

Reply to MacFarlane and Greco

Sarah Moss
ssmoss@umich.edu

I am deeply grateful to John MacFarlane and Daniel Greco for their insightful comments on my book. This reply is divided into six sections, starting with responses to the four objections raised by MacFarlane and ending with responses to the two objections raised by Greco.

1 Overgeneration of readings

John MacFarlane expresses two reservations about my semantics for epistemic vocabulary. The first is that it overgenerates readings of sentences such as:

- (1) If the die landed on a low number, it probably landed on an odd number.

The content of (1) is a set of probability spaces. Which set? At first, one might expect the answer to be, “the set of probability spaces that assign high conditional probability to the die having landed on an odd number, conditional on it having landed on a low number.” But this answer is not quite right. According to my semantics, ‘if’ in (1) can be interpreted relative to different partitions in different contexts. In most ordinary contexts, ‘if’ is interpreted relative to the question of whether the die landed on a low number, and (1) has the content mentioned above. But in principle, (1) could also have a stronger content, a content that you could fail to believe even if you have high conditional credence that the number rolled is odd if it’s low.

At this point, there are two objections about overgeneration that one might raise for my account. First, one might object that there just aren’t any actual conversational contexts in which (1) is used to express anything stronger than a simple constraint on conditional credences. If this were true, it would count as a significant strike against

my view. Advocates of my semantics would incur the burden of explaining why it's impossible to hear any stronger readings of (1), and it would be reasonable to prefer theories that did not allow for these readings to begin with.

The correct response to this objection is that there are indeed contexts where (1) is naturally read as having a stronger content. Consider the following example:

Kiddie Casino: Oliver, age 5, is playing a simple gambling game. The dealer rolls a die behind a screen, and the players bet at even odds on whether the number rolled is even or odd. Oliver has discovered a hole in the screen, through which he can see the die. Unfortunately, Oliver is not great at math. He can recognize the low numbers 1, 2, and 3. But he can't tell apart 4, 5, and 6, and he has no idea what it means for a number to be odd or even. After each roll, Oliver tells his teenage brother Liem what he can see, and Liem tells Oliver how to bet. The following conversation ensues:

Oliver: It's a high number! What do I do?

Liem: Okay, that's probably even. Bet on even.

Oliver: The next number is a 1.

Liem: Bet on odd.

Oliver, *starting to get lazy*: It's another low number.

Liem: Wait, what? Count the dots! Which low number is it?

Oliver: Why does that matter? If it's low, then I should just bet on odd, right?

Suppose that Oliver asserts the following conditional to Liem:

(2) If the number is low, then it's probably odd.

Liem could respond by correcting Oliver, saying, "No! One way that the number could be low is by being 2, in which case it's definitely even. So just because the number is low, that doesn't mean you should bet that it's odd." In this context, it is natural to interpret (2) as having a content that Liem does not believe, despite his having a high conditional credence that the number rolled is odd if it's low. Liem does not believe the content of (2) because 'if' is interpreted relative to an unusually fine-grained partition. Roughly speaking, Oliver is saying that whenever the die lands on a low number, it probably lands on an odd number, where the relevant ways of landing on a low number include each of the three specific outcomes that Oliver can discriminate between.

Our judgments about *Kiddie Casino* not only answer the first objection for my semantics; they also provide a positive argument in support of my view. A number of alternative semantic theories fail to account for the context sensitivity of (1) and (2).¹

1. For instance, according to the semantic theories developed by VELTMAN 1996, YALCIN 2012, and MAC-FARLANE 2014, epistemic expressions do not exhibit this sort of context sensitivity.

They fail to generate the stronger reading of (2) in the context of *Kiddie Casino*. By contrast, my semantics generates this reading of (2), while leaving open the possibility that there are pragmatic reasons why the reading is ordinarily difficult to access.²

This brings us to a second objection that one might raise for my semantics. Rather than objecting that my semantics generates a reading of (1) that is not available in *any* context, one might raise the more modest objection that my semantics generates a reading of (1) that is not available in an *ordinary* context where we are rolling dice and guessing how they landed. As MacFarlane points out, “Both *LowOrHigh?* and *Which-Face?* are natural partitions to think about” in the ordinary context that he describes. If available interpretations of ‘if’ depended only on what partitions it was natural to think about, then my theory would fail to explain why it is hard to interpret (1) relative to *WhichFace?* in an ordinary context.

The correct response to this second objection is that there are indeed other factors that influence which readings of epistemic expressions are available in a given context. For starters, even when *LowOrHigh?* and *WhichFace?* are both natural partitions to think about, the content that results from interpreting (1) relative to the former partition is considerably more natural than the content that results from interpreting it relative to the latter. In addition, the former content is considerably more useful. Speakers routinely have some reason to express their conditional credences in propositions. Conditional credences play an important role in reasoning, especially in ordinary gambling contexts where we can easily imagine ourselves making conditional bets. By contrast, other potential readings of (2) are obscure. It is rarely relevant whether each of several ways of being low is such that if a number is low in that way, it is probably odd. The context of *Kiddie Casino* is unusual precisely because the speakers in this context actually have some reason to entertain this complicated probabilistic content. More generally, it is unusual for a speaker to have any reason to express the state of having a high conditional credence in a proposition given any of a number of other propositions, each of which is such that if the speaker updates on it, she accepts a certain probabilistic content. But sure enough, as the case of *Kiddie Casino* illustrates, it is easier to interpret conditionals as expressing such beliefs in contexts where it makes sense for speakers to be reasoning with them. To sum up, my theory accounts for the fact that the available readings of (2) vary predictably in accordance with pragmatic features of contexts, and this is a distinctive advantage of my partition-based semantics for conditionals. In short, my semantics is not *overgenerating* readings, but simply *generating* them.

2. For further discussion of this relative advantage of my semantics, see Moss 2018b, p.72ff.

2 No independent motivation for the indices

MacFarlane worries that my semantics appears to posit “variables that can *only* be filled by context, and not given values explicitly or bound by quantifiers.” By contrast, one can easily bind variables over quantifier domains:

- (3) Whenever Jones is in charge of bringing drinks to a party, there is no beer.

and one can explicitly specify the values of such variables:

- (4) There is no beer *at this party*.

If epistemic vocabulary is also context sensitive, MacFarlane asks, then shouldn't we expect to find analogous constructions in which the partition used to interpret an epistemic expression is made explicit, or in which the variable over partitions is bound?

The correct response is that we should expect epistemic expressions to behave less like quantifiers, and more like context-sensitive expressions with more similar sorts of implicit arguments. For a closer comparison with my account, consider a contrastivist account of causal talk. According to contrastivists, 'cause' is context sensitive, and it is interpreted relative to a set of propositions—namely, a set of contrasting causal claims. For example, suppose that India is plagued by a terrible drought, and the government has no food reserves. A terrible famine ensues.³ According to the contrastivist, 'caused the famine' is context sensitive. In some contexts, we can truly say that the drought caused the famine; in other contexts, we can truly say that the lack of food reserves caused it. However, we do not have natural ways of specifying the implicit arguments of these causal claims. For instance, we do not normally say things like:

- (5) ?Relative to the question what destroyed the crops, the drought caused the famine.

At best, we have more roundabout ways of specifying the implicit argument of 'cause' in ordinary contexts, as in the following dialogue:

- (6) a. What caused the famine?
b. If you're asking what destroyed the crops, I'd say it was the drought. But if you're asking what the government could have done differently, I'd say it was the lack of food reserves that caused the famine.

For another example, consider the contrastivist account of 'voluntary' defended by SINNOTT-ARMSTRONG 2012. Suppose that Smith lights a cigarette on a nonsmoking

3. This example is due to HART & HONORÉ 1985, pp. 35–36.

flight. According to Sinnott-Armstrong, whether we can truly apply ‘voluntary’ to this action depends on our context. In order to clarify what sort of context we are in, we might have the following conversation:

- (7) a. Did Smith light that cigarette voluntarily?
- b. If you’re asking whether he could go a whole flight without smoking, I’d say it was involuntary. But if you’re asking whether he could have waited a few minutes, I’d say it was voluntary.

Again, we have roundabout ways of specifying the set of propositions used to interpret a context-sensitive term.

According to my semantics, ‘if’ and ‘probably’ are also interpreted relative to sets of propositions. As one might expect, we have roundabout ways of clarifying the interpretations of these expressions. For instance, suppose that Jones is standing on top of a tall building. There is no safety net along the roof, but Jones certainly does not have a death wish.⁴ We might have the following conversation about Jones:

- (8) a. What are the chances that Jones will live if she jumps?
- b. If you’re asking whether she could survive a ten-story fall, I’d say the chances are low. But if you’re asking whether she’s suicidal, I’d say the chances are high; Jones wouldn’t jump unless there was a safety net up there to stop her from falling.

To sum up, in order to interpret certain context-sensitive expressions, speakers must figure out roughly what alternatives are under discussion and what background beliefs they are supposed to be holding fixed. When the speakers in (6) hold fixed the fact that the government had no food in reserve, they can truly say that the drought caused the famine. When the speakers in (8) hold fixed the fact that there is no net along the roof, they can truly say that if Jones jumps, she will probably die. I agree with MacFarlane that ordinary speakers do not use expressions like ‘relative to the question how I should bet’ in order to specify what they are holding fixed in a given context. But this does not constitute a challenge for my semantics, since epistemic expressions act like other similar context-sensitive expressions in this respect.

In addition to indirectly specifying the values of partition pronouns, we can indeed bind them. Here again, epistemic vocabulary behaves like other similar context-sensitive vocabulary. For instance, one can easily imagine the above dialogues being continued by a third speaker as follows:

4. This example is inspired by example (15) in chapter 8 of LYCAN 2001, which builds on the discussion of subjunctive conditionals in SLOTE 1978 and LEWIS 1979. For further discussion, see §4.3 of MOSS 2018b.

- (10) c. Actually, the government was secretly depleting the groundwater, so however you look at it, the government caused the famine.
- (11) c. Actually, Smith can go for a few days without smoking, so however you look at it, his lighting the cigarette was voluntary.
- (12) c. Actually, the fire department just put the safety net back up, so however you look at it, there's a good chance that Jones will live if she jumps.

In §4.5 of Moss 2015, I argue that bound readings of sentences like (12-c) provide support for my semantics, since rival theories of conditionals have a hard time accounting for these readings. Here again, careful reflection on ordinary language turns the tables in favor of my view.

3 Factivity

Turning to epistemology, MacFarlane raises a worry for my account of factivity. Suppose that I know that it is .5 likely that a fair coin landed heads. Then since knowledge is factive, it is indeed .5 likely that the coin landed heads, and no one else can know otherwise. However, it seems compelling to say that after looking at the coin, I'll know that it's not .5 likely that it landed heads.

As MacFarlane points out, a relativist about epistemic vocabulary can say that my later belief is correct because it is true from my later perspective. Although I do not endorse relativism in the book, my theory is relativist-friendly. If you are a relativist about epistemic vocabulary, then you can—and probably should—be a relativist about knowledge ascriptions embedding that vocabulary. If the content that heads is .5 likely is true relative to my earlier perspective but not my later perspective, then it is something that my earlier self can know relative to my earlier perspective but not my later perspective.

That being said, we do not need relativism in order to make sense of our judgments about MacFarlane's example. To pull apart various concerns, let us recast the example in the third person:

Coin Flip: A subject—call her Agent Dumb—has flipped a coin that she knows to be fair, but she hasn't yet seen how it landed. At a later time, Agent Smart looks at the coin and sees whether it landed heads or tails.⁵

We can now distinguish two worries for my account of probabilistic knowledge. The first worry can be stated from the point of view of the narrator of *Coin Flip*. The worry is that both of the following ascriptions seem compelling:

⁵ In the story as told by MacFarlane, Agent Smart is Agent Dumb, but we need not make this assumption.

- (9) Agent Dumb knows that it is .5 likely that the coin landed heads.
- (10) Agent Smart knows that it is not .5 likely that the coin landed heads.

However, my account appears to entail that these ascriptions cannot both be true.

In response to this worry, it is important to keep in mind that (9) and (10) have multiple readings. In addition to using probability operators to express and ascribe thoroughly probabilistic beliefs, we also use them to express and ascribe beliefs about probability facts, such as facts about objective chances, subjective credences, and evidential probabilities. On one reading, (9) says that Agent Dumb knows that it is .5 likely on her limited evidence that the coin landed heads, whereas (10) says that Agent Smart knows that heads is not .5 likely given her richer body of evidence. These are facts about the *evidential probability* of heads for each agent in roughly the sense of WILLIAMSON 2000—that is, the probability of heads that results from feeding their propositional knowledge to an initial measurement of “the intrinsic plausibility of hypotheses prior to investigation” (211). On this reading, (9) and (10) are compelling—and they are also consistent. However the coin landed, Agent Dumb and Agent Smart can indeed know these evidential probability facts.

In addition to accommodating our intuition that (9) and (10) can both be true, my account can also accommodate the more theoretical intuition mentioned by MacFarlane—namely, that “the probabilistic beliefs formed by others with different information are in some sense right.” There are multiple respects in which credences can be good or bad. Having credences that constitute knowledge is one way for your credences to be good. Having credences that match your evidential probabilities is another. Agent Smart and Agent Dumb each have credences that match their own evidential probabilities. Accepting the existence of probabilistic knowledge does not force us to reject the familiar practice of saying that such agents are doing something right, epistemically speaking. Rather, it simply adds another dimension along which credences can be evaluated and potentially criticized.

Of course, (9) and (10) can also be used to ascribe knowledge of thoroughly probabilistic contents. On this second reading, (9) and (10) are inconsistent—but they are not both compelling. As we evaluate this reading, we must ask ourselves: what is true in the *Coin Flip* story? Is it indeed .5 likely that the coin landed heads? The answer to this probabilistic question will constrain facts about which subject in the story has knowledge. If it is .5 likely that the coin landed heads, then Agent Smart doesn’t know that it isn’t. If it’s not .5 likely, then Agent Dumb doesn’t know that it is. In other words, as you narrate or reflect on the story about Agent Dumb and Agent Smart, you must take sides. And as soon as you do, you will favor one knowledge

ascription over the other. Only one subject in the story has a true belief, and so at most one subject can have knowledge.

So far, we have examined *Coin Flip* from the point of view of the narrator. The example also raises a worry that we can state from the point of view of Agent Dumb. As Agent Dumb, we can easily imagine thinking, "I know that heads is .5 likely. But then again, Agent Smart seems to know that it is not .5 likely. And if I take myself to know the former content, then I have to deny that she knows otherwise." Something seems amiss in this reasoning. The same goes for the thought that MacFarlane imagines Agent Dumb having—namely, "I'll have to suppose that I'm going to go off the rails, when I examine the coin."

The correct response to this worry is that just as the narrator must take sides, the same goes for Agent Dumb herself. It is not a foregone conclusion that Agent Dumb should stick to her guns in the context of this example. She might instead think, "At most one of us has knowledge—myself or Agent Smart. Which one is it? Well, if Agent Smart were here in front of me, I would not be trying to argue her out of her probabilistic belief, or insisting that heads is indeed exactly .5 likely. Rather, I would be paying attention to some relevant alternatives to my .5 credence—namely, whatever credences Agent Smart might have about the coin. I would have to admit that I can't rule those alternatives out, which would lead me to conclude that as a matter of fact, my .5 credence isn't knowledge after all." As Agent Dumb delivers this monologue, something happens to her context. Her epistemic standards go up. Agent Dumb used to say 'I know that heads is .5 likely' in virtue of the fact that she could rule out certain ordinary alternatives to her probabilistic belief, such as the possibility that the coin was biased towards tails. After thinking about Agent Smart, Agent Dumb is no longer willing to utter that same sentence, since she can't rule out certain salient contents that Agent Smart might believe. Hence Agent Dumb is not forced to accept that she is about to "go off the rails," as MacFarlane puts it. It may be just as natural for her to accept that she herself lacks probabilistic knowledge.

At this point, readers may be wondering, "But can't we just play this game with any middling credence at all?" Yes, we can. In any context, there will be skeptical alternatives waiting in the wings, ready to shift our standards for interpreting knowledge ascriptions. But this is nothing new, nor is it anything specific to probabilistic knowledge. There have always been plenty of skeptical alternatives for traditional knowledge ascriptions, too. Brains in vats, evil demons, simulation hypotheses—there are plenty of possibilities that you can't rule out. A familiar response to this skeptical observation is that these alternatives are not relevant in ordinary contexts of knowledge ascription. The same goes for extreme credences in many ordinary contexts in which we assert that middling credences constitute knowledge.

4 Safety

According to moral emotivists, believing that murder is wrong involves having certain attitudes of disapproval towards murder. But what is it to believe that murder is *necessarily* wrong, or that it is wrong at some merely possible world? Here is a natural thought: when you believe that murder is wrong at some merely possible world, that involves having certain attitudes of disapproval towards merely possible murders. Like ordinary moral contents, *modal* moral contents can also be understood in terms of attitudes of approval and disapproval.

Just as you have moral attitudes about merely possible murders, you also have credences about merely possible lotteries. To say that your ticket probably lost the lottery at a merely possible world is to express your high credence that it lost at that world. That is, modal sentences such as (11-a) express probabilistic beliefs, as explicated in (11-b):

- (11) a. At w_0 , it is probably the case that I lost the lottery.
- b. It is probably the case that I lost the lottery at w_0 .

There is indeed a general algorithm for translating sentences like (11-a) into sentences like (11-b): delete the intensional operator \mathcal{O} that takes scope over epistemic vocabulary, and replace it with tokens of \mathcal{O} that take scope over each maximal simple sentential constituent of the original complement of \mathcal{O} , i.e. each sentential constituent that does not contain epistemic vocabulary and that is not a constituent of any larger sentential constituent that does not contain epistemic vocabulary. The partitions used to interpret epistemic vocabulary are shifted in just the same way, with each proposition p being replaced by result of applying the denotation of \mathcal{O} to p .

This translation strategy provides us with a useful interpretation of the safety condition for probabilistic beliefs. For example, (12-a) is explicated by (12-b):

- (12) a. For every close world w where I believe I probably lost,
 at w , it is probably the case that I lost the lottery.
- b. For every close world w where I believe I probably lost,
 it is probably the case that I lost the lottery at w .⁶

MacFarlane worries that this interpretation of safety has an unfortunate result. Suppose Euthydemus forces Crito to think about some close world as one where Crito definitely loses the lottery. Then my view appears to predict that Crito can't truly

6. For more on the formal semantics of these sorts of constructions—quantifiers that take scope over epistemic expressions, which themselves take scope over singular terms—see ALONI 2000, YALCIN 2015, and NINAN 2018, in addition to §7.5 of MOSS 2018b.

say that he knows that his ticket probably lost. But intuitively, Crito is still capable of truly self-ascribing this knowledge.

There are two possible responses to this worry. For present purposes, I will stay neutral between them. The first response is that in the context described by MacFarlane, Euthydemus does not actually succeed in getting Crito to think of possible worlds in the right way—that is, in a way that prevents Crito from truly self-ascribing knowledge that he probably lost. Euthydemus points out that there is a close world in which Crito wins the lottery. But Crito might continue the conversation by saying, “For sure, there is some world—call it *Lucky*—among the close worlds where I believe that I probably lost. But for any given such world, *that world* is extremely unlikely to be identical to *Lucky*. That is why my high credence that I lost is safe, and why it is knowledge.” In other words, Crito may accept that *Lucky* is among the close worlds where he believes he probably lost, even while still thinking of those same worlds as *arbitrary* close worlds where he believes he probably lost.

A second response to the worry about safety was suggested to me by Dilip Ninan in conversation. Although our interpretation of (12-b) can be affected by how we think about close possible worlds, it does not automatically follow that the safety condition on knowledge is hostage to the same pragmatic forces. Even though (12-b) has multiple readings, perhaps only some of these readings should be identified as stating anything like the familiar safety condition on knowledge. For example, in order to evaluate whether Crito knows that he probably lost, we must evaluate whether his high credence that he lost is safe in the traditional sense. And in order to evaluate that, we should be in a context where we are thinking of close possible worlds as arbitrary close worlds where Crito believes he probably lost the lottery. Because Crito probably loses in all such worlds, his probabilistic belief is indeed safe. In the face of the worry that my account of safety may “concede too much to the skeptic,” this second response allows knowledge ascribers to stand their ground.

5 Relevant alternatives and probabilistic ignorance

In his thoughtful comments, Daniel Greco challenges me to describe the normative significance of probabilistic knowledge. Greco imagines that knowledge norms of action are meant to explain why we have negative rather than positive evaluations of some actions, such as your crossing the street on the basis of the belief that one person is more likely than another to assault you, when this belief is based entirely on statistics about racial groups. As Greco sees it, my account explains the wrongness of this action in terms of two facts: actions of this sort must be based on knowledge,

and your belief in this case isn't knowledge. Greco objects that this explanation leaves two questions unanswered—namely, when does knowledge matter for our evaluation of an action, and why do you lack knowledge in this case?

In discussing the first question, Greco acknowledges that we often evaluate actions in terms of their decision-theoretic features, such as whether actions maximize expected utility given your credences, or given the credences that are justified by your evidence. As Greco puts it, the question is “why, in cases like the ones above, is it proper to evaluate actions by reference to (probabilistic) knowledge norms, rather than decision theoretic ones?” As I see it, this question is misguided. Asking this question is like asking why we should evaluate some actions by reference to social norms rather than legal norms, or why we should evaluate some food by reference to its nutritional content rather than its taste, or why we should evaluate a textbook by reference to its breadth rather than its depth. As a normative pluralist, I take it that any of these evaluations may be perfectly appropriate to make. There are multiple dimensions along which we can properly evaluate an action, including the action of crossing the street.

The correct response to Greco's first question, then, is that my argument for the normative significance of probabilistic knowledge is different from the argument mentioned above. Knowledge norms are not meant to explain why your crossing the street is bad rather than good. Rather, knowledge norms help explain the weaker intuition that there is at least some respect in which this action is bad. This intuition is explained in terms of two facts: there is some respect in which actions are good only if they are based on knowledge, and your belief in this case isn't knowledge. To put it another way, suppose we grant for sake of argument that crossing the street has higher expected utility for you than not crossing the street.⁷ There would still be something bad about your acting on the basis of merely statistical evidence in this case. Knowledge norms are powerful in part because they enable us to explain this normative judgment, which cannot be explained by decision-theoretic norms alone.

This leaves Greco's second question: when it comes to the probabilistic belief that causes you to cross the street, why does this belief fail to constitute knowledge? Greco elaborates on this question in the following passage:

On Moss's view, context supplies a partition of logical space, and depending on what that partition is, either **Ignorance** or **Knowledge** can truly describe your situation. . . Why is it that, when discussing cases like the ones above, the proper

7. To be clear, I am not endorsing this assumption. However, advocates of racial profiling often take this claim for granted (cf. LEVIN 1992 for a notable example), so it is useful to understand how your action could be wrong even if the claim were true. The same goes for the claim that crossing the street has higher expected utility given your actual utility function and your evidential probabilities.

way to partition logical space—the *right* context—involves relative fine-grained partitions, rather than trivial ones?

A quick clarification is in order. My semantics entails that the epistemic modal ‘might’ is context sensitive. My theory of knowledge allows that the verb ‘knows’ is context sensitive. The latter context sensitivity is not necessarily connected with the former. According to the epistemic contextualist, ‘knows’ is interpreted relative to a set of relevant alternatives. These alternatives are probabilistic contents, not propositions. In general, these propositions will not correspond to any partition of logical space, including the partition used to interpret ‘might’ at a context. It might turn out that if we interpret ‘might’ relative to a fine-grained partition and can therefore truly say ‘Jones believes the marble might be .95 likely to be white’ at a context, then the content that the marble is .95 likely to be white will be a relevant alternative for ‘knows’ in that same context, and hence ‘Jones knows the marble is probably black’ will be false. But relevant alternatives for ‘knows’ are also generated in many other ways. For instance, the factivity of knowledge means that true contents are always relevant, whether or not the subject believes those contents might be true. Additional rules described by LEWIS 1996, such as the Rule of Resemblance, will further augment the set of relevant alternatives. To sum up, it is far from straightforward to give an algorithm for interpreting ‘knows’ at a context, and my semantics for ‘might’ is not intended to provide substantive guidance with this endeavor.

Why are certain alternatives relevant in some contexts and not others? The correct response is that contextualists should address this question in just the same way that they address it in the propositional case. We can give some general guidelines for relevance, but one should not expect comprehensive instructions for identifying relevant alternatives at an arbitrary context. To take a more concrete example, many have argued that relevance depends partly on what is at stake for the subject, but more specific definitions of practical stakes have met with several serious challenges. These challenges hold equally when it comes to moral stakes.⁸ In the end, I agree with Greco that “we’d like some story about just what sorts of generalizations about people are ones whose exceptions must be kept in mind,” just as I think it would be nice to have some story about what sorts of bank schedules are ones whose exceptions must be kept in mind when saying whether someone knows that their local bank is open on Saturday. As far as my theory is concerned, however, the relevant upshot is just that there is nothing special about probabilistic knowledge that makes telling this story any harder—or any easier.

8. For critical commentary on several precise notions of stakes sensitivity, see WORSNIP 2015 and ANDERSON & HAWTHORNE 2019. For relevant discussion of moral encroachment, see WORSNIP 2019.

6 Encroachment on full and partial belief

Greco ends by reflecting on why one might accept that knowledge is subject to pragmatic encroachment. He imagines motivating this view by saying, “To be a perfect Bayesian agent, you should act to maximize expected utility. Unfortunately, this norm is impossible for mere mortals to follow. Fortunately, you may come close to perfection by following another norm: do whatever has the greatest expected utility in the possibilities left open by your knowledge.” In order for this norm to yield sensible results, whether a possibility is ruled out by your knowledge must depend on whether including that possibility in your expected utility calculations makes a difference to what you should do. To paraphrase WEATHERSON 2012, the point is that if knowledge structures decision problems, then knowledge must be interest-relative.

This line of thought can support some of our intuitions about racial profiling. For example, suppose you are deciding whether ask Smith for a drink, and your credence that Smith is a waiter is based entirely on statistical evidence about his race.⁹ If the cost of being wrong about Smith is extremely high, then you should not set aside the possibility that he is not a waiter, nor should you act as if he is. Assuming that Smith is a waiter is wrong in a very traditional sense. It is just as irrational as taking a bet where you lose one hundred dollars if you roll snake eyes, and gain one dollar otherwise. As Greco sees it, pragmatic encroachment on propositional knowledge can be “defended as part of a good strategy for approximating expected utility maximization.”

As Greco points out, pragmatic encroachment on probabilistic knowledge can’t be defended in this way, which appears to be a problem for my view. The correct response is that knowledge norms of action are better motivated on alternative grounds. This fact is already taken for granted by advocates of traditional propositional knowledge norms. To illustrate, suppose that you are justified in having an astronomically high credence that Jones is a waiter, but your justified belief just happens to be false. A perfect Bayesian agent would act on the basis of her astronomically high credence. But since knowledge is factive, the possibility that Jones is not a waiter is left open by your knowledge, and so any knowledge norm of action will forbid you from ignoring that possibility. Hence even traditional propositional knowledge norms of action are rather poor surrogates for perfect Bayesian reasoning, and should be motivated on their own terms.

As I see it, knowledge norms for action are actually motivated by our realizing the limits of traditional Bayesian norms. For illustration, consider the following excerpt

9. For further discussion of similar examples in connection with the phenomenon of moral encroachment, see BOLINGER 2018, GARDINER 2018, HELLMAN 2018, and BASU 2019.

from the poem “BANG!” by Daniel Beaty:

When a woman clenches her purse and crosses the street. . .
Actually, ma’am, Coach is not my brand. I prefer Hermès.
BANG! You don’t know me.

When the clerk follows me through the store at Macy’s. . .
What are you looking at? I can buy ten of these leather coats if I want them.
BANG! You don’t know me.¹⁰

A woman crosses the street to avoid the speaker. A clerk follows him through the store. Both agents are acting on thoroughly probabilistic beliefs. The pedestrian does not fully believe that she is going to be assaulted, nor does the clerk fully believe that the speaker is going to shoplift. Rather, these agents are acting on credences that are based on merely statistical evidence. For present purposes, the crucial contribution of the poem is the identification of these agents as having thereby made some epistemic error. The speaker does not fault these agents for having failed to maximize utility given their credences, but rather for having those credences to begin with. As the speaker puts it, “You don’t know me,” and the relevant criticism is that there is something wrong with forming probabilistic beliefs from this position of ignorance. In short, there is indeed a familiar and intelligible sense in which credences based on statistical evidence can be epistemically bad when the moral stakes are high. This provides a more direct motivation for moral encroachment, independent of the instrumental usefulness of knowledge in promoting broadly Bayesian reasoning.¹¹

This more direct motivation for moral encroachment is compatible with a range of views about the scope of knowledge norms of action. For instance, one might endorse the modest conclusion that only certain actions must be based on knowledge—such as actions related to blame or punishment, for instance, or assertion, or the giving of reasons. Alternatively, one might endorse comprehensive knowledge norms of action. As in the case of traditional knowledge norms, fans of probabilistic knowledge may choose their own adventure here. The important upshot for my account is that knowledge norms can indeed be motivated as genuine alternatives to traditional Bayesian norms, rather than mere approximations of such norms, and these motivations apply just as straightforwardly to norms involving probabilistic knowledge.

10. BEATY 2014, p.90–1.

11. For further discussion of this motivation for moral encroachment, see Moss 2018a.

References

- ALONI, MARIA. 2000. "Quantification Under Conceptual Covers." PhD dissertation, ILLC/Department of Philosophy, University of Amsterdam.
- ANDERSON, CHARITY & JOHN HAWTHORNE. 2019. "Knowledge, Practical Adequacy, and Stakes." Ms., Department of Philosophy, Baylor University and University of Southern California. Forthcoming in *Oxford Studies in Epistemology*, vol. 6.
- BASU, RIMA. 2019. "What We Epistemically Owe to Each Other." *Philosophical Studies*, vol. 176 (4).
- BEATY, DANIEL. 2014. *Transforming Pain to Power: Unlock Your Unlimited Potential*. Berkley Books, New York.
- BOLINGER, RENEE. 2018. "The Rational Impermissibility of Accepting (Some) Racial Generalizations." Ms., Department of Philosophy, Australian National University.
- GARDINER, GEORGI. 2018. "Evidentialism and Moral Encroachment." Ms., Department of Philosophy, Oxford University. Forthcoming in *Believing in Accordance with the Evidence: New Essays on Evidentialism*, ed. Kevin McCain. Springer.
- HART, H. L. A. & TONY HONORÉ. 1985. *Causation in the Law*. Oxford University Press, Oxford, second edn.
- HELLMAN, DEBORAH. 2018. "The Epistemic Commitments of Nondiscrimination." Virginia Public Law and Legal Theory Research Paper No. 2018-60. Available at SSRN: <https://ssrn.com/abstract=3273582>.
- LEVIN, MICHAEL. 1992. "Responses to Race Differences in Crime." *Journal of Social Philosophy*, vol. 23 (1): 5–29.
- LEWIS, DAVID K. 1979. "Counterfactual Dependence and Time's Arrow." In *Philosophical Papers*, vol. 2, 32–51. Oxford University Press, Oxford.
- . 1996. "Elusive Knowledge." In *Papers in Metaphysics and Epistemology*, 418–46. Cambridge University Press, Cambridge.
- LYCAN, WILLIAM. 2001. *Real Conditionals*. Oxford University Press, Oxford.
- MACFARLANE, JOHN. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford.

- MOSS, SARAH. 2015. "On the Semantics and Pragmatics of Epistemic Vocabulary." *Semantics and Pragmatics*, vol. 8 (5): 1–81.
- . 2018a. "Moral Encroachment." *Proceedings of the Aristotelian Society*, vol. 118 (2): 177–205.
- . 2018b. *Probabilistic Knowledge*. Oxford University Press, Oxford.
- NINAN, DILIP. 2018. "Quantification and Epistemic Modality." *Philosophical Review*, vol. 127 (2): 433–85.
- SINNOTT-ARMSTRONG, WALTER. 2012. "Free Contrastivism." In *Contrastivism in Philosophy: New Perspectives*, MARTIJN BLAAUW, editor, 134–53. Routledge, New York.
- SLOTE, MICHAEL. 1978. "Time in Counterfactuals." *Philosophical Review*, vol. 87 (1): 3–27.
- VELTMAN, FRANK. 1996. "Defaults in Update Semantics." *Journal of Philosophical Logic*, vol. 25 (3): 221–61.
- WEATHERSON, BRIAN. 2012. "Knowledge, Bets, and Interests." In *Knowledge Ascriptions*, JESSICA BROWN & MIKKEL GERKEN, editors, 75–103. Oxford University Press, Oxford.
- WILLIAMSON, TIMOTHY. 2000. *Knowledge and its Limits*. Oxford University Press, Oxford.
- WORSNIP, ALEX. 2015. "Two Kinds of Stakes." *Pacific Philosophical Quarterly*, vol. 96 (3): 307–24.
- . 2019. "Can Pragmatists Be Moderate?" Forthcoming in *Philosophy and Phenomenological Research*.
- YALCIN, SETH. 2012. "A Counterexample to Modus Tollens." *Journal of Philosophical Logic*, vol. 41 (6): 1001–24.
- . 2015. "Epistemic Modality *De Re*." *Ergo*, vol. 2 (19): 475–527.