Sarah Moss, *Probabilistic Knowledge*

New York: Oxford University Press, 2018. xi + 268 pp.

In her wide-ranging book, *Probabilistic Knowledge*, Sarah Moss presents a unified account of probabilistic content in theories of belief, assertion, and knowledge. The first part (chaps. 1-2.3) begins by arguing for a particular way of incorporating credence into an extant picture of belief—that is, by analyzing credence as a relation between an agent and a probabilistic content. The second part (chaps. 2.4-4) moves on to language, with a focus on epistemic vocabulary, indicative conditionals, and logical connectives. The third (chaps. 5-10) turns to knowledge. Here, the central thesis is that all grades of Bayesian credence—traditionally considered as mere partial belief—can in fact constitute knowledge in many of its central epistemological roles. Additional chapters turn to applications, including action and decision (chap. 9) and the role of statistical evidence in the law (chap. 10).

In each case, the role played by individual possible worlds on standard theories—as well as some of the roles played by worlds in compositional semantic treatments of logical operators, epistemic modals, and conditionals—is played in Moss's book by *probability spaces:* triples $\langle \Omega, F, m \rangle$ consisting of a sample space, a sigma-algebra, and a probability measure. *Sets* of probability spaces are Moss's choice for modeling content in the belief-assertion-knowledge trifecta, as sets of (mere) worlds (viz., propositions) model content in the theories she seeks to upgrade. Content satisfaction—for example, the content that it is .4 likely to rain—goes by *containment:* an agent's probability space satisfies this content iff it is contained in the set of *all* probability spaces that assign .4 to the proposition that it will rain. *Uptake* goes by simple

Thanks to Josh Dever, Dorothy Edgington, Michael Nielsen, and Rachel Rudolph for helpful discussion of this material.

intersection: perceptual content, for example, is intersected with one's probability space or representor; asserted content is intersected with an upgraded version of the common ground.

The book is written in a clear and accessible style, and while the overall picture coheres around the upgrade to probabilistic content, arguments in the language, mind, and knowledge strands are to a large degree self-contained.[1] Below, I focus on a selection of the main arguments, with an eye to issues that have already begun, and most likely will continue, to provoke the most discussion.

Language: Thorough Probabilism and Partition-Sensitivity

We begin with the thesis that assertions like (1)

(1) It is .9 likely that it is raining.

typically express, and aim to engender coordination on, credence (here, .9 credence) in the proposition *it is raining*.

At a first pass, Moss takes on board the equation (in Dummett's terms) between assertoric content and ingredient sense, and thus holds that the compositional semantic value of (1) in context is a probabilistic content of the kind described above: the set of probability spaces assigning .9 to *it is raining*.

This first-pass analysis does not eliminate a need for the more traditional notion of worldly content as the object of credence. The compositional contribution of the prejacent "that

[1] Moss makes some recommendations for selective paths through the material in the introduction (x).

it is raining" to (1) is a possible-worlds constraint on how things might be: that is, it is a (mere) proposition.

(1) thus contrasts with (a noncollapsed reading of) a sentence like (2):

(2) It is .9 likely that it might be raining.

which is assertable, Moss argues, in e.g. a case where there is a one-tenth chance we will get evidence that rules out that it is raining. In (2), the compositional contribution of the prejacent under "likely"—since it includes "might"—seems itself to be probabilistic. Hence there's a puzzle for the compositional semantics: what type of object—a set of worlds or a set of probability spaces—best models the argument of "it is .9 likely"?

Moss's response is to generalize to the worst (or, at least, highest-type) case: content embedded under, for example, "likely" is *always* probabilistic, though sometimes only trivially (or "nominally") so.[2] Returning to (1), Moss posits the presence of a hidden type-lifting operator '*C*' (think: *certainly*) with scope over the prejacent, yielding (1') as the true LF of (1):

(1') it is .9 likely that *C*(it is raining)

Now, a compositional semantics operating on (1') must assign a numerical probability (.9) to something that is itself epistemic: "certainly"-claims are accepted or rejected by probability spaces rather than individual worlds. This is "thorough probabilism" at work: semantic values are probability space-sensitive "all the way down". This choice lends itself to the general formal

[2] Another way to go is to hold that the compositional semantic values are (at least) world-probability space *pairs;* see, for example, Yalcin 2012 and MacFarlane & Kolodny 2010 for world-information state pairs playing the relevant compositional role.

picture of possibility, high probability, and *n*-valued credence sketched in Moss's appendix: a proposition's being (epistemically) possible is its being possibly *certain*; a proposition's being highly probable is its being highly probably *certain*, and a proposition's being assigned a credence *n* is the value *n*'s being associated with the region of the sample space throughout which that proposition is *certain*.

Beyond the type-shifting operator '*C*', there is additional syntactic enrichment of LFs in Moss's picture, in the service of a very flexible brand of contextualism. On her semantics, epistemic operators, conditionals, and logical operators carry hidden indices which, relative to $g_c(.)$, the contextually available assignment function, denote questions (formally, Hamblin (1973)-style partitions of logical space or Roberts (1996)-style "Questions under Discussion"). These partitions isolate some important consideration of interest in evaluating a probabilistic object-language claim. Different indices on epistemic operators correspond to cases in which different questions influence the communicated content of, for example, indicative conditionals (for which Moss favors a strict conditional analysis.)

As an illustration, consider the famous Sly Pete case of Gibbard (1981). In a game of poker, Pete has a worse hand than his opponent does. But by cheating, Pete has also gained knowledge of this fact, and hence will fold. Different indices *j* and *k* such that $g_c(j)$ = *did Pete cheat?* = {*Pete cheated, Pete didn't cheat*} and $g_c(k)$ = *did Pete have a winning hand?* = {*Pete had a winning hand, Pete did not have a winning hand*} on "likely" support the apparently dueling takes on the embedded conditional "if Pete called, he won/lost" in (3) and (4):

(3) It is likely$_j$ that if$_i$ Pete called, he won.

(4) It is likely$_k$ that if$_i$ Pete called, he lost.[3]

Formally, (3) is true because within {*Pete cheated, Pete didn't cheat*}, the agent considers the former to be likely. The region of her sample space that's settled on Pete's having cheated supports the *strict* conditional "if Pete called, he won". However—turning to (4)—within {*Pete had a winning hand, Pete did not have a winning hand*}, the agent considers the latter to be likely. The region of the sample space that's settled on Pete's having the weaker hand supports the strict conditional "if Pete called, he lost". Hence given the different indices, the *prima facie* incompatible (3) and (4) can both characterize the same credal state in the same context.

Knowledge: Closure and Impersistence

In epistemology, Moss's high-level thesis is that all degrees of Bayesian credence can constitute knowledge. Offered in support of her claim are the contention that states of credence are factive (given a minimalist construal of what "factive" amounts to), are Gettierizable, and play the knowledge role with respect to rational action.

The constitution thesis is simple. But it is also radical in a number of ways—for example when it comes to reasoned change in view. Extreme credences—of 0 or 1—are *persistent* in a Bayesian framework. But intermediate values can fluctuate. Indeed, we may be certain they will, such as when we anticipate observing the outcome of an experiment. As Moss puts it, "When you have a .6 credence that Jones smokes, you necessarily also believe it *might* be that

[3] Though (3) and (4) differ with respect to the index on "probably", I give them the same index, *i*, on the inner "if", as Moss does (236). I also assume, as she does, that indices on "if...then" constructions tend to be *decisive* with respect to the antecedent—viz., in (3) and (4), decisive with respect to whether Pete called (73).

Jones certainly *doesn't* smoke" (138).[4] If intermediate-valued credence can constitute knowledge, it is knowledge that has this changeling character built in.

Such a view contrasts with a more restricted, "Meno-inspired" way of conceptualizing knowledge for Bayesians, wherein the *stability* of extreme-valued credences in the light of new information is what makes them (uniquely) knowledgelike. On this view, my expressed .6 credence that (e.g.) *the dice will land even* can't be knowledge, because in a moment—after observing how the dice land—I will know something incompatible with that very claim: either that they certainly came up even or that they certainly didn't. As Gardenfors (1982: 749) puts it:

> A fundamental presupposition of the Bayesian doctrine is that, if a sentence *B* is accepted as **known** in a given state of knowledge and if a sentence *A* is consistent with what is accepted as **known**, then *B* should be accepted as **known** in the revised state of belief which is obtained by adding *A* as a new piece of knowledge. [emphasis added]

This claim is false on Moss's picture, instantiating *B* with "it is .6 likely that the dice land even", *A* with "the dice certainly don't land even", and glossing "*A* and *B* are consistent" as entailing "the claim '*B* and might *A*' is consistent."

Moss is, of course, aware of this "changeling" feature of credence and its ramifications for constitution thesis. Her response—in addition to motivating cases where epistemic contradictions like "it's .6 likely that *p* and it might be that certainly not-*p*" can sound true to ordinary speakers—is to follow an analogy with contextualist responses to the raising-to-salience of skeptical scenarios (chap. 7.3). The idea is that the presence of some incoming information—such as an impending observation of how the dice land—can act as a *challenging alternative* to the epistemic status of intermediate credence, in much the same way that raising possibilities

[4] The claim "it might be that Jones certainly doesn't smoke" is true in Moss's object language when "might" is indexed to a partition which is decisive with respect to the proposition that Jones smokes. As Moss notes (132), this makes the statement about Jones a type of *epistemic contradiction* (Yalcin 2007): it entails a sentence of the form *p and might not-p*.

regarding evil demons can challenge ordinary knowledge about, for example, having hands. It may generally be true in such skeptical contexts that knowledge "vanishes". But (the argument goes) that does not mean that knowledge fails to be present in unchallenged cases. Returning to the case of intermediate credence, an impending news bulletin, the presence of a more perfectly informed expert (138) or active consideration of the omniscient credences of God (141) may all successfully disqualify intermediate credence from knowledge-like status. Yet, Moss suggests, this is no argument that intermediate credence *generally* fails to be knowledge.

The gambit, of course, is unusual in casting situations where experts are present or observations are about to be made—intuitively, situations in which extra knowledge is there to be *gained*—in the role of situations where knowledge *vanishes* under skeptical pressure. But such a view may be inevitable given the broader picture at work in the book. As Moss's compositional semantics is probabilistic all the way down, so knowledge is probabilistic through and through: knowers are afloat on a fully Bayesian sea.

Melissa Fusco
Columbia University

References

Dummett, Michael (1973). *Frege: Philosophy of Language*. Harvard University Press.

Gardenfors, P. 1982. "Imaging and Conditionalization". *Journal of Philosophy* 79(12): 747-760.

Gibbard, Allan (1981). "Two Recent Theories of Conditionals." In W. Harper, R. Stalnaker, and G. Pearce (eds.) *Ifs: Conditionals, Belief, Decision, Chance, and Time*, D. Reidel, Dordrecht.

Hamblin, C. L. (1973). "Questions in Montague English." Foundations of Language, 10(1): 41–53.

MacFarlane, John, and Niko Kolodny (2010). "Ifs and Oughts." *Journal of Philosophy*, 107: 115–143.

Roberts, Craige (1996). "Information Structure in Discourse." In J. Yoon, and A. Kathol (eds.) *OSU Working Papers in Linguistics*, Ohio State University, vol. 49: 91–136.

Yalcin, Seth (2007). "Epistemic Modals." *Mind*, 116: 983–1026.

Yalcin, Seth (2012). "A Counterexample to Modus Tollens." *Journal of Philosophical Logic,* 41: 1001-1024.