Review of *Probabilistic Knowledge*, by Sarah Moss

Daniel Greco

In general, it's possible to speak your mind—to say what you believe. And in the good case, you can know whereof you speak. These seem like truisms. But from them, Moss is able to build a novel view about content, with implications for a wide range of recent debates; *Probabilistic Knowledge* is essential reading for philosophers of language, semanticists, and formal epistemologists.

Moss's view involves positing a kind of harmony between the contents of belief, the contents of assertion, and the objects of knowledge. On a popular view about the nature of assertion, its function is to take something you believe and add it to the common ground of a conversation. (Stalnaker, 1978) Suppose you and I are making plans for the weekend. Having seen a gloomy forecast, I believe it will rain on Saturday. So I say that it will rain on Saturday, and you believe what I say; as the conversation progresses, we take rain for granted and plan around that. Plausibly, what I believe, what I convey in speech, and what you come to believe are all the same thing—the proposition that it will rain on Saturday.[1] To this package Moss adds the Williamsonian (2000) idea that beliefs and assertions are normatively appropriate only when they amount to (in the case of beliefs) or express (in the case of assertions) knowledge.

This package works particularly cleanly when the propositions we believe, assert, and

For helpful comments and discussion, thanks to Kenny Easwaran, Sarah Moss, and Zoltán Szabó.

[1] In earlier work, Moss (2012) calls this the "package delivery" model of communication.

know, are understood as either identical to, or at least capable of determining, sets of possibilities. That way adding a proposition to the common ground can be modeled by set intersection; we can intersect the set of possibilities compatible with what is asserted with the set of possibilities compatible with the common ground prior to the assertion to obtain the new set of possibilities compatible with the common ground after the assertion has been made and accepted.

So far, so familiar. One way of understanding Moss's project is to see it as drawing out the consequences for this whole package of adopting a probabilistic conception of belief, of the sort standard in Bayesian epistemology. This isn't exactly how Moss argues; she offers independent arguments for moving to probabilistic conceptions of belief, assertion, and knowledge, and treats each of the parts of her package as mutually reinforcing in light of the views about the function and norms of assertion just discussed. Nevertheless, perhaps because of my own philosophical background, I find it easiest to understand Moss's project as driven by a view about the nature of belief.

Bayesian epistemologists often argue that if we try to describe someone's opinions by just listing which possibilities are compatible with what she believes, the description we give is too coarse grained. We believe some things more strongly than others. That it will rain tomorrow and that it will rain every day next week are both compatible with my beliefs (as I type, it's looking gloomy). But I'm more confident of the former than the latter. While some philosophers have speculated that strength of belief might turn out to be a kind of epiphenomenon of which propositions we accept in a binary fashion, in the Bayesian tradition strength of belief is modeled directly. An agent's state of opinion is modeled by a domain of doxastic possibilities, an algebra of subsets of that domain, and a probability measure on that algebra. How strongly an agent believes a proposition is modeled by what probability the measure assigns to the set of possibilities corresponding to it.

Even if this picture is an improvement over a simple binary approach to belief, it can

seem highly artificial, due to the precise levels of confidence it posits. I take seriously the possibility that I'll have grandchildren who'll live past 200, but I would have a very hard time quantifying my uncertainty on this matter. We can be unsettled as to just which possibilities we should be taking seriously, and also as to how seriously we should take them. For this reason, many writers, Moss included, favor a generalization of the standard Bayesian picture, on which our state of opinion is modeled not by a single probability space, but instead by a set of probability spaces. In this review I'll call it a "belief set", but remember that's not a set of possibilities but a set of probability spaces (each of which itself includes a set of possibilities). If I'm unsettled as to whether to take a possibility seriously *at all*—if I'm not sure whether something *might* be true—then it will be included in some of the probability spaces in my belief set, but not others. If I'm unsettled as to *how* seriously to take some possibility—if I'm not sure how likely it is—then not all of the measures in my belief set will assign it the same probability.

Most of the philosophical literature on imprecise probabilities has focused on epistemology and decision theory. How should they be updated? How should they interact with utilities (potentially imprecise themselves) to inform behavior? A much less familiar set of questions concern their implications for semantics and philosophy of language. Once asked, a number of natural applications suggest themselves. If we can have opinions modeled by sets of probability spaces, then it's natural to expect that we should be able to express those opinions. So a promising research program is to look for examples of assertion that are hard to make sense of if we must understand them as conveying traditional contents, but easier if we can understand them as conveying a probabilistic content, i.e., a content corresponding to a set of probability spaces.

Explicitly probabilistic assertions are perhaps the most obvious place to look. When I sincerely assert that it will probably rain tomorrow, and my conversational interlocutors accept my assertion, how should we model what happens? Moss's view suggests a natural answer. I believe that it will probably rain, which on Moss's view amounts to its

being the case that all the probability spaces in my belief set assign the (traditionally understood) proposition that it will rain some highish probability, say > 0.5. That is, there is a probabilistic content—the set of spaces that assign the proposition that it will rain probability > 0.5—that I believe, because every probability space in my belief set is contained in that content. I assert the very content I believe, and if it's accepted by my interlocutors, then the conversation will evolve accordingly. As with the traditional view, the effect of an accepted assertion is modeled by set intersection; the ultimate state of the common ground is obtained by intersecting the initial set of probability spaces in the common ground with the set corresponding to my assertion. So if my assertion is accepted, the common ground will evolve to contain only spaces according to which rain is probable. If, in addition to accepting my assertion for the purposes of conversation, my interlocutors believe what I say, then they too will be confident that it will rain.

Moss's view suggests a similarly straightforward treatment of epistemic modals. Suppose I say that it *might* rain, rather than that it will probably rain. I believe that it might rain, which is to say that all the probability spaces in my belief set include rain as a possibility. That is, there is a probabilistic content—the set of spaces that include the proposition that it will rain as a possibility—that I believe, because my belief set is a subset of that content. I assert the very content I believe, and if it's accepted by my interlocutors, then the conversation will evolve accordingly; the set of probability spaces corresponding to the common ground will contain only spaces according to which rain is possible. And if my interlocutors believe what I say, then they too won't rule out the possibility of rain. While I've focused on the applications that are simplest to describe, Moss also provides sophisticated treatments of nested epistemic modals, and the interaction of conditionals, modals, and explicitly probabilistic language.

While Moss provides specific arguments against extant attempts to model the phenomena she treats using traditional contents, in my view you don't really need to be convinced by her negative case to find the positive case compelling. I remain agnos-

tic, leaning skeptical, as to whether one *could* provide a similarly powerful and unified treatment of the topics she addresses without appeal to probabilistic contents. But I'm thoroughly convinced that, once one uses probabilistic contents to model states of opinion—which is, essentially, what many Bayesian epistemologists have been doing for decades—there's no good reason why one shouldn't *also* expect those states of opinion to find expression in language. And so when we look and see a number of linguistic phenomena that can be straightforwardly accounted for in those terms, bending over backwards to account for them using traditional contents can look unmotivated and *ad hoc*.

So far I've focused on the role of probabilistic contents in belief and assertion, but haven't yet said much about the titular topic of the book: probabilistic knowledge. In the remainder of this review I'll try to give a sense for what Moss means by positing probabilistic knowledge, as well as some of its putative epistemological implications.

Bayesian epistemology is often contrasted with traditional epistemology. Part of the apparent divide is that the two approaches involve different ways of characterizing our states of opinion. The Bayesian epistemologist works with credences, while the traditional epistemologist works with beliefs and knowledge. Moss aims at a reconciliation, albeit one that may be more congenial to those that start out with Bayesian sympathies. For Moss, to have a credence *is* to have a belief—e.g., an imprecise but highish credence that it will rain *just is* a belief that it will probably rain. What does it take for such beliefs to amount to knowledge? Following Williamson, Moss is skeptical that we can provide an analysis of what it takes for a belief (probabilistic or not) to amount to knowledge. But there are informative things we can say—e.g., that lucky guesses don't amount to knowledge, that only truths can be known, etc. This latter claim turns out to be trickier to make in the case of probabilistic beliefs than traditional ones—ordinarily, we think of credences as more or less accurate, but not true or false. In Chapter 6 Moss provides a deflationist-inspired account of truth for probabilistic contents that is remi-

niscent of similar strategies that meta-ethical quasi-realists and expressivists have used to make sense of truth attributions for normatie; very roughly, if P has a probabilistic content, then the claim that it's true that P will have the same probabilistic content that P does.

So we end up with the picture that some credences amount to knowledge. A highish credence that it will rain, if formed in the right way (e.g., not by a lucky guess), and true (in at least the deflationist sense), can amount to knowledge that it will probably rain. This lets Moss appeal to ideas familiar from knowledge-first epistemology—e.g., that actions are only appropriate when based on knowledge—and apply them to debates that seem more naturally cast in probabilistic terms. She defends Paul (2014)'s claim that we can't rationally decide to undergo certain high-stakes transformative experiences, on the grounds that evidence about how others similar to us have reacted to such experiences doesn't give us probabilistic knowledge concerning how we'll react. She argues that legal burdens of proof should be understood as requiring probabilistic knowledge. Racial/ethnic/gender profiling is wrong in part because the probabilistic beliefs it produces don't amount to probabilistic knowledge. In general, the move is to take some situation in which our evidence seems to make some claim probable, but it doesn't seem right for that probability to guide action, and to argue that in such cases even high credences based on solid evidence fail to constitute (probabilistic) knowledge.

Before I close, I want to sound a note of skepticism regarding these applications. For reasons I haven't gone into, Moss requires a hefty dose of contextualism to make her account work (see especially chapter 7). For example, she wants to secure the attractive result that it can be true to say, of a fair lottery, that each ticket is a probable loser, but also that the winning ticket is not a probable loser. And for this to work, those two claims require a context-sensitive parameter to be filled in differently. But this kind of flexibility, while attractive, creates difficulties for some of the normative implications Moss hopes to draw out of her account. The same contextualism that makes it possible—in

certain contexts—to truly say that each ticket is a probable loser, also makes it possible to truly ascribe probabilistic knowledge—e.g., concerning criminal guilt—on the basis of merely statistical evidence. Moss attempts to address this problem by positing moral requirements on how context-sensitive probabilistic language should be interpreted. While this seems to me the right general form for a solution to take, I wasn't satisfied by the particular solution offered. Moss proposes a "rule of consideration", according to which it's often morally required to keep in mind the possibility that an individual may be a counterexample to statistical generalizations she falls under. While undoubtedly true, this struck me as less explanatory than one might hope. Sometimes we think it *isn't* morally required to do so—e.g., when we're comfortable convicting somebody on "direct" but nevertheless fallible evidence—and the rule doesn't help us decide how to classify tricky cases. And even if it could be fleshed out, it probably wouldn't help with the non-moral applications. E.g., relative to some contexts, statistical information *will* let me know know how I'll probably experience parenthood. So if Moss wants to vindicate Paul's pessimism about the rationality of undergoing a transformative experience, she'll need a different kind of explanation of what's wrong with those contexts, or why the knowledge ascriptions that can be truly uttered in them shouldn't guide action.

I rate these as minor quibbles. *Probabilistic Knowledge* offers a persuasive case for a radical revision of some foundational ideas about the nature of content and communication. While the implications of this revision for downstream debates about statistical evidence strike me as unclear, any such implications would be icing on an already rich and rewarding cake.

Bibliography

Moss, Sarah. 2012. "Updating as Communication." *Philosophy and Phenomenological Research* 85:225–248.

Paul, L.A., 2014. *Transformative Experience*. Oxford University Press.

Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics (New York Academic Press)* 9:315–332.

Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford University Press.