

Probabilistic Knowledge by Sarah Moss

Reviewed by Bernhard Salow, Magdalen College, University of Oxford¹

Probabilistic Knowledge is ambitious, original, and extremely wide-ranging. From epistemic modals to imprecise probabilities, from the content of beliefs to the knowledge norm on action, from loose speech to peer disagreement to statistical evidence to the relation between credence and full belief, Sarah Moss discusses a huge array of cutting-edge issues in epistemology and the philosophy of mind and language. Tying the discussion together is the titular phenomenon of *probabilistic knowledge*: the knowledge we attribute with phrases such as “he knew that they might be home” or “she realized she probably liked him as more of a friend”. Moss argues that such knowledge involves not an ordinary belief, but a credence (roughly: a positive credence that they are home in the first, a high credence that she likes him as more of a friend in the second) – and that, so understood, the notion is both coherent and theoretically fruitful. The case is made in three parts. Chapters 1-4 develop and defend the underlying philosophy of mind and language, on which words such as “probably” and “might” let us directly express credences. Chapters 5-7 motivate and defend Moss’s account of sentences like “she knew she probably liked him as more of a friend” as attributing credences that constitute knowledge. And chapters 8-10 apply the idea that credences can constitute knowledge to a variety of topics.

Chapters 1-4

A proposition is something that is true or false, depending on how the world is; at a first pass, we can identify it with a “truth condition”, the set of possible worlds in which it is true. On the standard picture of the mind, attitudes like belief are relations to such propositions, their contents. And, assuming that people have credences (are e.g. 75% confident that P) that are not reducible to beliefs in propositions (e.g. about objective chances, or evidential probabilities), these credences are different relations to these same objects.

Chapter 1 alleges that this standard picture is mistaken: the contents of beliefs are not propositions, but probabilistic contents. Let a *probability space* be a pair consisting of a set of worlds O (the possibilities deemed ‘open’ by the space) and a function from propositions to real numbers obeying the axioms of the probability calculus and assigning non-zero probability only to propositions overlapping O . A *probabilistic content* is then a set of probability

¹ I’ve benefited greatly from discussing these topics with Julien Dutant, Kevin Dorst, Rachel Fraser, Matt Mandelkern, and especially Sarah Moss – though this should not be taken to suggest that they agree with what I say here.

spaces. Corresponding to any proposition P , there is the probabilistic content containing the probability spaces for which $O \subseteq P$; call this probabilistic content CP (for “it is certain that P ”). Where the standard picture analyses your believing that P as you bearing the belief relation to the proposition P , Moss analyses it as you bearing the belief relation to the probabilistic content CP . (Though 3.5-3.6 raise important complications.) So far, probabilistic contents are merely a redundant epicycle. But Moss then uses her more complicated contents to give a new account of your being 75% confident that P : rather than involving a new relation (the *being-75%-confident-that* relation), you bear the same old belief relation to the set of probability spaces that assign 75% probability to P . Complexity in the contents allows for simplicity in the relations.

Moss’s main argument for her picture appeals to the theoretical role of “belief content”, which includes: “The fact that agents believe the same content helps explain why they behave the same way” (p.5); “Subjects agree about something just in case they believe the same content, and disagree just in case they believe inconsistent contents” (p.6); and “Ideal rationality demands that your beliefs ... [have] consistent contents.” (p.6) Moss argues that, once (irreducible) credences are on the table, probabilistic contents play this role better than propositions do: subjects with the same credences behave the same way; people with different credences disagree; and the requirement to have probabilistically coherent credences drops out of the requirement to believe consistent contents only if credences are beliefs with probabilistic contents. So probabilistic contents, not propositions, are the contents we believe.

Chapter 2 adds that probabilistic contents are also the contents we assert – and that, therefore, semantics should assign probabilistic contents rather than propositions to our sentences. Others have argued for related conclusions, based on subtle conversational data about expressions such as “might”, “probably”, and “is x likely”. Most (in)famous, perhaps, are the “eavesdropper” judgements. I’m watching a crime drama, in which Jones commits a murder but makes it look like suicide. Inspecting the crime scene, the detective says: “Smith probably committed suicide”; and I shout: “No, you fool – Jones *killed* him”. My reaction indicates over-investment in the program, not semantic incompetence. This seems a problem for propositional theories. For a proposition expressed by the detective’s utterance would presumably be something like “the evidence available to me makes it likely that Smith committed suicide”. And I do not disagree with that proposition: perhaps Jones’s framing was flawless, and the evidence available to the detective really does indicate suicide. How then can I disagree with the assertion? One solution is that the content of the assertion is not a proposition, but something else, which I do disagree with – e.g. the probabilistic content *Smith probably committed suicide*, which I reject by accepting *Jones certainly murdered Smith*.

Moss takes these familiar arguments to be instructive – they are conversational analogues of some of her claims about belief contents – but does not want to rely on them, since both the conversational data and the prospects for explaining them in a propositional framework remain controversial. Instead, she offers a “foundational” argument, based on the previous conclusion that we believe probabilistic contents. The details are subtle, but the basic idea is straightforward: it’s obviously attractive to maintain that the contents of assertion and belief coincide.

Chapters 3 and 4 develop a sophisticated semantics for expressions such as “might”, “probably”, “is x likely”, and indicative conditionals, exploiting the idea that sentences express probabilistic contents. This account, and the data Moss marshals to support it, are interesting and important; but they are quite involved, and so Moss is, I think, right to recommend that readers with other interests skip ahead and return only as and when the details matter.

On Moss’s presentation, much rides on the argument of chapter 1. I had two worries about that argument. The first concerns the notion of consistency featuring in the theoretical role of “belief content”. Moss assumes that for two contents to be consistent is for them to intersect: it is in this sense that the probabilistic content *It’s 60% likely that P* is inconsistent with the probabilistic content *It’s 70% likely that P*, which, on Moss’s picture, explains why someone who assigns P 60% credence disagrees with someone who assigns it 70% credence. But I doubt that this is the notion of consistency we have in mind when we endorse things like “subjects [...] disagree just in case they believe inconsistent contents.” When we endorse this supposed platitude, we’re instead using a naïve understanding of consistency as something like “could be true together”. This is equivalent to the definition in terms of intersection if we identify contents with the set of possibilities in which they are true; but, of course, Moss rejects that identification. It thus isn’t clear, at least at this stage of the argument, that Moss’s theory really does vindicate the theoretical role of “belief content”, properly interpreted.

The worry feels particularly forceful when Moss discusses explanations of conflicts between attitudes. Mark Schroeder (2008) influentially claims that conflicts between beliefs that have inconsistent contents are less mysterious than other conflicts. Moss argues (pp.10-12) that this supports her picture over the traditional one, since her account classifies conflicting credences (e.g. 60% and 70% credence in P) as having inconsistent contents. But this seems to equivocate between the two interpretations of “consistency”. Schroeder’s idea that inconsistency in contents can explain conflicts has some plausibility when “consistency” is understood as “could be true together”: after all, beliefs aim at truth, so two beliefs which are inconsistent in *that* sense can’t both succeed in their aim. By contrast, I don’t see how the fact that two beliefs have non-

intersecting contents makes, all by itself, the conflict between them any less mysterious. Schroeder's idea thus seems plausible only on one reading; but it supports Moss's account only on the other.

(Admittedly, this all becomes more complicated if we adopt theories of 'can' and 'true' on which probabilistic contents can be true, but not all can be true together – as Moss does later in the book. I suspect that some of the worry – especially the part about explanatory force – survives this complication; but I can't make that case here.)

My second worry is more fundamental: I'm unsure what – aside from terminology – really distinguishes Moss's picture from the standard one. We can use "content" as Moss does, and uphold the theoretical role she describes. But we can also use it to refer to propositions, and refine our accounts of when people disagree or are irrational instead. Moreover – at least when, as Moss is happy to do for the purposes of chapter 1, we focus on people with precise, coherent, and complete credences – we can easily translate between the two ways of speaking: where Moss talks of someone bearing the belief-relation to the singleton set {Pr}, the traditional theorist talks of that person bearing, for each proposition P, the believe-to-degree-Pr(P) relation to P. The distinction feels like that between representing objects as bearing the mass-in-kg-and-velocity-in-meters-per-second relation to a pair of numbers, or representing them as bearing the mass-in-kg relation to the first member of that pair and the velocity-in-meters-per-second relation to the second. That is, it feels like a matter of bookkeeping.

(Moss does discuss whether the dispute she has set up is merely verbal in section 1.3; though her response emphasizes the Schroeder-style explanatory difference discussed above. It's also worth noting that, even if the theories ultimately are notational variants, it may still be that one represents the relevant facts – about consistency, disagreement, etc – more perspicuously; and considerations of this kind may well favour Moss's account.)

By contrast, propositional and probabilistic accounts of the contents of assertions really do seem different; it isn't obvious how to convert explanations between the two. The main reason is that propositional theories lack an analogue of the *being-x-confident-of* relations: they do not postulate infinitely many "forces" with which assertions can be put forward. If they did, it would be far less clear how they differ from probabilistic accounts (except, presumably, on issues arising from compositionality).

If I'm right about this, the contents of beliefs can't be crucial to the foundational argument: any tension between accepting probabilistic contents for beliefs but rejecting them for assertions should already be present in accepting *being-x-confident-of* relations to propositions without accepting corresponding *asserting-to-degree-x* relations. Since this is a standard combination, a serious tension here would be surprising.

That's not to say there isn't one – Moss's arguments deserve detailed consideration that I cannot give them here. But the above suggests that these arguments may rely less on her theory of content than she indicates.

Chapters 5-7

Chapters 1-4 argue that we can believe and assert probabilistic contents; chapters 5-7 argue that we can know them. Chapters 5-6 argue that such knowledge is intelligible; chapter 7 responds to sceptical arguments alleging that, while it is intelligible, we don't have much of it.

As I see it, the central positive argument is Moss's observation that we do, in fact, use 'know' (and other factive attitude verbs such as 'realize' or 'remember') to embed sentences featuring epistemic vocabulary: "She realized she probably liked him as more of a friend"; "He remembered that they might not be home yet"; "She gave him the money despite knowing full well that he'd probably never repay her". Moreover, the relevant epistemic vocabulary seems exactly the kind that resists propositional analysis: following up on "she realized she probably liked him as more of a friend" with "though in fact she was romantically attracted to him" feels odd, requiring us to revisit and reinterpret the initial claim. If that sentence were merely attributing knowledge of propositions about probabilities, this would be puzzling. So it must attribute knowledge of a non-propositional, probabilistic content. So such knowledge is intelligible.

Moss backs up this central argument with several reasons to take probabilistic knowledge attributions at face value. One is that they display the same subtle patterns as propositional knowledge attributions; in particular, that they admit of "Gettier cases". Alice believes that she is probably attracted to Bert, based on the flutters she just experienced when she saw him. Now, Alice is in fact probably attracted to Bert – but in this instance, her flutters were caused by a drug in the coffee she just drank. This certainly feels like a classic "Gettier case"; and Moss reports the intuition that Alice does not know that she is probably attracted to Bert, despite her justified true belief that she is.

Other considerations respond to objections. One might worry where probabilistic knowledge could come from; Moss argues that it could have any of the familiar sources, including perception. One might worry that credences cannot be true while knowledge always is; or, relatedly, that credences can't display putative features of knowledge such as reliability, safety, or sensitivity. Moss responds that truth, in anything but a highly deflationary sense, is a red herring. What matters for the factivity of knowledge is that *S knows that it's likely that P* entails *it's likely that P*; what matters for sensitivity or safety is that *S* wouldn't have believed *P* was likely if it hadn't been, or that it couldn't easily have happened that *S* believes *P* to be likely even though it's not. As long as we

can make sense of these, probabilistic beliefs can have all the features we associate with knowledge.

Moss is undoubtedly right that some probabilistic knowledge attributions are hard to explain as attributing knowledge of propositions. But this phenomenon generalizes beyond the epistemic realm. We also sometimes attribute knowledge (and related states) that things are tasty: “My son has finally realized that mushrooms are actually rather tasty” or “Although he knew that the cake would be tastier, he went for the fruit instead.” As Peter Lasersohn (2009) notes, such attributions are similarly hard to account for in a propositional setting. Which proposition does my son discover? That mushrooms are tasty *to him*? That wouldn’t explain why I can’t say “although mushrooms are disgusting, my son has finally realized that they are actually very tasty”. That mushrooms are tasty *to me* (or *to both of us*)? That seems too strong, since he could realize they are tasty while mistakenly thinking I abhor them. Moss’s account of probabilistic knowledge attributions, unlike Lasersohn’s, doesn’t generalize naturally to these.

More directly probative is the fact that, while probabilistic knowledge attributions are natural, they also feel odd on reflection. “She realized that she probably liked him as more of a friend” suggests to me that the speaker changed his mind about how much he wanted to say halfway through the sentence: that initially, he wanted to attribute knowledge that she liked him as more of a friend, but then became reluctant to commit to the truth of that belief himself. In that mindset, the sentence strikes me as slightly muddled; not completely uninterpretable, but puzzling nonetheless. More generally, I find that I often struggle making judgments about probabilistic knowledge. Does Alice know that she is probably attracted to Bert? Can you know, based merely on the fact that most iPhone thefts are committed by men, that it was probably Jake rather than Barbara who stole your phone? I find myself uncertain.

We can distinguish two challenges here. One claims that judgements about probabilistic knowledge are not sufficiently clear and robust to be taken at face value; the other that they are not sufficiently clear and robust to be used in further theorizing. Moss most explicitly responds to the former; but, given her ambitions in chapters 8-10, both are relevant.

Explaining how credences can satisfy the various theoretical conditions on knowledge could help with both challenges. Unfortunately, I experience similar limitations to my intuitions here, at least with the modal conditions involved in Safety and Sensitivity. If it hadn’t been likely that Jones is at home, would Smith still have believed it to be? Could it easily have happened that it isn’t probable that Jones is home but Smith still believes that it is? Very often, I struggle to hear the intended reading, on which the probability claims are not propositions about physical or evidential probabilities. Moss does give

examples (p. 106, p. 152-4) where she thinks that reading is clearly present. But I could have done with more, to develop a clearer sense of how to approach new cases.

Moss (p. 106) also draws an analogy with ethical expressivists, who think they can make sense of claims like “If taking that piece of bread wasn’t OK, Smith wouldn’t have said that it was”. But ethical expressivists usually explain in more detail just how these are to be evaluated. They might, for example, propose the following: consider various ways of changing the non-moral facts to get a scenario such that your actual values tell against taking the bread in that scenario, e.g. by including others who need the bread more, or a promise that the bread remain untouched. Then look at which of these changes are minimal. And then see whether the minimal changes would also result in Smith saying that it wasn’t OK to take the bread. It’s not clear what the analogous procedure is in the probabilistic case, since there is no equally intuitive distinction between the non-probabilistic facts (which you’re allowed to change) and the purely probabilistic ones (which you hold fixed).

None of this shows that probabilistic knowledge is unintelligible, or too vague to be theoretically important. I don’t mean to assert anything that strong – only to express uncertainty, and to indicate where I think future work on the notion might help.

Chapters 8-10

With the possibility of probabilistic knowledge established, chapters 8-10 put it to work. The applications are myriad, including issues as varied as pragmatic encroachment, peer disagreement, scepticism about perceptual knowledge, transformative experience, and the limits of statistical information, especially in morally charged situations. In chapters 8 and 9, the pace of the discussion is often quite fast, as Moss wants to get to the aspects of these debates where probabilistic knowledge can be helpful. This makes sense; but it means that readers unfamiliar with the relevant literatures may need to use supplementary readings to keep up and ensure they don’t come away with a slightly skewed impression of the central problems. (By contrast, chapter 10, dealing with statistical information, could well be used to introduce its topic to sufficiently advanced students who are familiar with Moss’s core theory.)

I’ll give one example, from the discussion of perceptual scepticism in section 8.5. Suppose you’re looking down at what seem to be your hands. Let *Seems* be the proposition that things seem to you like THIS; *Hands* the proposition that you have hands; *Fake* the conjunction of *Seems* with not-*Hands*. “Dogmatists” want to say that perceptual seemings as of a hand provide evidence against *Fake*, so that they provide strong evidence for *Hands* even if you had no prior reason to reject *Fake*. But it’s hard to see how. The evidence you get from your

perceptual seemings looks to be captured by *Seems*; but since *Fake* entails *Seems*, $\Pr(\text{Fake}|\text{Seems}) > \Pr(\text{Fake})$, and so *Seems* is evidence for, not against, *Fake*. Moss's response (part of it, anyhow) is that *Seems* does not capture the total evidence you get from your perceptual seemings: what you learn from them, the dogmatist should say, includes also the probabilistic content *if Seems, then probably not-Fake*. Once we include this probabilistic evidence, the total evidence provided by the perceptual seemings can easily tell against *Fake*.

This is a good point; but it hardly settles the matter. (Not that Moss says it does – but not mentioning related problems may, nonetheless, create that impression in readers not familiar with the literature.) Here is a follow-up worry. Suppose you know that dogmatism is true, that perceptual experience gives you evidence that it is unlikely to be misleading. Then, knowing only that you're about to have a perceptual experience, you already know that it will give you evidence that it is unlikely to be misleading. But it seems very plausible that if you know you'll get some evidence, you don't need to wait until you do to take it into account. So you already have justification to believe that your perceptual experiences are unlikely to be misleading – and so, by the dogmatist's own light, you have prior reason to reject *Fake*, which was the very thing she wanted to do without. Roger White (2006) already raises this problem in the paper Moss is responding to, and Moss's response to the first problem only brings it into sharper focus.

In presenting her applications, Moss also relies heavily on the framework developed earlier in the book. This, too, makes sense; but it, too, has downsides. It means that those who, like me, have uncertain intuitions about probabilistic knowledge will sometimes struggle. I mentioned that I feel rather at a loss about whether you can know, based merely on the fact that most iPhone thefts are committed by men, that it was probably Jake rather than Barbara who stole your phone. But it's important to Moss's applications that you cannot. Reliance on the framework also means that Moss doesn't always discuss in detail which aspects of her picture are required for particular applications, and which are merely optional add-ons. For example, Brian Miller (2016) proposes a defence of dogmatism similar to Moss's, without invoking probabilistic contents, much less probabilistic knowledge. I expect that we might well gain something by embedding this defence into Moss's larger framework (for example, probabilistic knowledge might play an essential role in the theory of probabilistic evidence), but that issue is not addressed here.

Concluding Remarks

There is no doubt that *Probabilistic Knowledge* is an impressive achievement. It contributes to a vast array of interesting issues, tying them together in a novel and exciting way. The discussion is consistently fresh and original – Moss doesn't just develop the obvious consequences of a central idea, but introduces

new and illuminating considerations at every turn. And the details of her argumentation substantially raise the bar for what linguistically informed epistemology can look like; the discussion of loose speech, belief, and credence in sections 3.5-3.6, and of quantified belief ascriptions and Safety in section 7.5, neither of which I have been able to discuss here, deserve special mention in this respect.

That's not to say I was always completely convinced; and I've tried, above, to articulate some reasons for this. But it's worth pointing out that these mostly boil down to requests for more of the same; more about the nature of the dispute in chapter 1, even more explicit theorizing about counterfactual and related constructions containing probabilistic vocabulary, more details on which theoretical ideas are essential to which applications. And, perhaps, this is exactly what a book aiming to be a "springboard for further research" (p.242), aiming to reorient a field towards a new phenomenon rather than exhaustively describing a familiar one, should leave its readers wanting.

References

Lasersohn, P. (2009) "Relative Truth, Speaker Commitment, and Control of Implicit Arguments" *Synthese* 166: 359-374

Miller, B. (2016) "How to be a Bayesian Dogmatist" *Australasian Journal of Philosophy* 94: 766-780

Schroeder, M. (2008) *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford UP

White, R. (2006) "Problems for Dogmatism" *Philosophical Studies* 131: 525-557