ON WHAT PROBABILISTIC KNOWLEDGE COULD NOT BE

Abstract. A critical reading of Sarah Moss's *Probabilistic Knowledge*.

1. Contextual and probabilistic theories of assertion and accuracy

In the first chapter of *Probabilistic Knowledge*[1], Sarah Moss makes two distinctions we must gloss. The first distinction hinges on whether expressions of the sort "the probability that P is .6" have the "simple content" P coupled with a tacked-on degree of belief .6, or whether they have "complex contents" that somehow amalgamate the proposition P and the attitude of believing to degree .6. Moss opts for the latter way of speaking, which I will stick to in this note. Specifically, she takes "the probability that P is .6" to be associated with the content $\{\mu : \mu(P) = .6\}$, and similarly for other sentences having apparent probabilistic content. More generally and to a first approximation, she takes the content of "*expression*" to be the set of probability measures that could correspond to the credence function of a speaker for whom "*expression*" is assertable. So, for example, "If Jones smokes, then Smith is at least .3 more likely to smoke than Brown" has, for Moss, content $\{\mu : \mu(S|J) \geq \mu(B|J) + .3\}$.

The second (and more substantive) distinction is manifested here:

> Suppose that Smith and Brown agree about exactly what credences are supported by any particular body of total evidence. Smith and Brown may nevertheless end up with different credences that Jones smokes, namely in virtue of having different evidence at their disposal. There is an intuitive sense in which Smith and Brown thereby count as disagreeing about the likelihood that Jones smokes. (...) The sense in which Smith and Brown disagree simply in virtue of having different credences is the sense that is relevant as you are forming your own credence that Jones smokes. It is the sense in which you...cannot agree with them both. [1, p. 15]

To put this passage in perspective, let us imagine that Smith and Brown know each other to be ideal rational agents, and have different credences in the proposition "Jones smokes". We will assume the uniqueness of rational priors, so Smith and Brown must have different evidence. By a result of Aumann, it is not the case that their credences are common knowledge. (Else these credences would coincide.)

[1]Though Moss's book inspired many replies, few of them (an exception is Edgington 2020, though punches are pulled there) make strong counter-arguments to Moss's claims. This note is unfinished (the author swore off writing philosophy mid-project), unedited, and the examples it resorts to are often not simple. This is a testament to the resourcefulness–albeit bent on nursing radicalities about the nature of probability, manifested in an onslaught against its through-and-through relativism in particular–of Moss. The result is tedious, and will test the patience of most readers. Good luck.

However, we can assume that, say, Brown knows Smith's credence. Note that the assertability conditions of $P =$ "the probability that Jones smokes is .6", for Smith, are the same as the assertability conditions of $Q =$ "Smith's total evidence is such as to yield a credence of .6 that Jones smokes"[2]. On one interpretation, then, P simply means that Q. But, says Moss, there is another interpretation on which P and Q are not interchangeable–a sense on which Brown can agree that Q, but disagree with P.

On the first interpretation Brown takes P to be assertive, whereas in the latter interpretation Brown takes P to be performative.[3] claim Compare $Z =$ "we find in favor of the plaintiff". One can read Z as an assertion with which one agrees ("they did, in fact, find in favor of the plaintiff"), or as a performative (establishing of a verdict) with which one disagrees ("knowing what I do, I would find in favor of the defendant"). So, here, one can read $P =$ "the probability that Jones smokes is .6" as an assertion with which one agrees ("Smith's credence is, in fact, .6"), or as a performative (establishing of a forecast) with which one disagrees ("knowing what I do, I would put the probability that Jones smokes at .0001".) Note, however, that in disagreeing with "the probability that Jones smokes is .6" one is not disagreeing with Smith's credences, but with the forecast established by Smith's utterance. Similarly, in disagreeing with "we find the defendant not guilty", one is disagreeing with the verdict, rather than with the collective private beliefs of the jurors, all of whom may in fact believe the defendant to be guilty. (Cf. the murder trial of O.J. Simpson.)

Moss appears to speak of the assertoric mood as that in which "speakers use epistemic modals to assert propositions about contextually determined bodies of evidence, such as the total evidence possessed by certain contextually relevant subjects". More generally, the assertoric mood may be used to assert facts about the ideal interpretation of one's evidence, or simply to assert facts about one's own interpretation of one's evidence, such as the distribution one holds for the ideal epistemic probability conditional on said evidence or expectations thereof (i.e. one's credences).

That the assertoric mood does not cover all uses, however, is shown by the following example, which Moss adapts from Egan (2007).

> (1) *Criminal in Paris*: It is unlikely that James Bond is in London.
>
> (2) *Eavesdropper in London*: No it's not–Bond is almost certainly in London.

The idea here is that if (1) were interpretted to mean saying something like "my evidence is such as to suggest that Bond is unlikely to be in London", then "No it's not" wouldn't be sanctioned. The intuition is that it is sanctioned, however, hence (1) is not being so interpretted. However, the aptness of "No it's not" degrades in nearby, nearly synonymous iterations.

[2] Or the assertibility conditions of "Smith's credence that Jones smokes is .6" if Smith is not ideal.

[3] Moss does not speak of "assertive" and "performative" utterances, and might in fact not agree with these characterizations. In particular, she might not agree that the factive reading of a report of one's credences has the nature of a forecast. To deny this, however, seems to commit one to a naive materialism about probabilities, giving them an objective "out there" status, rather than seeing them as an internal artefact of one's (partial) ignorance. This may be our whole disagreement.

(3) *Criminal in Paris*: It's .1 likely that James Bond is in London.

(4) *Eavesdropper in London*: #No it's not–Bond is almost certainly in London.

(4) sounds bad because Bond being in London is consistent with the criminal's forecast, which in fact explicitly allows for it. The following, in contrast, sounds fine:

(5) *Eavesdropper in London*: That's inaccurate–Bond *is* in London.

(Note: the inaccuracy of the criminal's forecast may be quantified via logarithmic scoring as -$\log_2(.1)$ bits should Bond be in London, and -$\log_2(.9)$ bits otherwise.) In (5), however, the eavesdropper professes absolute confidence that Bond is in London. Without this absolute confidence, the claim of inaccuracy requires hedging.

(6) *Eavesdropper in London*: #That's inaccurate–Bond is almost certainly in London.

(6) sounds bad because the eavesdropper should be saying "That's almost certainly inaccurate" instead. However, (7) seems fine despite the absence of hedging:

(7) *Eavesdropper in London*: I disagree–Bond is almost certainly in London.

Buoyed by (1) and (2), Moss writes: "Arguments against contextualist theories of epistemic vocabulary support probabilistic theories of assertion." She claims support for her own theory in particular, wherein one can "assert probabilistic contents". That we can assertnon-propositional contents is, however, a radical thesis. If we can avoid crisis without resorting to radical theses, we should do so. Happily, the above discussion demonstrates that (2)'s aptness as a response to (1) is an artifact of propositional (i.e. non-probabilistic) reasoning. The criminal's phrasing implicitly suggests the *proposition* that Bond is not in London, to which the eavesdropper says no–he *is* in London. That's why the relevant intuition degrades when the criminal formulates an explicit probability that Bond is in London. What was previously available for the eavesdropper to *contradict* wasn't the criminal's forecast (forecasts can't be contradicted), but the proposition that Bond isn't in London. The most the eavesdropper can say against the forecast is that it's *inaccurate* (if it is known that Bond is in London) or *disagreed with* (if the eavesdropper's credences are different).

To summarize: Moss takes the choice here to be between a "propositional content view", according to which all assertions have propositional contents, and a "probabilistic content view", according to which some assertions have probabilistic contents. Moss takes token utterances employing "likely", "probably" etc. to be instances of assertions having probabilistic content, for the reason that such utterances can be coherently contradicted by eavesdroppers, or for the reason that one may later wish to retract such an utterance, the idea being that if such utterances were assertions having (contextual) propositional contents, these phenomena would not be observed.

But while Moss is correct that (1) isn't an assertion about what attitude the criminal's evidence supports, that's because (1) isn't an assertion at all. To assert a content is to present that content as putatively true, and credences are not truth apt. They don't have truth values, but rather degrees of accuracy. (1) is a hedged prediction about what the world is like, a prediction that, if made precise–as in (3)–would constitute a forecast, which could be graded for accuracy. The default intuitions about (3)-(7) are decisive in favor of this way of thinking.

2. Some semantic funambulism

Moss seeks to establish that (1) one can believe probabilistic contents, (2) one can assert probabilistic contents, and finally (3) one can know probabilistic contents. The radicalness of her view is already evident in her defense of (1), as she writes: "there is no scale of attitudes such that maximal degree of belief is at one end of the scale" (Moss 2017 p. 87). In particular, credence is not degree of belief, and to have a full belief is not to have an extreme credence. To the contrary, according to Moss, there is simply the attitude of belief, and one may bear it toward propositional or probabilistic contents.

Although this notion of belief is radical, it is at least easy to understand. To believe that it's .6 likely that Jones smokes just is to have .6 credence that Jones smokes. When it comes to (3), however, that we can know probabilistic contents, Moss really needs to scramble. For knowledge is factive, and it isn't at all clear what it would even mean for Smith's .6 credence that Jones smokes to be true.

Indeed, by the time Moss offers up the following seeming concession, it seems that the game must surely be up: "If Jones doesn't smoke, then it isn't at least .6 likely that she smokes, and Smith doesn't know that it is at least .6 likely that she smokes" (Moss 2017, p. 182). If that's so, one reasons, then by the same token, if Jones does smoke, then it isn't at most .6 likely that she smokes, and Smith doesn't know that it is at most .6 likely that Jones smokes. So since either Jones smokes or doesn't smoke, Smith can't both know that it's at least .6 likely that she smokes and know that it's at most .6 likely that Jones smokes. So she can't, in particular, know that it's precisely.6 likely that Jones smokes. *Mutatis mutandis*, she can't know that it's x likely that Jones smokes for any $x \in (0, 1)$.

In Chapter 7 of her book Moss plays chicken with such arguments, armed with a novel semantics (presented in Chapter 3) for epistemic expressions. Taking a cue from S. Kaufmann's (2004) so-called "local interpretation" of indicative conditionals, Moss associates context-supplied partitions to epistemic expressions for use in unpacking their semantic contents. For example, consider the conditional $A \to C$, with contextual partition $\mathcal{P}$. Here A and C are probabilistic contents, i.e. sets of probability measures. Then the content of $A \to C$ consists in those probability measures ϕ such that $\phi|_p$ is in the content C for every $p \in \mathcal{P}$ with $\phi|_p$ in the content A. Here $\phi|_p$ is the normalized restriction of ϕ to p, i.e. the result of conditioning ϕ on p. Other expressions receive similar treatments, at times with counterintuitive results. E.g., (8) and (9) may be consistent if parsed with different partitions:[4]

[4] The basic idea is that one parses (8) using the trivial partition, and (9) with a non-trivial partition $\{P_1, P_2\}$, where $1 > \mu(P_1) > .5$ and $P_2 \cap J = \emptyset$. Note however that though (8) and (9) are consistent owing to the different contextual partitions, Moss's semantics does not admit their conjunction as consistent (which would require supplying them with the same contextual partition).

(8) It is possible that Jones smokes

(9) It is probably impossible that Jones smokes

It's even possible to construct incompatible-sounding yet consistent sets of expressions all of which are parsed with the same partition. For example suppose that exactly one of Jones, Smith has the first name "Mary" (we have no evidence which). Then Moss's semantics allows for a consistent same-partition reading of the following triad[5]

(10) It is probable that Jones is a likely hire

(11) It is probable that Smith is a likely hire

(12) It is improbable that the candidate named "Mary" is a likely hire

Let us return to the troubling example above, however, the first part of which Moss formulates as a putative instance of constructive dilemma, thus (Moss, p. 207):

(13) a. Either Jones smokes or Jones doesn't smoke
b. If Jones smokes, then it's not .6 likely that Jones smokes
c. If Jones doesn't smoke, then it's not .6 likely that Jones smokes
d. Therefore, it is not .6 likely that Jones smokes

Moss notes that, according to her semantics, this is not an instance of constructive dilemma, and is in fact not a valid inference. The reason she gives is equivocation: in (13) b. and (13) c., the contextual partition used in unpacking the contents of the conditional is taken to be at least as fine grained as $\{J, \neg J\}$, whereas in (13) d. this is not so for "likely" and "not". On the other hand, consider an alternate path, all of whose contextual partitions are trivial:

(14) a. Either Jones smokes or Jones doesn't smoke
b. If it's .6 likely that Jones smokes and either Jones smokes or Jones doesn't smoke, then either it's .6 likely that Jones smokes and Jones smokes or it's .6 likely that Jones smokes and Jones doesn't smoke
c. It is not the case that (either it's .6 likely that Jones smokes and Jones smokes or it's .6 likely that Jones smokes and Jones doesn't smoke)
d. Therefore, it is not the case that (it's .6 likely that Jones smokes and either Jones smokes or Jones doesn't smoke)
e. Therefore, it is not .6 likely that Jones smokes

This time equivocation on "or" is the saving grace. Not owing to a difference between absolute and conditional probabilities (only absolute probabilities are involved, all of the contextual partitions being trivial), but because, as Moss notes (Moss, p. 268):

[5]We can for example take context to supply a partition $\{R, T, S\}$ (for research, teaching and service) having equally likely cells. If research or teaching is most important, Jones is two-thirds likely to be hired, but is a definite non-hire if service is most important. If teaching or service is most important, Smith is two-thirds likely to be hired, but is a definite non-hire if research is most important. The candidate named Mary (note the *de dicto* reading) is two-thirds likely to be hired if teaching is most important, but only one-third likely otherwise.

> According to my semantics, logical operators are polymorphic, taking arguments of multiple semantic types. As operators on sentential constitutents denoting propositions, 'or', 'and' and 'not' have their traditional intensional semantic values. But as operators on sentences denoting probabilistic contents, these operators (map sets, or ordered pairs of sets of probability spaces to sets of probability spaces).

Hence the "or" appearing in a. operates on propositions, whereas in b.–d. "or" operates on probabilistic contents.

A third try circumvents the semantic acrobatics:

(15) { Case 1: a_1. Jones smokes
b_1. If Jones smokes then it isn't .6 likely that Jones smokes
c_1. Therefore, it isn't .6 likely that Jones smokes}
{Case 2: a_2. Jones doesn't smoke
b_2. If Jones doesn't smoke then it isn't .6 likely that Jones smokes
c_2. Therefore, it isn't .6 likely that Jones smokes}
d. Therefore, it isn't .6 likely that Jones smokes

It's going to cost Moss something to answer this putative proof that it isn't .6 likely that Jones smokes, because the meta-language in which the argument is framed...that used to express that this a proof by cases, that either Case 1 or Case 2 obtains, and that in either of the given cases, d. follows...is beyond the reach of her semantic tricks. I in fact was unable to find any satisfactory way to respond to the argument in a manner consistent with Moss's other commitments. She already endorses sentences such as b_1 and b_2, and is explicit that "whenever it would be right to assign a proposition as the content of a belief, we may simply assign the corresponding nominally probabilistic content as the strict content of that belief. (...) probabilistic contents should ultimately replace propositions as the fundamental contents of belief." (Moss, p. 22). Such passages commit her to nominally probabilistic content readings "Jones smokes' and "Jones doesn't smoke" in a_1 and a_2, precluding any claim of equivocation.

One might think that it would cost Moss little to simply ban the nominally probabilistic content readings of "Jones smokes" and "Jones doesn't smoke" in a_1 and a_2. She fights like hell against this, though, explicitly rejecting it in passages like "...one might be tempted to adopt a disjunctive semantics, according to which sentences can have either propositional or probabilistic contents. However tempting it might be, this disjunctive theory is ultimately unacceptable." (Moss, p. 83.) Why does she fight so hard against it? If she adopted such a disjuctive semantics she could at once claim claim equivocation from a_1 to b_1; the former instance of "Jones smokes" would indicate a proposition, while the latter instance would indicate a probabilistic content. Thus the inference to c_1 would not be a valid instance of modus ponens.

Another option would be to ban the propositional readings of a_1 and a_2. This would undercut the whole argument, as then the two considered cases would no longer be exhaustive.[6] I think that's what she has to do, but it's not an attractive option. For

[6]The issue is that as propositions, "Jones smokes" and "Jones doesn't smoke" are exhaustive, i.e. these are only two possible cases. (That's why the above "proof by cases" works.) But as

"If Jones smokes then it isn't .6 likely that Jones smokes" is a controversial, evocative and substantive claim when "Jones smokes" is interpreted propositionally, but a bland, uncontroversial truism when "Jones smokes" is interpreted as a probabilistic content. And, apart from taking a lot of the seeming bite away from her claims, such a concession would undercut much of what Moss has to say about the "factivity" of probabilistic knowledge. More about that in a later section, though.

3. An argument from closure under entailment

Curiously, Moss does not resort to a semantics-based defense for one of the skeptical arguments she considers. This is the "argument from closure" against probabilistic knowledge, which she formulates thus:

> (16) a. If Smith knows that it's .6 likely that Jones smokes, then Smith knows that it is not the case that Jones certainly doesn't smoke
> b. Smith doesn't know that it is not the case that Jones certainly doesn't smoke
> c. Therefore, Smith doesn't know that it's .6 likely that Jones smokes

A semantics-based defense was available. Indeed "it's .6 likely that Jones smokes", from the conclusion, is apparently being parsed with the trivial partition. So that has to be the parsing in the first premise's antecedent as well, else the conclusion doesn't follow. But in that case, in the consequent of the first premise "not" must be parsed with the trivial partition (else the first premise would not be accepted). But the "not" in the second premise requires that "not" be parsed with some partition at least as fine as $\{J, \neg J\}$. So we have equivocation; what Smith knows in the first premise's consequent doesn't agree, contentwise, with what Smith fails to know in the second premise. Hence this is not a valid instance of *modus tollens*.

What are we to make of the fact that Moss doesn't go for the semantics-based argument here, whereas everywhere else she does? Well, here the premises are complicated somewhat by the presence of "knows that" operators. Here is my guess at Moss's concern. Observe that "If Smith believes that it might be the case that Jones certainly doesn't smoke, then Smith doesn't know that it is not the case that Jones certainly doesn't smoke" sounds innocuous. However, if we parse "not" in the consequent of this conditional using the trivial partition (just what allows us to reject premise b. in the semantic defense just glossed), then we have to parse "might" in the antecedent with the trivial partition as well in order to get this seemingly innocuous conditional to contain our credences. This seems wrong, though, because it makes "it might be the case that Jones certainly doesn't smoke" equivalent to "Jones certainly doesn't smoke", meaning that the conditional becomes "If Smith believes Jones certainly doesn't smoke, then Smith doesn't know that it is not the case that Jones certainly doesn't smoke", to which any epistemologist worth her salt would surely say "yes...but in fact, the same conclusion follows even if Smith only believes that it *might* be the case that Jones certainly doesn't smoke".

---

nominally probabilistic contents, they do not exhaust. Indeed, the disjunction of these contents is $\{\mu : \mu(J) \in \{0, 1\}\}$, which does not contain any function assigning J a non-extreme credence.

The upshot seems to be that the "knows that" operator in "Smith knows that it is not the case that Jones certainly doesn't smoke" automatically opens up the possibility that "not" should be parsed with some non-trivial partition. Which non-trivial partition? That's where Moss draws on a perceived analogy between (16) and Dretske's zebra thought experiment:

> (17) a. If Smith knows that the caged animial in front of him is a zebra, then Smith knows that it is not the case that the animal is a cleverly disguised mule.
> b. Smith doesn't know that it is not the case that the animal is a cleverly disguised mule.
> c. Therefore, Smith doesn't know that the caged animial in front of him is a zebra.

The arguments (16) and (17) are not, however, analogous. "Smith knows that P" has a strict content and a loose content. The strict content is that Smith's credence in P is 1, that Smith's credence is justified, and that P is true.[7] The loose content, by contrast, is based on a loose concept of knowledge that is not closed under entailment. In particular, you can't loosely say you know that P in a case where your credence in P is both less than 1 and not interestingly greater than some baseline, i.e. if you have no interesting evidence that P. It's perfect consistent that you could have a lot of interesting evidence that the animals are zebras and yet (because all of this evidence is consistent with the skeptical hypothesis) no interesting evidence that they aren't cleverly disguised mules.

To see the disanalogy, suppose one tiles the proposition "the animal is not a zebra" by skeptical hypotheses. One of these could be "the animal is a cleverly disguised mule". Perhaps another is "the animal is a mutant donkey that looks just like a zebra, mistakenly placed in the zebra exhibit by a rookie zoo worker". Suppose that we had no interesting evidence against any of these skeptical hypotheses. Then, since the hypotheses fully tile, we would have no interesting evidence against "the animal is not a zebra". But we do have such evidence. So we can't tile the original hypothesis by such skeptical hypotheses. The fact is general. One cannot tile a proposition one knows to be false by hypotheses, each of which one does not know to be false. Moreover, *this is true whether one actively entertains the tiling hypotheses or not.*

But that is just what Moss claims is going on in (16). She writes "Even though it is perfectly consistent to believe that it is .6 likely that Jones smokes, Smith cannot rule out that this content is something that God does not believe." However, this understates the matter. In fact, Smith *is certain that this content is something that God does not believe*. For either God believes that it is certain that Jones smokes, or God believes that it is certain that Jones doesn't smoke. So once Smith so much as entertains this fact, i.e. adopts $\{J, \neg J\}$ as textual partition, Smith becomes certain that it is either certain that Jones smokes, or certain that Jones doesn't smoke. *A fortiori*, she becomes certain that it is not the case that Jones is .6 likely to smoke.

[7] Where P is a finitary event (the sort whose truth could be established in bounded expected time) it's plausible that truth of P is entailed by Smith's justified full credence; there are, that is to say, no finitary miracles. (Note that these views preclude existence of true defeaters as well.)

Moss concedes that the mere fact that Smith cannot rule out that Jones certainly doesn't smoke precludes her knowing that it is .6 likely that Jones smokes in contexts where the prospect that Jones certainly smokes is "relevant". She writes:

> If you are talking with an expert who knows whether Jones smokes, then you will often pay attention to the content that Jones certainly smokes, as well as the content that Jones certainly doesn't smoke, since you realize that the expert believes one of these two contents. As a result, you will fail to satisfy 'knows that it is .6 likely that Jones smokes' as uttered in this context. The same goes for contexts where you believe that an expert might be eavesdropping on your conversation, and even for some contexts where you simply imagine that an expert might say that your probabilistic beliefs are false.

If correct, this would be bizarre, for it would imply that in order to know thoroughly probabilistic contents, one would have to "blind oneself to the elephant in the room", namely, that the propositions concerning which one has thoroughly probabilistic beliefs *do have truth values*. Fortunately, it isn't correct. For again, it isn't merely that Smith, upon turning her gaze to the elephant, cannot rule out some skeptical content. Smith is actually certain, whenever she does so, that either Jones certainly smokes or that Jones certainly doesn't smoke. These scenarios tile the tautology, and as a result it doesn't matter whether she looks at the elephant or not.

4. Moss's argument against an error theory

As mentioned earlier, Moss's use of contextual partitions recalls Kaufmann's (2004) local interpretation of conditionals. And indeed, Moss devotes some space to arguing against "The Thesis", i.e. against the principle that the "probability" (ersatz or otherwise) of "if A, then B" is just $P(B|A)$. Like the local interpretation, however, Moss's position is subject to strong objections. She considers a case where Jones, a thrill-seeker who likes to jump into the painters' net from the roof of her apartment building, is currently on the roof of the building. We think it likely (though aren't sure) that the net is not currently up. We think it very likely (though aren't completely sure) that Jones isn't suicidal, and assume for present purposes that she will die if and only if she jumps and there is no net. Moss considers the restriction that we have promised the painters not to tell anyone directly whether or not the net is up. Someone asks us whether there is currently a net. Given the restriction on our possible replies, Moss sanctions the following, where the conditional is parsed using contextual partition $\{\text{net}, \text{no net}\}$:

> (20) I can't answer your questions directly. But I can tell you this much: it's highly probable that if Jones jumps off the roof, she'll die.

Moss writes:

> Expressions of confidence in conditionals are context sensitive. In some contexts, 'it is highly probable that if Jones jumps, she will die' may have a content that you believe just in case you have high conditional credence that Jones will die if she jumps. But in other contexts, this sentence has a different probabilistic content.

This argument is specious. When we are prevented from saying something semantically equivalent to what we mean to say, as here, it is natural to resort to pragmatic considerations in deciding which to utter among various things that are not semantically equivalent to what we mean to say. In this example the audience is primed for information confirming or disconfirming "there is a net". But even a Thesis defender will agree that "it's highly probable that if Jones jumps off the roof, she'll die" confirms "no net". For either it's highly probable that if Jones jumps off the roof she'll die or it isn't, and it's clearly the former that leads one to assign a higher posterior probability to "no net". So given the restrictions we are under and the pragmatic aim of casting light on "net or no net", (20) looks fairly assertible to the Thesis defender as well, even if it isn't at all true. (It wrongly implies, says the Thesis fan, that a suicidal dive is more likely than a leap taken just for kicks.)

We can emphasize the point by using a traditional proposition rather than a conditional. Suppose Bill lives and works as maintenance man in the building we're talking about. Bill gets unhappy when the painters leave their net up. Suppose it's well know that Bill is across town in a meeting with his boss, where he will probably lose his job, and that this would make him unhappy, as well as cause him to become indifferent about the net. Consider two responses to the query about the net:

> (21) I can't answer your questions directly. But I can tell you this much: it's highly probable that Bill will be happy when he gets home.
> (22) I can't answer your questions directly. But I can tell you this much: it's highly probable that Bill won't be happy when he gets home.

In the current context, where we believe there's probably no net and want to communicate this indirectly, "Bill will be happy when he gets home" is more assertible than "Bill will not be happy when he gets home". The latter, in particular, seems willfully misleading. That the former is probably the case, however, isn't true. Notwithstanding that, the thing (21) gets wrong (it's wrongly optimistic about Bill's mood) isn't something our audience cares about. All they care about is the net. The thing (20) gets wrong (it wrongly indicates greater credence in a suicidal dive than a thrill-seeking leap) is, similarly, something our audience doesn't care about.[8]

Elsewhere I've introduced the following principle, ostensibly somewhat weaker than the full-blown Thesis, that can be use to critique some purported counterexamples.

> *Equivalence*. Let A, B, C and D be propositions such that A if and only if B and, conditional on A, C if and only if D. Then the conditionals $A \to C$ and $B \to D$ must be assigned the same "probability".

According to *Equivalence*, we must assign equal "probability" to "If Jones jumps, she'll die" and "If Jones jumps, there's no net." It isn't plausible that we should assign the latter high probability, however, so *Equivalence* precludes assigning the former high probability. Moss requires us to give up this highly plausible principle. The

[8]Further support for Moss's intuitions can be found in the fact that our audience likely knows that we know that they know that Jones isn't more likely to be suicidal than the net is to be up anyway. In short...they know what we mean despite our choice of words.

cost, however, is excessive. More economically, we should insist that it's pragmatic considerations, not semantic considerations, that underlie her stated intuitions.

5. The Sly Pete Argument

Not content with her previous arguments against The Thesis, Moss writes:

> ...if our assertion of sentences like ("it's highly probable that if Jones jumps off the roof, she'll die") is a cognitive error, then sentences like these pose no threat to Stalnaker's Hypothesis or Adams' Thesis. To forestall error theories...I want to introduce a novel problem for the alleged connection between probabilities of conditionals and conditional probabilities, namely a problem concerning probabilistic inferences. Imagine that in the Sly Pete case, you listen to each of your henchmen and you trust them up to a point. On the basis of their reports, you reason:
>
> (23) a. It is at least .9 likely that if Pete called, he won.
> b. It is at least .9 likely that if Pete called, he lost.
> c. Therefore, it is as least .8 likely that Pete did not call.

Recall the setup, from (Gibbard 1981):

> Sly Pete and Mr. Stone are playing poker on a Mississippi riverboat. It is now up to Pete to call or fold. My henchman Zack see's Stone's hand, which is quite good, and signals its content to Pete. My henchman Jack sees both hands, and sees that Pete's hand is rather low, so that Stone's is the winning hand. At this point, the room is cleared...Zack knows that Pete knew Stone's hand. He can thus appropriately assert "If Pete called, he won." Jack knows that Pete held the losing hand, and thus can appropriately assert "If Pete called, he lost."

What are we to make of (23)? Moss's semantics sanctions the inference from a. and b. to c., basically because Moss's semantics interprets "if Pete called he won" as a strict conditional (so that in particular, it comes out true if Pete didn't call), and the inference is obviously correct under that interpretation.[9] Moss promises to expose a problem for the thesis that the ersatz probability of an indicative conditional is a conditional probability...i.e. a problem for the Adams conditional. This argument is more or less that some non-Adams conditional, namely the strict one, accounts for the argument (23), whereas the Adams conditional does not.

But is the argument any good? The move from "trusting your henchmen to a point" to assigning .9 credence to the contents of their avowals is mysterious. Would this move be available if said contents were thoroughly probabilistic? Assume for example that the henchmen are more cautious, claiming only that "If Pete called, he probably

[9] In cases where the contextual partition is not trivial, the conditional need only be "piecewise strict".

won" and "If Pete called, he probability lost".[10] It's tempting to then reason as follows:

(24) a. It is at least .9 likely that if Pete called, he probably lost.
b. It is at least .9 likely that if Pete called, he probably won.
c. Therefore, it is as least .8 likely that Pete did not call.

The inference from (24) a. and b. to (24) c. is valid on Moss's semantics, so this argument is sound if the premises are justified. Are they? They had better not be. For suppose instead that Zack has seen Stone's hand and Jack has seen Pete's hand, and that both hands are extremely high. Assume moreover that Pete called before the room was cleared, but that the players hadn't yet revealed their cards. Letting T denote the tautology, Zack can appropriately assert "If T is true, Pete probably lost", and Jack can appropriately assert "If T is true, Pete probably won." Buying into (24), you could reason:

(25) a. It is at least .9 likely that if T is true, Pete probably lost.
b. It is at least .9 likely that if T is true, Pete probably won.
c. Therefore, it is as least .8 likely that T is not true.

Obviously (25) is a disaster. Apparently, then, we have to reject (24) a. and b. So on what principle should we accept (23) a. and b.? The only notable difference between (23) and (24) is that "If Pete called, he won" and "If Pete called, he lost" have nominal probabilistic contents, whereas "if Pete called, he probably lost" and "if Pete called, he probably won" have thoroughly probabilistic contents. Citing this as justification for acceptance of (23) a. and b. is, however, problematic. Consider a situation in which a fair coin has been tossed. Zack says that the coin didn't land heads. Jack says that the coin didn't land tails. These messages don't contradict one another (the coin may have landed on edge, or flown out of the atmosphere), and in any event there is nothing too implausible about receiving them...we only trust Zack and Jack to a point, after all. So reasoning on the model of (23), one has:

(26) a. It is at least .9 likely that the coin didn't land heads.
b. It is at least .9 likely that the coin didn't land tails.
c. Therefore, it is at least .8 likely that the coin landed neither heads nor tails.

Clearly this is a bad conclusion. So it can't be that we accept (23) a. and b. because "If Pete called, he won" and "If Pete called, he lost" have nominal probabilistic contents. Why, then, should one accept these premises? Moss doesn't give any promising justification for them at all. Absent this, the claim that (23) makes trouble for the Adams interpretation looks completely impotent.[11]

[10]A first blush recourse for Moss: maybe this is not the correct way for the henchmen to hedge. Perhaps instead they should say "Probably, if Pete called he won" and "Probably, if Pete called he lost". But according to Moss's semantics one believes "Probably, if Pete called he won" whenever one knows Pete to have the worst possible hand. ("If Pete called he won" is material on Moss's semantics, hence true whenever Pete doesn't call...and he's unlikely to call, given that he can't possibly win.) So this idea can be rejected out of hand.

[11]Edgington (2020) summarizes the situation beautifully: "what is obvious on the Ramsey-Adams view, but lost on Moss's theory: Zack and Jack disagree. Each rejects a conditional the other accepts.

6. Factivity and Relativism

Moss takes pains to distance herself from the charge of relativism in passages such as the following, where she stresses that knowledge is not mere justification.

> Here is a third proposal: you know a content just in case it contains every probability space that represents what you may rationally believe, *and you believe that content*. ...the third proposal fails to predict that knowledge entails truth. (...) According to the third proposal, it is perfectly consistent to say that Smith knows that Jones probably drinks, and that Brown knows that she probably doesn't. (...) These results are unacceptable. Knowledge is not false. Knowledge is not inconsistent. The intended subject of this book is a factive attitude....

So if two agents believe inconsistent contents, says Moss, it must be the case that one of these beliefs fails to be knowledge. Which? A second passage sheds some light:

> Suppose that Bond is not in London. Then the content that Bond is probably in London is false. But according to the relativist, the content may nevertheless be true for a criminal in virtue of her limited epistemic reach. By contrast, my probabilistic theory of epistemic vocabulary does not make use of any relation of relative truth. The criminal may be justified in believing the content that Bond is probably in London. But the content of her belief is not true for her, or indeed true in any sense. The content is inconsistent with the true content that Bond is not in London, and that alone is enough to guarantee that the content is false.

Together, the two passages give an answer. For if two probabilistic contents are inconsistent, then at least one of them (perhaps both) is inconsistent with a true content, namely the content assigning full probability to the actual state of the world. That alone is enough to guarantee, according to the second passage, that the content is false. And according to the first passage, knowledge is not false. So that content cannot be knowlege. There is our answer. And for Moss, who wants some "intermediate" probabilistic contents to constitute knowledge, it's not tenable, for no intermediate credence is consistent with the content assigning full probability to the actual state.

How did Moss come to defend such an untenable collection of views? Part of the answer can be found in her conviction that relativism falls prey to the eavesdropper arguments. Another is her conviction that rightness is an epistemic good. Consider for example the following, from (Moss 2020):

There is not way you can simply take over their conditionals. The disagreement is faultless, but disagreement it is."

> To see what the factivity condition rules out, consider the following utterances:
>
> a. *Jones*: It will 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 does not 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.

Unless her credence in rain was zero, how can Smith have "got it right"? Suppose her credence in rain was .25. Then compared to Brown, who had credence .1 in rain, she got it *wrong*. Indeed Brown got it wrong, too. For "if something is false, then it is not more than 05 likely, and no one ever knew that it was". Is this reasonable? Plainly not. Brown got it more right than Smith, who in turn got it more right than Jones. Yes, there is a factive reading of "the probability that Jones smokes is .6"...a reading whose veracity depends on whether or not Jones actually smokes, irrespective of the ideal credences. However, as I stressed way back in the first section of this paper, these "thoroughly probabilistic contents" aren't assertions but forecasts, and their veracity (their *accuracy*) isn't an all-or-nothing affair.

The difficulties faced by the view that these contents are rather assertions, subject to all-or-nothing factivity, are likely unanswerable. Moss's ingenuity and resourcefulness in marshalling an heroic defense notwithstanding, that view doesn't work. Probabilistic knowledge, if there is such a thing, has its factivity underwritten not by truth, but by accuracy. A proper investigation of such a notion would begin with the question "when are credences justified", progressing to "how can we measure degree of justification", eventually arriving at "is justified credence subject to epistemic defeat, and if so, how is it measured". Failing to address these matters, Moss hasn't properly begun the project that the title of her book would appear to have promised.[12]

References

[1] Edgington, Dorothy. 2020. Credence, Conditionals, Knowledge and Truth. *Analysis*. 80(2):332–342.

[2] Moss, Sarah. 2016. *Probabilistic Knowledge*. Oxford University Press.

[3] Moss, Sarah. 2020. Replies to Edgington, Pavese, and Campbell-Moore and Konek. *Analysis* 80(2):356-370.

[12] The plan was to write additional sections on Chapters 8-10, the applications sections of Moss's book. I can't imagine what could induce me to type a single character of philosophy ever again, however. For the peculiar, rare reader who just has to know what these sections might have looked like, I have posted (unlisted) my live, cold readings of these chapters on YouTube, beginning at https://www.youtube.com/playlist?list=PLZboVlfdXLv4nIL6DRL3cGVqYEjUicR54