centred psychology marries well with a broadly intellectualist picture of know-how. In this section, I want to show that Moss's probabilistic knowledge considerably helps with a defence of this general picture. In particular, several objections to *(Intentionality/Knowledge)* and to *(Intellectualism about Know-How)* can be overcome if one accepts Moss's notion of probabilistic knowledge.

The first sort of objection comes from the case of basic actions. Take an action φ that is basic for a subject s at a time t . What knowledge, if any, must s have in order to intentionally φ at t ? Suppose it was the knowledge that some further action ψ -ing was required for φ -ing. Then, φ would not be basic for s at t , for there would be another action ψ that s must perform in order to perform φ . That means if any knowledge is involved in explaining s 's performing a basic action φ intentionally, it must be knowledge that ψ -ing is a means for one to φ , where ψ -ing is not an action itself but rather a matter of *intending to perform one* (cf. Goldman 1970: 66, Setiya 2012: 291).

With this in the background, consider the case of Mary suggested by Setiya (2012: 291–5). Mary's arm has been paralyzed for some time. Although Mary has in fact healed, she does not fully believe that she has. Hence, at the time of trying to clench her fist, she does not believe that she is able to clench her fist nor does she believe that she can clench it by intending to do so. Call this case 'Clenching'. According to Setiya, this does not stop Mary from clenching her fist intentionally, if she intends to do so. If one accepts the

assumption that Mary's clenching is intentional on that occasion, it follows by (*Know-How/Intentionality*) that Mary knows how to clench her fist – even though she does not believe that she is able to clench her fist by intending to do so. The example is at once a putative counterexample to (*Intellectualism about Know-How*) and a putative counterexample to (*Intentionality/Knowledge*), for it supposedly shows that knowing how to perform a basic action does not require full belief, let alone knowledge; and that intentionally performing a basic action does not require full belief, let alone knowledge.¹⁰

Probabilistic knowledge to the rescue! If one appeals to probabilistic knowledge, one can accept that Mary's performance is intentional and yet insist that her performance is guided by knowledge.¹¹ After all, as Setiya (2008: 391) acknowledges, a more fine-grained psychology encompassing degrees of beliefs may provide a more accurate description of this case: Mary plausibly has *some* degrees of belief that she will succeed at clenching her fist that makes the effort of trying worthwhile. Standard epistemology tells us that this credence cannot amount to knowledge: if so, Setiya's (2012) objection stands. But if Moss's *probabilistic knowledge claim* is correct, then this credence, though less than a full belief, can amount to probabilistic knowledge, then Mary's case can no longer be used as a case of intentional action or of know-how without knowledge.

Hence, probabilistic knowledge helps fend off the objection from basic actions. Consider a different sort of objection, relying on cases of subjects differing in their full beliefs but not in their know-how nor in the intentionality of their performances. Two basketball players, Emma and John, attempt a three-points shot. Suppose they have the same chances of success, the same track record of successes and the same evidence that they can succeed. But suppose Emma is just slightly less confident than John that she can do it. In fact, John has a full belief that he will succeed at sinking the basket by shooting as he usually does, whereas Emma fails to have such a full belief. First question: if they both succeeded, is it not plausible that they would have done so intentionally, despite their slight difference in confidence? However, if knowledge requires full belief, a knowledge-centred psychology incorrectly predicts that only John intentionally sank the basket. Second question: is it not intuitive that both John and Emma can know how to sink the basket on this occasion? However, if know-how required full belief about what one can do, only John knew how to sink the basket on that occasion.¹²

Probabilistic knowledge to the rescue! Although John has a full belief that Emma lacks, both have some degree of confidence that they will succeed, and

10 An earlier version of this argument appeared in Setiya 2008: 405.

11 In Pavese 2018, I consider a different response, according to which Mary does not intentionally clench her fist.

12 I owe to Littlejohn (personal communication) a version of this challenge.

since in both cases their confidence is based on the same evidence, in both cases their degree of confidence might amount to knowledge. As soon as we follow Moss in allowing doxastic attitudes other than outright beliefs to be candidates for knowledge, this second putative objection also vanishes.

Thus far, I have considered objections to a knowledge-centred psychology that are also objections to a full-belief requirement on intentional action. Let me consider an objection that distinctively targets the knowledge requirement. Consider Archie who attempts to shoot a target. An expert archer such as he is, he does what he would do in any similar situation in which the wind is quiet. But suppose a gust of wind *could* easily divert his shot, although it does not. His shot can be intentional, it seems, and he might even know how to hit the target, even though his belief about how he can hit the target in that circumstance is unsafe. After all, the way he shot in that circumstance might easily not have been successful, if the gust of wind had intervened.

Probabilistic knowledge to the rescue! The proponent of probabilistic knowledge could object that Archie's belief is probabilistic: it is the probabilistic belief that shooting in a certain way is sufficiently likely, given certain normal circumstances, to lead to successfully hitting the target. We might think of this probabilistic belief as a *conditional credence*: a credence that assigns sufficiently high probability to hitting the target by shooting in a certain way, conditionally on certain normal circumstances being in place. This conditional credence can be safe in this circumstance, and even be knowledge, for in normal circumstances it would not easily be the case that shooting in that way would not likely be successful.

More generally, it is independently plausible that the content of one's knowledge in (*Intentionality/Knowledge*) and (*Intellectualism about Know-How*) ought to be spelled out *in probabilistic terms*. To see this, start with (*Intentionality/Knowledge*). It requires that one has knowledge, of some means, that they are means *for one to ϕ* . What does it mean for some means to be means for one to ϕ ? Not that, for some way ψ of ϕ -ing, one *will* ϕ by ψ -ing: that is far too strong, for one might intentionally ϕ even though one has some doubts about whether one will succeed (Goldman 1970, Harman 1976). For example, consider Davidson's (1971: 50, 1980: 91–94) example of a man who is intentionally making 10 carbon copies as he writes but does not know that he will succeed in so doing. Or consider Bratman's (1987: 37–38) case where one intends to mail the bills on the way to work but one knows that they are forgetful and so doesn't know that they will actually do it. If knowledge is involved in explaining the intentionality of the action, it cannot be the knowledge that one *will* succeed at it.

Should the relevant knowledge be that, for some means ψ of ϕ -ing, one would in *most cases* succeed at ϕ -ing by ψ -ing? This is also too strong: one might intentionally ϕ even though one might fail in most circumstances, as the baseball player who fails at batting 19 times out of 20 may nonetheless

intentionally have batted the one time they succeeded. That suggests that the relevant knowledge is that, for some means ψ of φ -ing, one *could* φ by ψ -ing. But what does it mean that one could φ by ψ -ing, if not that one is *sufficiently likely* to φ by ψ -ing, where what counts as sufficiently likely may vary from task to task?¹³ This gives us:

(*Intentionality/Probabilistic Knowledge*): If s successfully and intentionally φ s at t , then at t s knows, for some means ψ of φ -ing, that oneself is sufficiently likely to φ by ψ -ing.

Finally, consider *Intellectualism about Know-How*. In the original formulation, it is the view according to which knowing how to φ is at least in part a matter of knowing that certain means are means for one to φ . But what does that mean? We do not want to require, for some means to be means for one to φ that one's ψ -ing *invariably* result in one's successfully φ -ing; nor that it result in one's successfully φ -ing *in most cases*. That would be too demanding: after all, Babe Ruth does know how to hit a home run and yet fails at successfully hitting a home run in many circumstances. In order for ψ to be a way for one to φ , all that is required is for one to be sufficiently likely to successfully φ by ψ -ing, where what counts as 'sufficiently likely' may vary with the task at hand (and the circumstances under which the task is being performed). This gets us to:

(*Intellectualism about Know-How*): s knows how to φ only if, for some means ψ for s to φ , s knows that it is sufficiently likely for *oneself* to succeed at φ -ing by ψ -ing.

Hence, if one unpacks the clause *means for one to φ* , a plausible upshot is that both (*Intellectualism about Know-How*) and (*Intentionality/Knowledge*) should be stated by use of a *probabilistic language*: they both require knowledge that one is sufficiently likely to succeed at φ -ing through certain means. These ascriptions are the sort of ascriptions that, according to Moss, can be used to ascribe probabilistic knowledge.

Appealing to probabilistic knowledge in understanding (*Intentionality/Knowledge*) and (*Intellectualism about Know-How*) anticipates a final objection. As Setiya (2012: 294) observes, it seems that one could intentionally perform an action, or know how to perform it, without possessing the concept *could*. And it seems plausible that one can know how to φ and that one can intentionally φ even if one is not able to assess the objective chances of one's succeeding at φ -ing or even if one is not able to assess the likelihood of that outcome on one's evidence. For example, Fido can know how to bring the bone back and can do it intentionally without being able to make any of these sophisticated assessments.

13 Cf. Beddor and Pavese (2018), who suggest understanding abilities in probabilistic terms.

Probabilistic knowledge to the rescue! By appeal to probabilistic knowledge, the proponent of a knowledge-centred psychology can allow for not very sophisticated agents to be able to perform intentional actions as well as to possess the relevant know-how. For example, Fido might have a high degree of confidence that they will bring the bones back and that credence can count as knowledge, even if Fido cannot entertain a proposition about the likelihood of their success.

5. Probabilistic knowledge-centred psychology

In conclusion, the appeal to probabilistic knowledge helps fend off some outstanding objections to a knowledge-centred picture of intentional action. This suggests that a knowledge-centred psychology might best be understood as a probabilistic knowledge-centred psychology. Let me conclude by briefly mentioning two potential issues that distinctively arise for a probabilistic knowledge-centred psychology.

In §3, we have seen that satisfactory psychological explanations of behaviour, *qua* explanations, ought to be *modally robust*. That means that for probabilistic knowledge to figure in psychological explanations of intentional actions, what it means ought to be intelligible for the corresponding belief to be modally robust. This imposes additional pressure on the issue of the intelligibility of the safety of probabilistic beliefs.

As we have seen, Moss (104) gives a deflationistic answer to this question. For one's probabilistic belief that it is sufficiently likely that it will rain to be safe is simply for one not to easily believe that it is sufficiently likely that it will rain while it is not – that is, for it to be the case that in nearby worlds where one has that belief, it is sufficiently likely that it will rain. A worry with this deflationistic account is that, as Moss herself acknowledges (36–37), sentences and ascriptions containing probabilistic expressions can sometimes be used to assert or ascribe probabilistic contents, while at other times they can be used to assert or ascribe contents about objective chances or about contextually determined bodies of evidence. If so, what guarantees that a safety clause for a belief ascription is understandable and intelligible *not* only when it receives a non-probabilistic readings *but also* when it receives a probabilistic reading?

Relatedly, on a probabilistic knowledge-centred psychology, the question of what is the intended *explanandum* becomes urgent. One motivation for a knowledge-centred psychology, as we have seen, is that intentional successes, and not simply attempts, need explaining. Now consider a probabilistic knowledge-centred psychology. Would such a psychology explain the success itself? Or only that the success is sufficiently likely? Probabilistic knowledge might explain probabilistic facts – or *likely-successes*. It is unclear that it can explain the successes themselves. This issue needs more discussion than I can provide here. But the general thought is that replacing a knowledge-centred

psychology with a probabilistic knowledge-centred psychology might demand a retreat from an explanation of intentional successes to an explanation of something less than intentional successes. Whether this retreat is really undesirable or whether it can be avoided are issues that I have to leave for further discussion.

Cornell University
322 Goldwin Smith Hall
Ithaca, New York 14853 USA
cp645@cornell.edu

References

- Aarnio, M. L. 2010. Unreasonable knowledge. *Philosophical Perspectives* 24: 1–21.
- Audi, R. 1986. Intending, intentional action, and desire. In *The Ways of Desire*, ed. J. Marks. 17–38. Chicago: Precedent.
- Beddor, B. and C. Pavese. 2018. Modal virtue epistemology. *Philosophy and Phenomenological Research*. Early View, doi: 10.1111/phpr.12562.
- Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Leland Stanford Junior University, USA: Center for the Study of Language and Information.
- Cath, Y. 2015. Revisionary intellectualism and gettier. *Philosophical Studies* 172: 7–27.
- Davidson, D. 1971. Agency. Reprinted in Davidson (1980), pp. 43–61.
- Davidson, D. 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.
- Dretske, F. 1988. *Explaining Behavior*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, vol. 2. Cambridge MA: MIT press.
- Gibbons, J. 2001. Knowledge in action. *Philosophy and Phenomenological Research* 62: 579–600.
- Ginet, C. 1990. *On Action*. Cambridge: Cambridge University Press.
- Goldman, A. 1970. *A Theory of Human Action*. Englewood Cliff, NJ: Prentice Hall Inc.
- Greco, D. 2016. Safety, explanation, iteration. *Philosophical Issues* 26: 187–208.
- Harman, B. 1976. Practical reasoning. *Review of Metaphysics* 79: 431–63.
- Hawthorne, J. and J. Stanley. 2008. Knowledge and action. *Journal of Philosophy* 105: 571–90.
- Hawley, K. 2003. Success and knowledge-how. *American Philosophical Quarterly* 40: 19–31.
- Hornsby, J. 2004. Agency and actions. *Royal Institute of Philosophy Supplement* 55: 1–23.
- Hornsby, J. 2011. Ryle's Knowing-how, and knowing how to act. In *Knowing How: Essays on Knowledge, Mind, and Action*, eds J. Bengson and M. Moffett, 80–100. New York: Oxford University Press.
- Hyman, J. 2015. *Action, Knowledge, and Will*. Oxford: Oxford University Press.
- Levy, Y. 2013. Intentional action first. *Australasian Journal of Philosophy* 91: 705–18.
- Lewis, D. 1974. Radical interpretation. *Synthese* 27: 331–44.

- McDowell, J. 1995. Knowledge and the internal. *Philosophy and Phenomenological Research* 55: 877–93.
- Mele, A. and P. Moser. 1994. Intentional action. *Noûs* 28: 39–48.
- Nagel, J. 2013. Knowledge as a mental state. *Oxford Studies in Epistemology* 4: 275–310.
- Pavese, C. 2015. Practical senses. *Philosophers' Imprint* 15: 1–25.
- Pavese, C. 2017. Know-how and gradability. *Philosophical Review* 126: 345–83.
- Pavese, C. 2018. Know-how, action, and luck. *Synthese*, doi: 10.1007/s11229-018-1823-7.
- Pavese, C. forthcoming. Knowledge, action, and defeasibility. *Reasons, Justification, and Defeaters*.
- Pritchard, D. 2005. *Epistemic Luck*. New York: Clarendon Press.
- Ryle, G. 1949. *The Concept of Mind*. London: Hutchinson & Co.
- Railton, P. 1981. Probability, explanation, and information. *Synthese* 48: 233–56.
- Setiya, K. 2008. Practical knowledge. *Ethics* 118: 388–409.
- Setiya, K. 2012. Know-how. *Proceedings of the Aristotelian Society* 112: 285–307.
- Sosa, E. 1999. How to defeat opposition to Moore. *Philosophical Perspectives* 13: 141–53.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: The MIT Press.
- Stanley, J. 2011. *Know How*. New York: Oxford University Press.
- Stanley, J. and T. Williamson. 2001. Knowing how. *Journal of Philosophy* 98: 411–44.
- Stich, S. 1978. Autonomous psychology and the belief-desire thesis. *Monist* 61: 573.
- Strevens, M. 2008. *Depth: An account of Scientific Explanation*. Cambridge: Harvard University Press.
- Velleman, D. 1989/2007. *Practical Reflection*. CSLI Publications.
- Weisberg, J. 2013. *Knowledge in Action*. *Philosophers' Imprint*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2017. Acting on knowledge. In *Knowledge First: Approaches in Epistemology and Mind*, ed. A. Carter, E. Gordon, and B. Jarvis, 163–83. Oxford: Oxford University Press.
- Woodward, J. 2000. Explanation and invariance in the special sciences. *British Journal for the Philosophy of Science* 51: 197–254.

Replies to Edgington, Pavese, and Campbell-Moore and Konek

BY SARAH MOSS

1. Reply to Edgington

Edgington raises concerns about four central subjects of my book – my complex content account of credences, my treatment of nested epistemic