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

Sarah Moss
ssmoss@umich.edu

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 modals and probability operators, my semantics for indicative conditionals, and the factivity of probabilistic knowledge. In what follows, I discuss and answer her concerns in this same order.

1.1 Credences

Even dogs have probabilistic beliefs. At dinnertime, Fido sits near the children, not the adults, because he believes that the children are more likely to drop food onto the floor. Edgington worries that this fact makes trouble for my complex content account of probabilistic belief. As she puts it, believing a probabilistic content appears to require more "self-consciousness and conceptual apparatus" than believing a proposition, and Fido "doesn't have the concept of probability" that seems required to believe a probabilistic content.

However, Edgington is mistaken to think that you must have a concept of probability in order to believe probabilistic contents. When it comes to concept possession, the complex content account is no more demanding than traditional accounts of belief. You do not need to grasp the concept of a possible world in order to believe the traditional content represented by the set of worlds where Jones smokes. For just the same reason, you do not need to grasp probability concepts in order to believe the

1. Thanks to Catrin Campbell-Moore, Dorothy Edgington, Scott Hershovitz, Jason Konek, Carlotta Pavese, and Eric Swanson for helpful discussion.

content represented by the set of probability spaces that assign .6 to Jones smoking. Whether the theorist models ordinary beliefs using possible worlds or using probability spaces, subjects can have ordinary beliefs without grasping the concepts that the theorist uses to talk about them.

A second worry raised by Edgington is that if belief does not come in degrees, then "we are left in the dark about what credence is." The idea here, I take it, is that we are all familiar with what it is to fully believe something. On the complex attitude account that I reject, credence can be explained in terms of the degree to which you have this familiar attitude towards a proposition. Edgington worries that if belief does not come in degrees, then we will be unable to elucidate credence in terms of belief.

However, I don't think we should aim to elucidate credence in terms of belief to begin with. Say you flip a fair coin. As you flip it, you do not believe that the coin will land heads. When it comes to the attitude of full belief, you simply *do not have* that attitude towards the proposition that the coin will land heads, even a little bit. The same can be said when you flip a coin that you know to be biased towards tails. Having .5 credence that a coin will land heads, or .3 credence, or .2 credence—these attitudes are totally unlike the attitude of fully believing that the coin will land heads. Credence is not a measure of the degree to which you have or lack the attitude of fully believing a proposition.[2]

Rather than trying to elucidate credence in terms of belief, we should elucidate credence by reference to its role in rational decision making. There are several ways that one can go about this. Some theorists say that disposing an agent to take certain bets is merely part of the functional role of credences, whereas others consider the link to be definitional. Whatever exactly we say about the relation between credences and decision making, it will be just as easy to say it on the complex content account as on the complex attitude account.

A third worry raised by Edgington is that my account distinguishes contents of belief from contents of other degreed attitudes—such as desire, hope, and fear—and that this makes trouble for certain models of practical reasoning. For example, say you want to get a can of soda, and you believe that if you open the fridge, you will get one. According to a simple belief-desire model of practical reasoning, this explains why you should intend to open the fridge. But if contents of belief are sets of probability spaces, whereas contents of desire are sets of worlds, then it isn't clear how these contents can be combined in one seamless model of practical reasoning.

2. A terminological point: Edgington suggests that on my account, your .3 credence that Jones smokes is a "full belief in a probability judgement." But on my account, such credences are not full beliefs. Full beliefs are beliefs expressed without using epistemic vocabulary—such as, for instance, the belief that Jones smokes. As I explain in §3.6, 'full' is used to regulate pragmatic slack in belief ascriptions. For further discussion, see Moss (2019a), where I argue that we should elucidate belief in terms of credence.

As I see it, though, practical reasoning fundamentally involves credences and utilities, not propositional beliefs and propositional desires. On the account of simple belief ascriptions developed in §3.6 of the book, 'Smith believes Jones smokes' says that for relevant purposes, it is as if Smith has the highest possible credence that Jones smokes. Analogously, 'Smith desires to win the race' says that for relevant purposes, it is as if Smith assigns the highest possible utility to winning the race.[3] Both simple belief ascriptions and simple desire ascriptions are loose speech. Belief-desire reasoning is best understood as a rough-and-ready shorthand for credence-utility reasoning, which does not depend on credences and utilities having the same sort of content.[4]

Edgington says that in decision theory, "it is important that the same content can be assigned a credence and a measure of desirability." But in fact, what is important for decision theory is merely that the same *object* can be assigned a credence and a utility—not that this object is identified as the *content* of a credence, or as the content of a desire. Suppose you assign an extremely low negative utility to the proposition that Jones smokes. That does not mean that you have a degreed desire that has the proposition *that Jones smokes* as its content. The complex content says just the same thing about credences: when you assign .3 credence to the proposition that Jones smokes, that does not mean that you have a degreed belief that has the proposition *that Jones smokes* as its content. This account of belief contents does not preclude credence and utility functions from playing their usual role in decision theory.

1.2 Epistemic modals and probability operators

Edgington suggests that the belief that Jones might smoke is "represented by the set of probability spaces which give at least a certain minimum positive value to the probability that Jones smokes." As a preliminary point, I should mention that this is not my view. Suppose you are throwing a point-sized dart at a dartboard. You can believe that you might hit the bullseye without assigning positive probability to that proposition. As I explain in §2.4 of the book, the content of your belief that you might hit the bullseye is a set of probability spaces that are distinguished not by the probabilities that they assign to propositions, but rather by the fact that their domains intersect some relevant possibility that accepts that the dart hits the bullseye.

3. See Moss 2019d for further development of this account of simple desire ascriptions.
4. What about contents of other attitudes and speech acts? Wishing is relevantly similar to desiring. Hope and fear are traditionally taken to be amalgamations of belief and desire, and hence should ultimately be understood in terms of both credence and utility. Edgington also objects that questions do not have probabilistic contents. However, we do ask probabilistic questions. *How likely is it that Jones smokes? Who might come to the party?* An adequate semantics for such questions will be probabilistic in nature.

For Edgington, the main worry for my account of probability operators is that it overgenerates readings of sentences such as the following:

(1) P might not be probable.

For example, Edgington argues that if we set aside objective readings of probability operators, the following inference is invalid:

(2) a. It is .6 likely that P is .4 likely, and it is .4 likely that P is .8 likely.
b. Hence P might not be probable, since P might be merely .4 likely.

According to Edgington, one must infer as follows:

(3) a. It is .6 likely that P is .4 likely, and it is .4 likely that P is .8 likely.
b. Hence it is false that P might not be probable, since P is .56 likely overall.

Insofar as (2) appears to be valid, the conclusion must be expressing some opinion about objective chance facts, or facts concerning some other probability function. Embedded probability operators are not used to denote thoroughly probabilistic contents.

Edgington's proposal about embedded probability operators is a close cousin of the shifty proposal about epistemic modals discussed in §3.3 of my book, according to which every epistemic expression embedded under epistemic vocabulary is interpreted relative to some contextually determined body of evidence. The arguments that I give against the shifty proposal apply equally to Edgington's account of (2). In a nutshell, the main idea of §3.3 is that insofar as we have good reasons to adopt a thoroughly probabilistic semantics for unembedded probability operators, we have a similar range of reasons to adopt a thoroughly probabilistic semantics for embedded probability operators, such as those in (2).

In addition to the arguments that I give in the book, there is a further reason to reject semantic theories that invariably collapse nested probability operators into all-things-considered probability judgments. It is widely agreed that semantically speaking, expressions such as *probably* and *probable* have a lot in common with modal adverbs and adjectives such as *intermittent, prevalent, occasionally, usually*, and *frequent*. To learn more about the former expressions, then, it can be helpful to examine the latter. And on examination, it is not hard to find non-collapsing readings of nested adverbial phrases:

(4) That grinding noise is sometimes intermittent, and sometimes constant.

(5) This weed is prevalent in some locations, and sparse in others.

(6) His whining is occasionally frequent, but usually infrequent.

An especially contrarian philosopher might reject (4), insisting, "If the grinding noise is only sometimes constant, then that just means it's *intermittent*. It's not *sometimes intermittent* and sometimes not." But we can and should interpret (4) more charitably than this. It is not hard to figure out what (4) means: 'sometimes' and 'intermittent' are being used to talk about different temporal partitions, one more fine-grained than the other. The same goes for (5), (6), and other uses of non-collapsing adverbial phrases. These observations about non-modal adjectives support my semantics of epistemic vocabulary. On my semantics, nested epistemic modals and probability operators are similarly interpreted relative to different partitions, one more fine-grained than the other. Just as speakers can use 'sometimes intermittent' to express a useful belief in the right sort of context, speakers can use '.6 likely to be .4 likely' to express a useful probabilistic belief. For instance, if one expert says that P is .4 likely and another says that P is .8 likely, and you know that the first expert is .6 likely to be correct, then you may well conclude that P is .6 likely to be .4 likely.

To sum up, my semantics for probability operators is flexible enough to generate the valid reading of (3) that Edgington describes. In addition, my semantics generates readings of nested epistemic operators as used in contexts that are predictably much more rare—namely, contexts in which we use complex epistemic constructions to express especially complex probabilistic thoughts. Along these lines, my semantics delivers a nicely unified account of epistemic expressions and modal adverbial phrases, which is an additional reason to prefer my theory over its many rivals.

1.3 Indicative conditionals

On my semantics, indicative conditionals are context sensitive. Context contributes a partition to the interpretation of a conditional, and you believe the conditional just in case every partition element that accepts the antecedent relative to your credences also accepts the consequent. Against this semantics, Edgington worries that in many contexts, "there is no relevant partition in play" that could determine the content of a conditional at that context.

As I see it, though, the partitions used to interpret epistemic vocabulary are roughly like questions under discussion. It is easy to put them into play. Just uttering a conditional with a simple antecedent is often sufficient to raise the question of whether the traditional propositional content of that antecedent is the case, and such a conditional is often interpreted relative to a partition that is decisive with respect to that content. For instance, consider (7) as uttered in the context of a medical consultation:

(7) If you have the operation, you will probably be cured.

This conditional easily raises the question of whether you will have the operation. As long as the conditional operator in (7) is interpreted relative to this question, and the probability operator is interpreted relative to the question of whether you are cured, then you believe the content of this conditional just in case you have more than .5 conditional credence that you will be cured, given that you have the operation. This is exactly the result that Edgington is looking for. Hence my semantics does indeed generate readings of conditionals that vindicate the Ramsey-Adams thesis, as Edgington desires.

1.4 Knowledge and truth

Edgington ends with a worry for my interpretation of the factivity condition for probabilistic knowledge. She points out that the following sentences sound fine:

(8) I knew the die was unlikely to land 6, though it turned out to land 6.

(9) The doctors knew it was very probable that you had the disease, though you didn't have it.

(10) We knew at halftime that it was virtually certain that team *A* would win, though they lost.

(11) At 11:30, you knew that it was likely that Jones would reach the center of the maze by noon. At 11:45, you knew that it was unlikely that Jones would reach the center by noon.

According to Edgington, these conjunctions demonstrate that you can know that something is probable, even when it is false.

As we examine these sentences, we must keep in mind that probability operators are used to express two sorts of contents—not only thoroughly probabilistic contents, but also nominally probabilistic contents about probability facts. The first conjuncts of (8)–(11) have natural readings on which they ascribe knowledge of the latter contents. For instance, the speaker of (8) did know that there was a low objective chance of the die landing 6. The doctors in (9) did know that they had a high evidential probability for the proposition that you had the disease. These claims are perfectly consistent with the factivity of probabilistic knowledge.

To see what the factivity condition rules out, consider the following utterances:

(12) a. *Jones:* It'll probably rain tomorrow.
b. *Smith:* It probably won't rain tomorrow.

There is something that Jones and Smith disagree about. Jones believes that it is more than .5 likely to rain tomorrow, while Smith believes that it is less than .5 likely. Suppose that as a matter of fact, it doesn't rain. Then there is at least one important sense in which, when it comes to the above disagreement, Smith is the one who got it right, while Jones got it wrong. This is true even if Jones had just the right credences given her evidence at the time, and even if her credences matched the objective chance that it would rain. This sense of correctness is the interpretation of the factivity condition that I endorse. If something is false, then it is not probable, and no one ever knew that it was.

Edgington worries that by saying that falsehoods are not probable, I have thereby committed myself to the implausible claim that only truths can be probable. However, we must be careful to distinguish the following rules of inference:

(13) a. P is false.
b. Therefore, P is not probable.

(14) a. P is probable.
b. Therefore, P is true.

The first rule is fine. The second is valid only in exceptional contexts—roughly speaking, those in which every proposition is considered to be certainly true or certainly false. In ordinary contexts, we think about propositions in more ordinary ways, or as some might put it, under more ordinary modes of presentation. As in §7.5 of my book, it is helpful to compare our ordinary assertions about lotteries. From the premise that you lost the lottery, it follows that you did not probably win. But from the premise that you probably lost, it does not follow that you did indeed lose. This latter inference is valid only in exceptional contexts in which each ticket, including yours, is considered to be a definite loser or a definite winner.

Edgington maintains that your low credence that the die will land 6 is "right" and "correct," even if the die does end up landing 6 after all. To be sure, there is a sense in which your low credence in 6 is the just the right credence for you to have. For instance, it satisfies the norm that says that you should believe a probabilistic content if and only if it contains your evidential probabilities. The knowledge norm is not the only useful standard for evaluating probabilistic belief. Different norms govern our evaluation of beliefs in different contexts. In her context, Edgington evaluates probabilistic beliefs using a non-factive norm of belief. Her positive evaluations of false beliefs are consistent with my thesis that false beliefs do not constitute knowledge.

A recurring theme of the later chapters of my book is that we have compelling uses for factive norms of probabilistic belief. Even if your credences match your evidential probabilities, they may fall short of other important ideals. As you read a thrilling mystery novel, you may adjust your credences as you gather evidence, and you may be blameless for failing to realize until the very end that the butler and the gardener might both be innocent. But even if you are blameless in your belief, avoiding blame is only one epistemic goal. As subjects gather evidence in an actual criminal trial, for instance, it is natural to evaluate their beliefs according to externalist epistemic norms. A judge may be perfectly blameless for falsely believing that a defendant is probably liable for trespassing, and her credences may match her evidential probabilities, but the defendant may naturally judge her belief by another standard. As I argue in Moss 2019c, in order to sustain a verdict of guilt or liability, probabilistic beliefs must constitute knowledge.

2 Reply to Pavese

Pavese argues that intentional action requires probabilistic knowledge—in particular, that you intentionally ϕ only if you know that it is sufficiently likely that you will ϕ.[5] Pavese also defends a broader knowledge-centered psychology according to which knowledge plays a central role in explanations of intentional action. I hope that Pavese is correct on both counts, since her view offers exciting potential applications for probabilistic knowledge. In what follows, I play devil's advocate. First, I give an example that challenges the claim that intentional action requires probabilistic knowledge. Second, I give an example that challenges the broader thesis that knowledge plays a central role in explanations of intentional action.

2.1 Intentional action without probabilistic knowledge

Let us briefly review the dialectic up to this point. According to Setiya 2012, examples like the following show that intentional action does not require knowledge:

> *Surgery Notes:* Mary has just emerged from a surgical operation on her paralyzed arm. She reads the notes on the operation written up by her surgeon, and she comes to believe that it is at least .3 likely that the operation restored her ability to clench her fist. She is under anaesthetic and cannot see or feel her hand. Hoping for the best, she attempts to clench her fist. She succeeds.

Setiya suggests that Mary clenches her fist intentionally, even though she does not know that she is clenching it. In response, Pavese points out that Mary does have

5. Actually, Pavese defends a stronger condition—namely, that you intentionally ϕ only if you know, for some means ψ of ϕ-ing, that you yourself are sufficiently likely to ϕ by ψ-ing. For our purposes, we can abstract from some nuances of her account.

some knowledge—namely, she knows that it is at least .3 likely that she will succeed in clenching her fist. Perhaps intentional action does not require propositional knowledge, then, but only knowledge of this sort of modest probabilistic content.

Unfortunately for Pavese, this revised knowledge requirement seems to be challenged by another example:

> *Gettier Notes:* Gary has just emerged from a surgical operation on his paralyzed arm. He reads some surgery notes and comes to believe it is at least .3 likely that the operation restored his ability to clench his fist. As it happens, due to a filing mishap, Gary is actually reading the notes for the operation performed on Mary. He has no other justification for believing that his own operation might have restored his motor control. Gary is under anaesthetic and cannot see or feel his hand. Hoping for the best, he attempts to clench his fist. He succeeds.

It seems intuitive to say that Gary clenches his fist intentionally, even though he does not know that he is sufficiently likely to succeed in clenching it. If that is right, then intentional action does not require even modest probabilistic knowledge.

Where do we go from here? Fans of the knowledge requirement might attempt to respond by arguing that Gary does indeed know that he is sufficiently likely to clench his fist, despite the dubious origins of his probabilistic belief. The strength of this response depends on foundational questions about exactly what sort of luck is incompatible with knowledge. When a belief is initially true by luck, could it come to have the status of knowledge simply in virtue of your successfully using it to guide your action? Fans of the knowledge requirement might hang on to the claim that intentional action requires knowledge, taking my example to illustrate that knowledge itself requires significantly less than we might have thought.[6] It is beyond the scope of this paper to adjudicate this debate. For present purposes, my conclusion is simply that it is far from obvious that intentional action requires probabilistic knowledge.

2.2 The role of knowledge in explanations of intentional action

Pavese defends a knowledge-centered psychology, a view on which knowledge takes center stage in explanations of intentional action. According to Pavese, facts about knowledge explain why Bobby and Cindy fail to act intentionally. As Pavese puts it,

> Bobby does not know that he can provoke the explosion through his plan. That is why his success is too coincidental to count as intentional. Cindy does not know that she can win the lottery by purchasing that particular ticket. That is why her victory is too coincidental to count as intentional. . . A knowledge-centered psychology can explain why luck can undermine the intentionality of Bobby's and Cindy's successes: by undermining their knowledge.

6. For sympathetic discussion of this line of thought, see Stanley 2011, p.190; Stalnaker 2012, p.755; and Baumann 2014.

The examples of Bobby and Cindy teach us that intentional action is incompatible with luck. As some might put it, you intentionally ϕ just in case you intend to ϕ, you do indeed ϕ, and your success is caused by your intention in some special non-lucky way. At first glance, it is tempting to understand Pavese as suggesting that the relevant anti-luck condition on intentional action can be understood in terms of knowledge. When she says that a knowledge-centered account explains why Bobby and Cindy fail to act intentionally, it is tempting to conclude that you intentionally ϕ just in case you intend to ϕ, you do indeed ϕ, and you have some sort of knowledge that guarantees that your intention causes your success in just the right way. If this were correct, knowledge would indeed play an important role in explaining intentional action. After all, according to this analysis, the absence of knowledge is precisely what distinguishes Bobby and Cindy from agents who act intentionally, since Bobby and Cindy satisfy the other conditions for intentional action.

However, knowledge does *not* play this role in explaining intentional action. Consider the following example:

> *Driving Practice:* You plan to kill your uncle by running him over with your car. For weeks ahead of time, you practice running over animals in order to make sure you know how to kill them. A reliable oracle informs you that you will almost certainly kill your uncle in just the way you have been practicing. As you drive home, you intend to practice your skills by running over a moose in the road. As it happens, the animal you kill is not a moose at all, but your uncle disguised in a very realistic moose costume.

In this example, you satisfy the probabilistic knowledge requirement on intentional action proposed by Pavese. On the basis of the oracle's reliable testimony, you know that you will almost certainly kill your uncle by running him over with your car. Furthermore, you grasp this proposition under a practical mode of presentation. As you step on the gas pedal, you may say to yourself, "According to the oracle, *this* is exactly how I will kill my uncle!" And yet, despite having this knowledge, you fail to kill your uncle intentionally. Although your beliefs are connected to the world in the special way that knowledge requires, your knowledge is not connected with your success in the special way that intentional action requires.

At this point, one might introduce epicycles to preserve a knowledge-based analysis of intentional action. For instance, one might add that in order to intentionally kill your uncle, you must succeed in killing him *partly in virtue of* your knowledge that you will likely kill him. Further counterexamples abound; for instance, we could add that you succeed in hitting your uncle partly in virtue of the confidence that you gain when the oracle informs you that your practice will likely pay off. As further epicycles and counterexamples are generated, it will seem more and more reasonable to abandon this project of analysis.[7]

Back to Bobby and Cindy. Bobby fails to intentionally set off his bomb. Cindy fails to intentionally win the lottery. According to Pavese, the fact that luck precludes their actions from being intentional is explained by the fact that these subjects fail to know that they are sufficiently likely to succeed. But the latter fact is a bad explanation of the former. To be specific, the explanation is bad because it is not sufficiently general. Luck can preclude your action from being intentional, even when you satisfy the probabilistic knowledge requirement that Pavese describes, and a lot of other knowledge requirements besides. To sum up, knowledge is incompatible with luck, and intentional action is incompatible with luck. But neither of these facts explains the other. This is an important sense in which knowledge does not take center stage in explanations of intentional action.[8]

At first, one might expect me to find this conclusion disappointing. After all, my project would gain further motivation if intentional action could be reductively analyzed only in terms of probabilistic knowledge. But as I see it, the spirit of my book is actually reinforced by our reflections on intentional action. An action is intentional when it exhibits a fundamental sort of causal connection between mind and world, where this sort of connection has a distinctive modal robustness. According to traditional epistemologists, this special sort of connection also holds between full beliefs and the world—namely, when full beliefs are knowledge. According to my book, the same sort of connection can hold between probabilistic beliefs and the world—namely, when those probabilistic beliefs are knowledge. By expanding our understanding of the fundamental sort of connection to the world that knowledge requires, and by recognizing that states other than full beliefs can connect with the world in just this way, philosophers of action pave the way for my project of extending this same special status to probabilistic beliefs.

3 Reply to Campbell-Moore and Konek

Campbell-Moore and Konek defend a number of claims about probabilistic belief.[9] Unfortunately, I do not have space to discuss all of their arguments here. I will focus

7. For arguments against reductive analyses of intentional action, see Levy 2013 and Williamson 2017. Similar arguments have been discussed by criminal law scholars as they evaluate reductive analyses of the concurrence requirement on intentional criminal acts. For an overview of relevant legal literature, see Simons 2002 and chapter 2 of Sarch 2019.
8. For similar arguments against intellectualist accounts of skill and know-how, see p.741ff. of Dickie 2012 and p.543 of Glick 2015, respectively.
9. A terminological point: by contrast with Campbell-Moore and Konek, I do not use 'probabilistic belief' and 'belief with a probabilistic content' interchangeably. As used in the book, 'probabilistic belief' is a general term for credences and other subjective probability judgments, where this term is neutral about whether the contents of such beliefs are sets of worlds or sets of probability spaces.

on the central argument of §2.3, which can be illustrated by the following example:

> *Coin Beliefs:* Jones and Smith have tossed a coin, but they haven't yet looked to see how it landed. Jones believes that the coin is double-sided, but she has no opinion about the likelihood that it has two heads or two tails. Smith believes that the coin is not double-sided. He has no opinion about the bias of the coin.

Campbell-Moore and Konek make three claims about *Coin Beliefs*. The first is that Jones and Smith have importantly distinct doxastic states. As one might put it, Jones believes (15), whereas Smith does not:

(15) Either the coin definitely landed heads, or it definitely landed tails.

Second, Jones and Smith desire exactly the same gambles about the outcome of the coin toss. Third and finally, Campbell-Moore and Konek claim that Jones and Smith differ in their probabilistic beliefs, since Jones believes the following set of probability spaces, whereas Smith does not:

(16) $\{p \mid p(\mathit{Heads}) = 1 \textit{ or } p(\mathit{Tails}) = 1\}$

Together these claims appear to favor models that ascribe probabilistic beliefs to Jones and Smith, over and above models that represent belief states entirely in terms of desirable gambles.[10]

I agree that Jones and Smith differ in whether they believe (15). But I do not think that we should use (16) to model this difference in belief. To see this, consider a third subject, Brown, who believes that the coin is double-sided. Suppose that Brown has a hunch about the coin, believing that the coin is .8 likely to have two heads. Like Jones, Brown believes (15). Brown agrees that the coin definitely landed heads if it is two-headed and definitely landed tails if it is two-tailed, and she believes that these are the only two ways the world could be. However, Brown does not believe the probabilistic content (16). In fact, Brown believes a content that is inconsistent with (16)—namely, the set of probability spaces that assign .8 to *Heads*. Hence the content of (15) is not (16), as Campbell-Moore and Konek suggest.[11]

What is the content of (15)? According to my semantics, (15) is context sensitive. In the context of *Coin Beliefs*, 'or' in (15) is naturally interpreted relative to a

10. As I see it, these three claims do not additionally favor the *complex content account* of probabilistic beliefs. Advocates of the complex attitude account will appeal to differences in complex attitudes to distinguish the total belief states of Jones and Smith. In other words, it is important to distinguish the following theses: (a) that the contents of probabilistic beliefs are sets of probability spaces, as opposed to mere sets of possible worlds, and (b) that total belief states are modeled by sets of probabilistic contents, as opposed to mere sets of probability spaces. I defend only (a) in my book, whereas Campbell-Moore and Konek defend only (b) in §2.3.

11. For further details about why Brown can consistently believe (15) while believing that heads is .8 likely, see my discussion of the argument from disjunction in §7.4 of my book.

partition containing three propositions: that the coin has two heads, that it has two tails, and that it has one of each. Jones believes (15) because only the first two of these propositions are possible for her, and because relative to her credences, each of these propositions either accepts that the coin definitely landed heads or accepts that it definitely landed tails. The same goes for Brown. Meanwhile, Smith does not believe (15), since the third partition proposition is possible for Smith, and relative to his credences, that proposition does not accept the content of either disjunct of (15). In other words, according to my semantics, speakers generally use the disjunction (15) to express a combination of conditional beliefs. In the context of *Coin Beliefs*, the relevant combination of conditional beliefs is equivalent to the simple belief that the coin has two heads or two tails. The strict content of this belief is a convex set of probability spaces—namely, all and only those probability spaces according to which it is not possible that the coin has one heads side and one tails side.

In §2.3, Campbell-Moore and Konek set out to construct a natural example where subjects disagree about a non-convex set of probability functions. Unfortunately, the natural reading of (15) in *Coin Beliefs* does not have such a set as its content. But that is not to say that no such examples can be found. For instance, perhaps (17) is commonly used to assert a non-convex content:[12]

(17) Whether Jones smokes is independent of whether Smith smokes.

Hence our discussion in epistemology points the way to a valuable project in formal semantics—namely, identifying and giving a semantics for natural language expressions used to assert non-convex probabilistic contents. This project is not part of my book. But to the extent that ordinary speakers express non-convex probabilistic beliefs, that only adds to the interest of the broader research program of developing a probabilistic semantics for natural language expressions.

12. For further discussion of potential non-convex contents of independence judgments and other beliefs, see §2 of Moss 2019b.

References

Baumann, Peter. 2014. "No Luck With Knowledge? On a Dogma of Epistemology." *Philosophy and Phenomenological Research*, vol. 89 (3): 523–51.

Dickie, Imogen. 2012. "Skill Before Knowledge." *Philosophy and Phenomenological Research*, vol. 85 (3): 737–45.

Glick, Ephraim. 2015. "Practical Modes of Presentation." *Noûs*, vol. 49 (3): 538–59.

Levy, Yair. 2013. "Intentional Action First." *Australasian Journal of Philosophy*, vol. 91 (4): 705–18.

Moss, Sarah. 2018. *Probabilistic Knowledge*. Oxford University Press, Oxford.

—. 2019a. "Full Belief and Loose Speech." *Philosophical Review*, vol. 128 (3): 255–91.

—. 2019b. "Global Constraints and Imprecise Credences." Ms., Department of Philosophy, University of Michigan. Forthcoming in *Philosophy and Phenomenological Research*.

—. 2019c. "Knowledge and Legal Proof." Ms., Department of Philosophy, University of Michigan. Forthcoming in *Oxford Studies in Epistemology*, vol. 7.

—. 2019d. "Simple Desire and Loose Speech." Ms., Department of Philosophy, University of Michigan.

Sarch, Alex. 2019. *Criminally Ignorant: Why the Law Pretends We Know What We Don't*. Oxford University Press, Oxford.

Setiya, Kieran. 2012. "Knowing How." *Proceedings of the Aristotelian Society*, vol. 112 (3): 285–307.

Simons, Kenneth. 2002. "Does Punishment for "Culpable Indifference" Simply Punish for "Bad Character"?: Examining the Requisite Connection Between *Mens Rea* and *Actus Reus*." *Buffalo Criminal Law Review*, vol. 6 (1): 219–315.

Stalnaker, Robert. 2012. "Intellectualism and the Objects of Knowledge." *Philosophy and Phenomenological Research*, vol. 85 (3): 754–61.

Stanley, Jason. 2011. *Know How*. Oxford University Press, Oxford.

Williamson, Timothy. 2017. "Acting on Knowledge." In *Knowledge First: Approaches in Epistemology and Mind*, Adam Carter, Emma Gordon & Benjamin Jarvis, editors, 163–81. Oxford University Press, Oxford.