2019

From outside of ethics: Moss, Sarah. Probabilistic Knowledge

Timothy Smartt
The University of Notre Dame Australia, timothy.smartt@nd.edu.au

This book review in a scholarly journal was originally published as:
Smartt, T. (2019). From outside of ethics: Moss, Sarah. Probabilistic Knowledge. *Ethics: An International Journal of Social, Political, and Legal Philosophy, 129* (2), 430-438.

Original book review in a scholarly journal available here:
https://www.journals.uchicago.edu/doi/full/10.1086/700042

This book review originally published in *Ethics: An International Journal of Social, Political, and Legal Philosophy* https://www.journals.uchicago.edu/doi/full/10.1086/700042

Smartt, T. (2019). From Outside of Ethics: Moss, Sarah, *Probabilistic Knowledge. Ethics: An International Journal of Social, Political, and Legal Philosophy, 129*(2), 430-438. doi: 10.1086/700042

FROM OUTSIDE OF ETHICS

One of an occasional series of reviews of books outside the bounds of moral, political, legal, and social philosophy that may nevertheless be of deep interest to people working in those fields.

Moss, Sarah. *Probabilistic Knowledge.*
Oxford: Oxford University Press, 2018. Pp. 288. $54.00 (cloth).

I. FORMAL AND TRADITIONAL EPISTEMOLOGY

Contemporary epistemologists talk about our epistemic lives in two very different ways. According to one way of talking about belief, our fundamental doxastic attitudes are relatively simple: we believe a proposition, disbelieve it, or suspend judgment on it. According to a different way of talking about belief, one that has grown in prominence since the early twentieth century, our fundamental doxastic attitudes are more complex. Essentially, they come in degrees. The standard way of cashing this out is in terms of credences, whereby each attitude is assigned a real number between 0 and 1 which measures a subject's confidence that the content of the attitude is true.

These different ways of talking about belief are central to two different approaches to epistemology. On the one hand, traditional epistemology tends to traffic in outright belief (and its siblings), and on the other, formal epistemology tends to stick to credences.

Here's a serious problem for the formal approach: isn't epistemology the philosophical study of knowledge? If credences represent subjective degrees of confidence, then it doesn't seem like they are the sorts of attitudes that could count as knowledge. Knowledge is factive, and factive attitudes can only have true contents. It's natural to speak of beliefs being true, but it's strange to speak of credences being true. In this regard, it seems that beliefs but not credences can constitute knowledge. If one knows that Amy is in Amsterdam, it must in fact be the case that Amy is in Amsterdam. If one has .6 credence that Amy is in Amsterdam, it needn't be the case that Amy is in Amsterdam. Furthermore, credences unavoidably involve some degree of subjective uncertainty, which seems incompatible with knowledge. Whatever else one might think about knowledge, it seems odd to think that we might both know that Amy is in Amsterdam and be uncertain whether Amy is in Amsterdam.

By describing our epistemic lives in terms of credences, it has been difficult for formal epistemology to find a place for knowledge. This has resulted in the development of an approach that, while making important contributions to a number of epistemological debates, has had very little to say about the topic traditionally taken to be at the center of the subject.

Sarah Moss's *Probabilistic Knowledge* proposes a novel solution to this problem: credences can be knowledge.

A major attraction of Moss's view is that it promises to build new connections between formal and traditional epistemology. If credences can be knowledge, it becomes a lot easier to see how traditional theories of, say, knowledge and justification are of relevance to formal theories of, say, evidential support and belief revision, and vice versa. In this sense, *Probabilistic Knowledge* is groundbreaking. Insofar as one thinks that it would be desirable to have a unified account of core epistemological issues, Moss's proposal deserves serious consideration.

The view also promises to do important work beyond epistemology. Philosophers working in ethics should be especially interested in the book for two reasons. First, knowledge has practical consequences. If credences can be knowledge, a number of aspects of our moral, social, and political lives start to look different. Moss teases out a number of these applications in areas as diverse as racial profiling, standards of legal proof, and the diminishment of women's speech. Second, the book defends an account of epistemic language which draws deeply on expressivism. Metaethicists working on questions related to normative language will find much of interest in Moss's application of the expressivist program to the epistemological domain.

I'll briefly summarize the central arguments of the book and then turn to the relevance of probabilistic knowledge for moral, social, and political philosophy.

II. PROBABILISTIC KNOWLEDGE

The big-picture view is that credences can be knowledge. A .6 credence that Amy is in Amsterdam can constitute knowledge in the same way that an outright belief that Amy is in Amsterdam can.

One main argument for the view is that a certain conception of probabilistic content can play a number of underappreciated roles. Specifically, we can believe, assert, and know probabilistic contents. Many philosophers might agree with this, so long as the contents in question are, say, propositions about objective chances or evidential probabilities. As we'll see, Moss defends a very different proposal about these contents.

In a nutshell, Moss defends her three claims in the following way. In regard to belief, she develops a novel theory of credences whereby the content of one's credence is a set of probability spaces rather than a proposition. For example, the content of my .6 credence that Amy is in Amsterdam is the set of probability spaces containing just those spaces that assign a probability of .6 to the proposition that Amy is in Amsterdam. In regard to assertion, Moss develops a nonstandard theory of assertion according to which the content of many sentences employing epistemic vocabulary is a set of probability spaces rather than a proposition. Her theory of assertion has much in common with recent expressivist semantics about epistemic vocabulary but claims to avoid a number of objections raised against those views. In regard to knowledge, she draws on the idea that we can believe and assert probabilistic contents to develop a number of arguments that various features of knowledge obtain in just the same way for credences with probabilistic contents as for outright beliefs with propositional contents. The three claims

are independently motivated and do not depend on one another. However, when taken as a package, they present a strong case for new work that can be done by a theory of probabilistic content.

It is worth briefly unpacking Moss's concept of probabilistic content. Moss stipulates that a probabilistic content is a set of probability spaces. A probability space is a mathematical model of a possible scenario consisting of an ordered triple: a domain of possibilities, an algebra of propositions containing those possibilities, and a probability function defined on the elements of that algebra. Since a probability space determines a probability function, informally we might say that a probabilistic content corresponds to a set of probability functions, which in turn corresponds to a constraint on precise credence distributions. So the less mathematically inclined reader can simplify talk of probabilistic contents to talk of properties of one's credence distribution. This simplification is not perfect, but it preserves much of the view that will be of interest for present purposes.

Let's look at each of Moss's three claims in more detail.

A. Belief

It is uncontroversial that, as a matter of folk psychology, we have probabilistic beliefs. These can be simple, such as a belief that a coin is equally likely to land heads as tails. They can be more complex, such as a belief that it's possible but not likely that if Amy is in Europe then Amy is in Zürich. According to a standard convention, the two examples can be formally represented as credences in the following way: a .5 credence that the coin will land heads, and a conditional credence below a certain threshold (perhaps .5) that Amy is in Zürich.

Transposing descriptions of beliefs from natural language constructions into the formal notation of credences relies on some theoretical assumptions. These have not always been made as explicit as one might want, which has resulted in a debate about how exactly credences should be interpreted.

Moss helpfully explains the fundamental distinction within this debate and advances a unique position.

The distinction is between interpreting credences as complex attitudes with simple contents and interpreting them as simple attitudes with complex contents. Each interpretation locates the fundamental difference between credences and outright beliefs in a different putative feature of credences. On the one hand, according to the complex attitude account, credences differ from outright beliefs by involving a more complex attitude toward the content of belief. They involve partially believing, or believing to a graded degree of confidence, or taking up some other belief-like attitude. On the other hand, according to the complex content account, credences differ from outright beliefs by having a content which is more complex than the propositional content of outright beliefs.

The two accounts provide subtly different interpretations of the slogan "belief comes in degrees." The complex attitude account builds the gradation into the attitude of believing itself, whereas the complex content account puts the gradation in the content of the belief.

By way of illustration, consider the example of my .6 credence that Amy is in Amsterdam. According to the complex attitude interpretation, I have a complex attitude toward a simple content: I partially believe to degree .6 that Amy is in

Amsterdam. According to the complex content interpretation, I have a simple attitude toward a complex content: I believe that it is .6 likely that Amy is in Amsterdam.

Moss rejects the complex attitude account and defends a unique version of the complex content account. As mentioned earlier, the crux of Moss's account is that the content of a credence is not a propositional content but a probabilistic content, cashed out as a set of probability spaces. The complex attitude account is widely presupposed, so her choice is significant.

Moss presents a number of arguments to support her theory of credences, all of which emphasize the elegance and simplicity of the complex content account. Her central argument concerns the role played by the contents of belief. It is common to identify the contents of belief with whatever objects play several important theoretical roles. Moss picks out four roles and argues that probabilistic contents can play all these roles. The four roles are (1) explaining rational action; (2) grounding relations of agreement and disagreement between different subjects; (3) grounding facts about a subject's beliefs over time, such as whether they have changed their mind; and (4) grounding rational relations between the elements of a subject's belief set at a time. Moss argues that there are prominent credence-based theories of all these topics which receive their most plausible articulation when interpreted according to her preferred complex content account.

For instance, consider theoretical role (2): the contents of belief ground relations of agreement and disagreement between different subjects. If I have .6 credence that Amy is in Amsterdam and you have .2 credence that Amy is in Amsterdam, it seems that we disagree. However, according to the complex attitude account, we have attitudes toward the same content. That is, we both have attitudes toward the simple content that Amy is in Amsterdam, but we adopt different complex attitudes toward this content. On this interpretation of credences, it seems that the contents of belief are unable to play theoretical role (2).

Moss doesn't belabor this point, but I think this presents a significant problem for the complex attitude account. A friend of the complex attitude account might reply that the attitudes themselves are enough to ground disagreement. That is, the difference between believing to degree .6 and believing to degree .2 is enough to provide us with the result that we disagree. But this is implausible. It's not clear that subjects who adopt different attitudes toward the same content necessarily disagree. For instance, if I outright believe that Amy is in Amsterdam and you hope that Amy is in Amsterdam, I don't think it's accurate to say that we disagree about whether Amy is in Amsterdam. It doesn't seem that we can ground agreement and disagreement in the attitudes themselves.

On the other hand, according to Moss's complex content interpretation, we believe fundamentally different probabilistic contents. I believe that it is .6 likely that Amy is in Amsterdam, and you believe that it is .2 likely that Amy is in Amsterdam. On her preferred interpretation of credences, our intuitive disagreement is explained simply by our believing inconsistent contents. Moss proposes similar arguments for the other roles traditionally assigned to contents of belief, showing not only that probabilistic contents can play these roles but also that credence-based theories of rational action, interpersonal disagreement, and intrapersonal rational requirements on credences are most plausible when interpreted according to her theory.

Chapter 1 goes into further detail on both the complex content and complex attitude interpretations of credences. This discussion is an especially valuable section of the book, and philosophers sympathetic to a complex attitude interpretation, or skeptical of the philosophical or psychological plausibility of credences whatsoever, will find a number of challenging new arguments to consider. The chapter will also make an excellent resource for credence-curious newcomers interested in exploring the idea for the first time, since Moss illuminates a number of philosophical issues which affect how one thinks about the objects of credences that tend to be passed over too quickly in much of the work on this subject.

B. Assertion

Probabilistic Knowledge argues that we can believe, assert, and know probabilistic contents. Chapters 2, 3, and 4 make the case that we can assert probabilistic contents.

Moss argues that the contents of assertion can be sets of probability spaces. This departs from the standard theory which holds that the contents of assertions are propositions. The application of the standard theory to epistemic vocabulary has been criticized by recent work on epistemic modals and probability operators, which has argued that sentences containing epistemic vocabulary need not assert propositions (see, e.g., Seth Yalcin, "Epistemic Modals," *Mind* 116 [2007]: 983–1026; and Eric Swanson, "How Not to Theorize about the Language of Subjective Uncertainty," in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson [Oxford: Oxford University Press, 2011], 249–69).

Moss follows these writers in holding that at least some assertions have probabilistic contents. However, the existing arguments for this conclusion are controversial. Moss proposes several new arguments, each of which makes the case that assigning probabilistic contents to assertions provides a simpler and more unified account of many fundamental communication practices. For instance, it allows us to account for the successful communication of probabilistic beliefs in the same way that we account for the communication of outright beliefs. Moreover, if you accept Moss's theory of credences, there is a strong motivation to accept her view about the contents of assertion, since it would be problematic if we were unable to assert some of the contents that we believe.

Moss goes into more detail than I can do justice to here, including outlining a formal probabilistic semantics for epistemic modals, probability operators, and indicative conditionals. These chapters will be especially of interest to philosophers of language working on the language of subjective uncertainty and epistemic modality. The upshot is a systematic defense of the claim that many sentences involving epistemic vocabulary are often best understood as asserting probabilistic contents.

C. Knowledge

The centerpiece of the book is chapters 5–7, which make the case that the probabilistic contents of credences can constitute knowledge. Throughout these chapters, Moss's chief strategy is to demonstrate that credences can behave just like traditional instances of knowledge. Covering a remarkable amount of ground, Moss

lays out a thorough case that there is no principled reason to treat opinions with propositional content as knowledge but not opinions with probabilistic content. In doing so, Moss relies on only the most well-established features of knowledge.

Moss devotes much of chapter 5 to presenting a number of arguments for the conclusion that we can acquire credences with probabilistic contents in structurally identical ways to how we acquire knowledge of propositional contents, namely, through testimony, perception, inference, memory, and a priori thought. When we come to believe propositional contents in these ways, in many instances the outright beliefs constitute knowledge. Moss argues that when we come to believe probabilistic contents in these ways, in many instances the credences amount to knowledge too. Although Moss is clear about how radical her view is, throughout these chapters she draws quite successfully on work from different corners of philosophy to suggest that a nascent conception of probabilistic knowledge has already done some work. Moss argues for this particularly persuasively in regard to recent work on perceptual experience.

Another argument Moss advances for probabilistic knowledge is that credences can be "Gettiered"—that is, a subject's credence can be justified yet epistemically deficient owing to the presence of intervening or environmental luck.

Consider this case described by Moss: "Alice enters a psychology study with her friend Bert. As part of the study, some participants are injected with a heavy dose of adrenaline, while others are injected with a saline solution. All participants are then sent to meet their friends. Alice is not told anything about the nature of the injection or the experiment. As it happens, Alice receives the adrenaline injection. As she meets Bert, Alice reflects on her fluttering nerves and comes to have high credence that she finds Bert attractive. And indeed, she probably does find Bert attractive" (102). In the literature on the Gettier problem, the standard way of explaining what's epistemically wrong with luck is that it is incompatible with knowledge. That is, in traditional Gettier cases, subjects have justification but lack knowledge. Moss argues that we should accept the same conclusion in credence-based Gettier cases. We can explain the epistemic deficiency of merely justified credences, such as Alice's, in the same way that we explain the deficiency of merely justified beliefs: they fail to constitute knowledge. This parallel strengthens Moss's case that credences behave just like traditional instances of knowledge. Just as credences can be acquired in the same way we traditionally acquire knowledge, so too credences can be epistemically defective in the same way beliefs traditionally fail to constitute knowledge.

The absence of luck is widely taken to be a necessary condition on knowledge. Various other conditions have been proposed by prominent theories of knowledge. Such theories have received much attention in traditional epistemology for a number of decades. Moss argues that all the conditions proposed by these theories can be applied equally well to credences as outright beliefs. She examines a number of popular conditions for knowledge—including reliabilist conditions, a sensitivity condition, warrant conditions, and skill-based conditions—and argues that some probabilistic contents can satisfy each condition. Indeed, she suggests that we ought to endorse one of these theories if we judge that it provides the most plausible conditions for probabilistic knowledge, just as we would be inclined to endorse a theory were it to provide the most plausible conditions

for propositional knowledge. In showing how these theories might accommodate credences, Moss strengthens her case that probabilistic knowledge can narrow the divide between formal and traditional approaches to epistemology.

Part of the persuasiveness of probabilistic knowledge is the way that Moss demonstrates the sheer usefulness of such a concept, and in these chapters she draws out a number of implications. She uses her account to develop a new probabilistic knowledge norm for belief, whereby one's credences should constitute knowledge. She proposes two new knowledge norms for action which spell out the connections between probabilistic knowledge, reasons for action, and decision-making. She also shows how probabilistic knowledge provides new ways to think about the difficulty of decision-making about transformative experiences.

Before we move on, there is one objection to probabilistic knowledge that is worth considering. As I mentioned earlier, knowledge is factive, and it doesn't seem obvious that probabilistic beliefs are truth apt. Can probabilistic knowledge satisfy a factivity condition? Moss devotes chapter 6 to arguing that it can, relying on a deflationary theory of truth to do much of this work. According to Moss's theory of truth, my .6 credence that Amy is in Amsterdam is true just in case it is .6 likely that Amy is in Amsterdam. This aspect of her account has clear parallels with how expressivists in ethics maintain that moral attitudes might be true.

In endorsing a deflationary account of truth, Moss explicitly rejects a straightforward sense in which credences might be true, namely, by matching evidential probability facts or the objective chances. This is a significant choice, as it parts ways with how some theorists have suggested that we develop a probabilistic analog of truth conditions.

To believe, *pace* Moss, that credences are made true by objective chances or evidential probabilities is to believe that there is more of philosophical interest to the connection between our probabilistic beliefs and the world than given space for by Moss. I think we can ask for more details here. For instance, consider two subjects who have the same .6 credence that Amy is in Amsterdam. Let's further suppose that the credences satisfy a number of Moss's conditions for knowledge, that is, they are acquired in an appropriate way, justified, non-"Gettiered," sensitive, and so forth. In this case, we should say that the two people know that it is .6 likely that Amy is in Amsterdam. According to Moss's deflationary theory of truth, the credences are true if it is .6 likely that Amy is in Amsterdam. Now let's further suppose that one of these people visits a wizard and becomes instantly omniscient. It seems that from the perspective of our newly all-knowing subject, she should resist saying that her friend knows that it is .6 likely that Amy is in Amsterdam, since she herself knows whether Amy is in Amsterdam. The unhappy upshot is that on Moss's theory of truth, it seems that we can know contents that we wouldn't know if we were omniscient.

To my lights, Moss's account of the way in which probabilistic knowledge satisfies factivity highlights that there are some gaps to be filled in the account. I suspect that many will find the concept of probabilistic knowledge only as strong as its account of truth. Although the deflationary account might work as a formal theory of truth, it seems reasonable to hope for more in terms of a metaphysical story about how credences might be true. Such a story would develop a more positive proposal of precisely what we know when we have probabilistic knowledge.

Developing probabilistic knowledge along these lines would, I suspect, help the concept be of greater use to a wider audience of philosophers.

III. IMPLICATIONS FOR ETHICS

Knowledge plays a substantive role in our practical lives. In the final chapters of the book, Moss picks out three topics to demonstrate that probabilistic knowledge is of importance to moral, social, and political philosophy. The topics are women's speech, standards of legal proof, and practices of racial profiling. I'll focus on just the application to legal proof.

Probabilistic knowledge opens up the possibility of understanding legal proof in general as requiring knowledge. Drawing on interdisciplinary literature, Moss demonstrates that when writers attempt to define legal proof they often come awfully close to defining knowledge. For instance, the following conditions have all been proposed as necessary for legal proof: sensitivity to the truth, incompatibility with luck, capable of serving as a reason for action, reliably safe from error, and truth. However, despite these motivations, accounts of legal proof that require knowledge have faced the problem that the standards for legal proof can be set in such a way that they do not seem to be compatible with knowledge. For instance, in many legal systems civil matters require proof by a balance of the probabilities, and criminal matters require proof beyond a reasonable doubt. It has seemed that proof by a balance of the probabilities can't require knowledge, since it seems to merely require that the judge or jury have a credence greater than .5 in the verdict.

Moss's notion of probabilistic knowledge is perfectly suited to solve this problem. She argues that probabilistic knowledge provides us with a unique way of defending a knowledge requirement for legal proof. Moreover, she puts this to work by developing new interpretations of many different standards of legal proof. Moss proposes that each standard of proof corresponds to a certain probability threshold. A standard is met when a judge or jury knows a probabilistic content which meets or exceeds this threshold. For example, a defendant is proved liable by a preponderance of the evidence only if the judge or jury has greater than .5 credence that the defendant is liable, and that credence constitutes knowledge. Likewise, a defendant is proved liable by clear and convincing evidence only if the judge or jury knows an even stronger probabilistic content, and guilty beyond a reasonable doubt only if the judge or jury knows a still stronger probabilistic content. Moss further argues that her probabilistic knowledge requirement for legal proof provides us with new ways to identify legal shortcomings of merely statistical evidence, one of which is that in many cases such evidence might justify a credence which nevertheless fails to constitute knowledge.

The applications Moss makes to moral, social, and political philosophy are creative and insightful. These need only be the beginning. For instance, probabilistic knowledge might provide act utilitarians with a new defense against the objection that subjects cannot know which act will lead to the best outcome. Likewise, probabilistic knowledge might provide a way for credences to play a much greater role in the philosophy of action, especially in work which emphasizes the relationship between knowledge and intention.

IV. CONCLUSION

Probabilistic Knowledge is an important and deeply original book. It will no doubt find a wide readership among epistemologists of different stripes, philosophers of language, and philosophers working on expressivism in various domains. It deserves to find a wider readership still. Moss's excellent discussion of the interpretation of credences should be required reading for students and researchers new to the subject, and her concept of probabilistic knowledge has the potential to reframe our thinking about how knowledge and uncertainty affect our practical lives.

Acknowledgment: In preparing this review I benefited greatly from participating in a master class on *Probabilistic Knowledge* at King's College London, instructed by Sarah Moss and organized by Julien Dutant, and from generous comments from Mark Colyvan, Brian Hedden, and Sarah Moss.

Tim Smartt
The University of Notre Dame Australia and the University of Sydney