Philosophers' Imprint

VOLUME 21, NO. 26
NOVEMBER 2021

EPISTEMIC MODAL CREDENCE

Simon Goldstein

Australian Catholic University

Abstract

Triviality results threaten plausible principles governing our credence in epistemic modal claims. This paper develops a new account of modal credence which avoids triviality. On the resulting theory, probabilities are assigned not to sets of worlds, but rather to sets of information state-world pairs. The theory avoids triviality by giving up the principle that rational credence is closed under conditionalization. A rational agent can become irrational by conditionalizing on new evidence. In place of conditionalization, the paper develops a new account of updating: conditionalization with normalization.

1. Introduction

Ordinary language includes many claims that describe what the world is like: *it will rain later*, *the die will land on 6*, etc. Rational agents have various attitudes towards these descriptive claims. These include coarse-grained attitudes such as belief and certainty. They also include fine-grained attitudes such as having a .2 credence that the die will land on 6.

But ordinary language also includes various claims that do not seem to directly describe the way the world is, but instead convey something about our epistemic relation to it. For example, we have conditional claims like *if the die is rolled, it will land on 6*, as well as modal claims such as *the die might land 6*.

Ordinary agents adopt propositional attitudes toward these claims as well. Moreover, it seems that an agent's attitudes toward these claims should bear some relation to her attitudes toward the descriptive claims. For example, if an agent is certain that the die will land hands, she should also be certain that the die must land heads. This paper develops systematic principles governing the relationship between our attitudes towards descriptive and epistemic modal claims.

The direct focus of this paper is a well studied puzzle about the probability of epistemic modal claims. There are many controversial inferences involving epistemic modals, including the Direct Argument

($\neg A \vee B$ implies $A \rightarrow B$), Lukasiewicz's principle (A implies $\Box A$), and right nested Modus Ponens (A and $A \rightarrow (B \rightarrow C)$ imply $B \rightarrow C$). Each inference has a pair of puzzling properties: the inference preserves certainty, but fails to preserve probability, so that a rational agent can be less confident in the conclusion than the premises. But these two properties seem incompatible.

This paper develops a theory of epistemic modal credence to resolve this puzzle within an information sensitive semantics for epistemic modals, where meanings are sets of information state-world pairs. We first develop a theory of epistemic modal credence which provides a formal characterization of what it takes for various inferences to preserve probability or preserve certainty. In order to divorce these two properties, the resulting theory gives up the assumption that conditionalization preserves rationality: a rational agent may become irrational if she conditionalized on a new piece of information. To give up this assumption, we go on to develop a new method of updating: conditionalization with normalization. When an agent learns a claim, she doesn't just narrow down her credence function to the state-world pairs at which the claim is true. In addition, she must normalize her credence function, by removing any world where the claim is false from any state she takes to be possible. With our new update procedure in place, we go on to resolve the puzzle with which we begin. The problematic inferences above are shown to preserve certainty but not probability. To avoid triviality, the proposal rejects the principle that the set of rational probability functions are closed under conditionalization upon evidence. Rather, the set of rational probability functions is closed under the new update procedure of conditionalization with normalization. In the final sections of the paper, we explore a particular thesis about modal credence: that agents should be maximally opinionated in modal claims, assigning any such claim a credence of 0 or 1.

2. A puzzle about modal credence

Consider the inference in (1):

(1) a. Either the gardener or the butler did it.
b. Therefore, if the gardener didn't do it, the butler did.

The inference in (1) is called 'Or-to-If', or 'The Direct Argument'.[1]

(2) **The Direct Argument.** $\neg A \vee B \models A \rightarrow B$

The reasoning seems impeccable. But the inference is also puzzling. Its validity seems to require that the indicative conditional $\neg A \rightarrow B$ is no stronger than the material conditional $A \vee B$. Given the validity of Modus Ponens, this in turn implies that the indicative conditional and the material conditional are logically equivalent.

As observed by Edgington (1995), Schulz (2017), and Santorio (2018a), the Direct Argument is also puzzling from a probabilistic perspective. Consider a scenario in which we have three suspects: butler, gardener, and cook. The gardener is the most likely suspect, followed by the cook. By contrast, the butler loved his master and so was very unlikely to commit the crime. In this case, the premise (1-a) is very likely, while the conclusion (1-b) is very unlikely. This suggests that the Direct Argument is invalid.

But now suppose that we discover that the cook possessed an airtight alibi. In this setting, we are certain of the premise (1-a), and so the conclusion (1-b) follows with certainty. In short, the Direct Argument preserves certainty, but not probability. This is puzzling, since the former property suggests that the Direct Argument is valid, while the latter suggests it is invalid.

Schulz (2010), Beddor and Goldstein (2018), and Santorio (2018a) note that an analogous paradox affects epistemic modal claims. Consider now:

(3) a. The house is empty.
b. Therefore, the house must be empty.

1. For discussion, see Stalnaker (1975), Jackson (1987), Edgington (1995), Bennett (2003), and Cariani (Forthcoming) among others.

This inference is called Lukasiewicz's principle.[2]

(4) **Lukasiewicz's Principle.** $\mathsf{A} \models \Box\mathsf{A}$

Its validity is controversial. As with the Direct Argument, this inference fails to preserve probability. Suppose Ari the burglar has been casing the house for hours. As far as she can tell, not a mouse is stirring. Consequently, Ari believes the house is empty (or at least has a quite high credence). Still, Ari is an experienced burglar; she knows that even the most thorough reconnaissance is fallible. Thus she allows that there's some possibility—albeit very remote—that an inconspicuous resident is still inside. For this reason, she does not believe that the house must be empty (and her credence in this claim is low). Here we again have an inference where an agent's credence in the premise is lower than her credence in the conclusion. On the other hand, imagine that Ari became certain of the premise (3-a), by accessing security footage of the interior. In this case, she would also become certain of the conclusion (3-b).[3]

For a third example, let's turn to instances of Modus Ponens containing right nested conditionals. Suppose a fair six-sided die is about to roll:

(5) a. The die landed even.
b. If the die landed even, then if it didn't land 2 or 4, it landed 6.
c. Therefore, if the die didn't land 2 or 4, it landed 6.

Santorio (2018a) observes that this case is a potential probabilistic counterexample to Modus Ponens.

(6) **Right Nested Modus Ponens.** $\mathsf{A}, \mathsf{A} \rightarrow (\mathsf{B} \rightarrow \mathsf{C}) \models \mathsf{B} \rightarrow \mathsf{C}$

2. For initial discussion see Yalcin (2007).
3. Here we abstract from the role of *must* in signaling indirect evidence. For more on this feature of *must*, see for example von Fintel and Gillies (2010).

In this case, we have a fairly high credence in (5-a) and (5-b), but a low credence in the conclusion (5-c). Again, this counterexample disappears under conditions of certainty. If we become certain of (5-a), we will immediately become certain of the conclusion (5-c).

We now have three examples of inferences which preserve rational certainty but not rational credence. In each case, Santorio (2018a) precisifies the puzzle above by developing a triviality result, showing that if the inference above preserves certainty, it must also preserve rational credence.

Let Pr be a probability function defined over a simple propositional language enriched with the epistemic modals *might* and *must*, and the indicative conditional. Above, we saw that three inferences concerning epistemic modals seem to be certainty preserving for any rational agent's credence function:

(7) Where Pr is a rational probability function:
a. If $Pr(\neg\mathsf{A} \vee \mathsf{B}) = 1$, then $Pr(\mathsf{A} \rightarrow \mathsf{B}) = 1$.
b. If $Pr(\mathsf{A}) = 1$, then $Pr(\Box\mathsf{A}) = 1$.
c. If $Pr(\mathsf{A} \wedge (\mathsf{A} \rightarrow (\mathsf{B} \rightarrow \mathsf{C})) = 1$, then $Pr(\mathsf{B} \rightarrow \mathsf{C}) = 1$.

Now assume that the class of rational credence functions is closed under conditionalization, so that whenever a probability function is rational, the result of conditionalizing it on a piece of evidence is also rational:

(8) **Closure Under Conditionalization.** If Pr is a rational probability function, then $Pr(\cdot \mid \mathsf{A})$ is a rational credence function.

Given these assumptions, Santorio (2018a) shows that each of the inferences above is probability preserving.[4] First consider the Direct Argument.

(9) a. $Pr(\mathsf{A} \rightarrow \mathsf{B}) = Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) \times Pr(\neg\mathsf{A} \vee \mathsf{B}) + Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg(\neg\mathsf{A} \vee \mathsf{B})) \times Pr(\neg(\neg\mathsf{A} \vee \mathsf{B}))$

4. For a generalization of this style of argument, see Santorio (2021).

b. $Pr(\mathsf{A} \to \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) = 1$
c. $Pr(\mathsf{A} \to \mathsf{B}) = Pr(\neg\mathsf{A} \vee \mathsf{B}) + Pr(\mathsf{A} \to \mathsf{B} \mid \neg(\neg\mathsf{A} \vee \mathsf{B})) \times Pr(\neg(\neg\mathsf{A} \vee \mathsf{B}))$
d. So $Pr(\mathsf{A} \to \mathsf{B}) \geq Pr(\neg\mathsf{A} \vee \mathsf{B})$

We start with an application of the Law of Total Probability to some rational probability function Pr: the probability of $\mathsf{A} \to \mathsf{B}$ is a weighted sum of its probability conditional on $\neg\mathsf{A} \vee \mathsf{B}$ and $\neg(\neg\mathsf{A} \vee \mathsf{B})$, weighted by the prior probability of $\neg\mathsf{A} \vee \mathsf{B}$. But given Closure under Conditionalization, we know that $Pr(\cdot \mid \neg\mathsf{A} \vee \mathsf{B})$ is itself a rational probability function. Since $Pr(\neg\mathsf{A} \vee \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) = 1$, we can infer that $Pr(\mathsf{A} \to \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) = 1$ by the assumption that the Direct Argument preserves certainty. But then applying the Law of Total Probability, we know that the current probability of $Pr(\mathsf{A} \to \mathsf{B})$ is $Pr(\neg\mathsf{A} \vee \mathsf{B})$ plus some value. So the Direct Argument's conclusion $\mathsf{A} \to \mathsf{B}$ must have a probability at least as high as the probability of its premise $\neg\mathsf{A} \vee \mathsf{B}$.

Similar arguments show that Lukasiewicz's Principle and Right Nested Modus Ponens are probability preserving if they are certainty preserving.

(10) a. $Pr(\Box\mathsf{A}) = Pr(\Box\mathsf{A} \mid \mathsf{A}) \times Pr(\mathsf{A}) + Pr(\Box\mathsf{A} \mid \neg\mathsf{A}) \times Pr(\neg\mathsf{A})$
b. $Pr(\Box\mathsf{A} \mid \mathsf{A}) = 1$
c. So $Pr(\Box\mathsf{A}) \geq Pr(\mathsf{A})$

(11) a. $Pr(\mathsf{B} \to \mathsf{C}) = Pr(\mathsf{B} \to \mathsf{C} \mid \mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C}))) \times Pr(\mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C}))) + Pr(\mathsf{B} \to \mathsf{C} \mid \neg(\mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C})))) \times Pr(\neg(\mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C}))))$
b. $Pr(\mathsf{B} \to \mathsf{C} \mid \mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C}))) = 1$
c. So $Pr(\mathsf{B} \to \mathsf{C}) \geq Pr(\mathsf{A} \wedge (\mathsf{A} \to (\mathsf{B} \to \mathsf{C})))$

To solve these problems, we need to develop a general theory of epistemic modal probability. The rest of this paper does exactly that, in a way that explains how each of the inferences above can preserve certainty without preserving probability. Crucially, the theory will reject the principle of Closure Under Conditionalization, by developing a precise alternative to conditionalization.

The main goal of this paper is to develop an account of epistemic modal credence that explains the puzzles above. But before continuing, it is worth flagging another reason one might want such an account. Moss (2015) and Charlow (2020) develop theories of epistemic modals that are designed to explain failures of the automatic iteration of epistemic vocabulary. For example, following Moss (2015), Charlow considers the following triad:

(12) a. It is probably the case that Trump might be impeached.
b. It is probably the case that Trump will be impeached.
c. Trump might be impeached.

Charlow suggests that (12-a) is weaker than (12-b) or (12-c), for example because believing (12-b) licenses betting at fair odds that Trump will be impeached, while believing (12-a) does not.

The goal of the current paper is to model directly how agents assign credences to epistemic modal claims, rather than to provide a semantics for the probabilistic language under which epistemic modals embed. Nonetheless, the data above provides at least some reason to think that agents can have interesting degrees of belief in epistemic modal claims.[5] On the other hand, we'll see below (in sections 8 and 9) that one promising way of developing our theory requires that epistemic claims always have a probability of 0 or 1, in which case sentences like (12-a) would require further study.

5. Charlow (2020) differs from the current paper by offering a polymorphic semantics for epistemic modals, on which epistemic modal claims are assigned semantic values of a different type than ordinary claims. For this reason, the probability of an epistemic modal claim cannot directly be compared with the probability of an ordinary claim. As a result, Charlow (2020) does not model the kinds of substantive bridge principles connecting these states that is the subject of this paper.

3. Information sensitivity

In order to make sense of modal credence, we'll help ourselves to an information sensitive semantics for modal expressions.[6] On this proposal, the meaning of any sentence is not a set of worlds, but rather a set of pairs of information states and worlds. While ordinary claims are sensitive to the world component of meaning, modal claims are sensitive to the information state.

We will interpret a propositional language enriched with three epistemic modals: a possibility modal $\Diamond$, a necessity modal $\Box$, and an indicative conditional operator $\rightarrow$.

Definition 1. $\mathsf{L} ::= \mathsf{p} \mid \neg\mathsf{L} \mid \mathsf{L} \wedge \mathsf{L} \mid \mathsf{L} \vee \mathsf{L} \mid \Diamond\mathsf{L} \mid \Box\mathsf{L} \mid \mathsf{L} \rightarrow \mathsf{L}$

In Yalcin (2007), information states are represented as sets of possible worlds. Then an interpretation function assigns each sentence a set of state-world pairs as its meaning. We can then define $s + \mathsf{A}$, the update of state s with sentence A, in terms of intersecting s with $\mathbf{A}$, the set of state-world points where A is true. Finally, A is supported by s, or true throughout it, just in case $s = s + \mathsf{A}$, so that A is true at $\langle s, w\rangle$ for every $w \in s$.

Definition 2.

1. W is a set of possible worlds which assign a truth value $w(\mathsf{p})$ to each atomic sentence p.
2. $S = \mathcal{P}(W)$ is a set of information states, all subsets of W.
3. An interpretation function $[\![\cdot]\!]$ assigns a set of pairs of information states and worlds to every sentence in L.
4. $\mathbf{A} = \{\langle s, w\rangle \mid [\![\mathsf{A}]\!]^{s,w} = 1\}$
5. $s + \mathsf{A} = s + \mathbf{A} = \{w \in s \mid \langle s, w\rangle \in \mathbf{A}\}$
6. A is true throughout s ($[\![\mathsf{A}]\!]^{s} = 1$) iff $\forall w \in s : [\![\mathsf{A}]\!]^{s,w} = 1$.

The next step is to define our interpretation function. Atomic sentences like *it's raining* are true at an $\langle s, w\rangle$ pair if and only if they are true at w. The negation $\neg\mathsf{A}$ is true at $\langle s, w\rangle$ iff A is false there. The conjunction $\mathsf{A} \wedge \mathsf{B}$ is true at $\langle s, w\rangle$ iff each conjunct is true there. The disjunction $\mathsf{A} \vee \mathsf{B}$ is true at $\langle s, w\rangle$ iff either disjunct is true there. Epistemic modals are sensitive to the information state parameter. Possibility and necessity modals existentially and universally quantify over worlds in an information state. The indicative conditional universally quantifiers over the worlds in the information that results from updating with the antecedent.

Definition 3.

1. $[\![\mathsf{p}]\!]^{s,w} = 1$ iff $w(\mathsf{p}) = 1$
2. $[\![\neg\mathsf{A}]\!]^{s,w} = 1$ iff $[\![\mathsf{A}]\!]^{s,w} = 0$
3. $[\![\mathsf{A} \wedge \mathsf{B}]\!]^{s,w} = 1$ iff $[\![\mathsf{A}]\!]^{s,w} = 1$ and $[\![\mathsf{B}]\!]^{s,w} = 1$
4. $[\![\mathsf{A} \vee \mathsf{B}]\!]^{s,w} = 1$ iff $[\![\mathsf{A}]\!]^{s,w} = 1$ or $[\![\mathsf{B}]\!]^{s,w} = 1$
5. $[\![\Diamond\mathsf{A}]\!]^{s,w} = 1$ iff $\exists v \in s : [\![\mathsf{A}]\!]^{s,v} = 1$
6. $[\![\Box\mathsf{A}]\!]^{s,w} = 1$ iff $[\![\mathsf{A}]\!]^{s} = 1$
7. $[\![\mathsf{A} \rightarrow \mathsf{B}]\!]^{s,w} = 1$ iff $[\![\mathsf{B}]\!]^{s+\mathsf{A}} = 1$

In order to apply this semantics to the puzzles about probability with which we began, we must formulate a notion of entailment. Within this system, several different notions are natural. First, one might define entailment as simple preservation of truth: whenever a state-world pair makes the premises true, it also makes the conclusion true. Second, one might restrict this definition to state-world pairs that are *proper*, so that the world is itself a member of the state. For example the inference from $\Box\mathsf{A}$ to A does not preserve truth at all state-world pairs; but it does preserve truth at proper points. Finally, one could instead define entailment as preservation of truth at a state, so that whenever every world in a state makes the premises true relative to that state, those

6. There are several different kinds of information sensitive frameworks. For representative work within update semantics, see Veltman (1996), Groenendijk et al. (1996), Beaver (2001), Gillies (2004), and Willer (2013), along with a technical appendix of this paper, where the information sensitive semantics in the main text is compared to update semantics. For other information sensitive accounts, see Yalcin (2007), Swanson (2011), Klinedinst and Rothschild (2012), Swanson (2012), Moss (2015), Yalcin (2015), Ninan (2018), and Mandelkern (2019). For useful overviews of the literature, see von Fintel and Gillies (2007) and Willer (2015).

worlds also make the conclusion true relative to that state. Summarizing:

Definition 4.

1. $\mathsf{A}_1, \ldots, \mathsf{A}_n$ classically entail B ($\mathsf{A}_1, \ldots, \mathsf{A}_n \models_c \mathsf{B}$) iff for every s, w, if $[\![\mathsf{A}_1]\!]^{s,w} = 1, \ldots,$ and $[\![\mathsf{A}_n]\!]^{s,w} = 1$, then $[\![\mathsf{B}]\!]^{s,w} = 1$.
2. $\mathsf{A}_1, \ldots, \mathsf{A}_n$ properly entail B ($\mathsf{A}_1, \ldots, \mathsf{A}_n \models_p \mathsf{B}$) iff for every s, w where $w \in s$, if $[\![\mathsf{A}_1]\!]^{s,w} = 1, \ldots,$ and $[\![\mathsf{A}_n]\!]^{s,w} = 1$, then $[\![\mathsf{B}]\!]^{s,w} = 1$.
3. $\mathsf{A}_1, \ldots, \mathsf{A}_n$ informationally entails B ($\mathsf{A}_1, \ldots, \mathsf{A}_n \models_i \mathsf{B}$) iff for every s, if $[\![\mathsf{A}_1]\!]^{s} = 1, \ldots,$ and $[\![\mathsf{A}_n]\!]^{s} = 1$, then $[\![\mathsf{B}]\!]^{s} = 1$.

To illustrate these different notions of entailment, let's return to our three data points. Each one is an informational entailment but not a classical or proper entailment.

Observation 1.

1. (a) $\neg\mathsf{A} \vee \mathsf{B} \not\models_c \mathsf{A} \to \mathsf{B}$
   (b) $\neg\mathsf{A} \vee \mathsf{B} \not\models_p \mathsf{A} \to \mathsf{B}$
   (c) $\neg\mathsf{A} \vee \mathsf{B} \models_i \mathsf{A} \to \mathsf{B}$
2. (a) $\mathsf{A} \not\models_c \Box\mathsf{A}$
   (b) $\mathsf{A} \not\models_p \Box\mathsf{A}$
   (c) $\mathsf{A} \models_i \Box\mathsf{A}$
3. (a) $\mathsf{A}, \mathsf{A} \to (\mathsf{B} \to \mathsf{C}) \not\models_c \mathsf{B} \to \mathsf{C}$
   (b) $\mathsf{A}, \mathsf{A} \to (\mathsf{B} \to \mathsf{C}) \not\models_p \mathsf{B} \to \mathsf{C}$
   (c) $\mathsf{A}, \mathsf{A} \to (\mathsf{B} \to \mathsf{C}) \models_i \mathsf{B} \to \mathsf{C}$

For illustration, consider Lukasiewicz's Principle. Let $s = \{w, v\}$ with p true at w and false at v. p is true at proper $\langle s, w\rangle$, but $\Box\mathsf{p}$ is false there. So Lukasiewicz's Principle does not preserve truth at proper points. However, when $\Box\mathsf{A}$ is true throughout s, A is true at every world in s, and so A is true throughout s as well.

Our target principles are informationally valid but classically and properly invalid. For this reason, Santorio (2018a) suggests that informational validity preserves certainty, while classical and proper validity preserves probability. The rest of this paper develops exactly this suggestion, by showing how to assign probability to epistemic modal claims.

4. Modal credence

The basic idea is simple. Instead of assigning probability to worlds alone, we assign probability to state-world pairs. The probability of a claim, modal or otherwise, is just the probability of the state-world pairs at which it is true.

We are about to develop a series of rational constraints on how an agent should assign probabilities to sets of pairs of information states and worlds. One natural question before beginning this task is what such probabilities even are. What is it for an agent to assign credences to objects of this kind? Sections 8-10 of this paper develop one answer to this question. On that proposal, an agent's credences in sets of state-world pairs supervene on her credences in sets of possible worlds. If we accept this claim, we can rely on any old understanding of ordinary credence (say, as betting dispositions) to interpret credence in epistemic modal claims.

On the other hand, much of what we say below is also compatible with the thesis that an agent's credences in sets of state-world pairs do not supervene on her credences in sets of possible worlds. Indeed, the reductive account in §8-9 requires that any given agent assign credence to only one information state at a time (although see §10 for a way of avoiding this commitment). Holding fixed the semantics for epistemic modals we just reviewed, this requires that an agent's credences in epistemic modal claims are extreme. Those who wish to reject the reductive account below must find another interpretation of what it is to assign credences to sets of information state - world pairs. Here, one option would be to rely on recent work in Charlow (2020). To make sense of credence in epistemic modal claims, Charlow suggests that rational agents have primitive conditional preferences to act in certain ways given various epistemic modal hypotheses. Rational agents can prefer to drink water from a lake if it must not be infected; but can prefer not to drink the water if it might be infected. At any rate, we now set aside questions of interpretation and return to our main task

of studying the rational constraints on epistemic modal credence.

To assign credences to information state-world pairs, we introduce the notion of an epistemic space. An epistemic space is just a probability function defined over state-world pairs.

Definition 5. $E = \langle W', F', Pr \rangle$ is an epistemic space, where:

1. $W' \subseteq W$ is a set of worlds.
2. $F' = \mathcal{P}(W' \times \mathcal{P}(W'))$ is the set of all state-world pairs constructed from W'.
3. (a) Pr is a probability measure over F'.
   (b) $Pr(\langle s, w \rangle)$ abbreviates $Pr(\{\langle s, w \rangle\})$.
   (c) For any sentence A in L, $Pr(\mathsf{A})$ abbreviates $Pr(\mathbf{A})$.

We sometimes use W'_E, F'_E, and Pr_E to abbreviate the W', F', and Pr component of epistemic space E.

An epistemic space provides a measure over different state-world pairs. Given such a measure, it will be important in what follows to extract a notion of a state, world, or state-world pair being live. Given an epistemic space, the live points in the space are those state-world pairs which the space assigns positive probability. The live worlds in the space are exactly those assigned positive probability by that space when paired with some state. The live states are those assigned positive probability when paired with some world.[7]

Definition 6.

1. $\mathsf{Live}(E) = \{\langle s, w \rangle \mid Pr_E(\langle s, w \rangle) > 0\}$ is the set of live points in E.
2. $\mathsf{LiveWorlds}(E) = \{w \mid \exists s : \langle s, w \rangle \in \mathsf{Live}(E)\}$ is the set of live worlds in E.
3. $\mathsf{LiveStates}(E) = \{s \mid \exists w : \langle s, w \rangle \in \mathsf{Live}(E)\}$ is the set of live states in E.

Throughout, we'll illustrate the main ideas with a pair of working examples. In the first example, the agent treats two incompatible states as live.

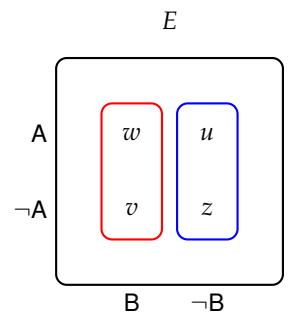

state	world	Pr_E
red	w	1/4
red	v	1/4
blue	u	1/4
blue	z	1/4

Here, the agent distributes her credence over combinations of information states and worlds. In particular, she considers w, v, u, and z to be the live worlds, and considers red and blue to be the live states. One state implies B and one implies ¬B. Our agent is indifferent between which state and which world obtains; her credences over various combinations of states and worlds is given by the stochastic truth table above. For example, her credence in $\Box\mathsf{B}$ is equal to her credence in $\{\langle \mathsf{red}, w \rangle, \langle \mathsf{red}, v \rangle\}$, which is $\frac{1}{2}$. This is equal to her credence in B.

In our second example, the agent again treats two states as live; but now the states overlap.

7. Here we assume that an agent is certain of a claim (so that it holds at every live point) iff her credence in that claim is 1. But there are a variety of counterexamples to this thesis involving infinite partitions (see Hájek (2017)). For example, imagine throwing an infinitely small dart at the real number line between 0 and 1. The probability that the dart lands at $\frac{1}{2}$ is 0. But an agent should treat this possibility as live, since she can't be certain it won't obtain. To solve this problem, we could enrich our models that an epistemic space includes not only a probability function Pr_E but also a set of live points L_E where $Pr_E(L_E) = 1$. Our major results below could then be recast by systematically replacing $\mathsf{Live}(E)$ below, defined in terms of Pr_E, with L_E. For simplicity, we suppress this complication in what follows.

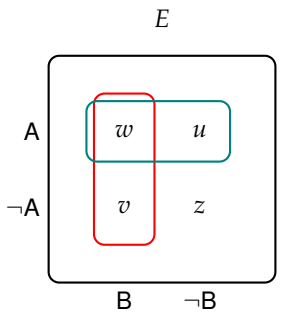

state	world	Pr_E
red	w	1/4
red	v	1/4
teal	w	1/4
teal	u	1/4

In this example, our agent considers only w, v, and u to be the live worlds, and considers red and teal to be the two live information states. One state implies B and one implies A. Our agent is indifferent between which state obtains, but not about the world. Since world w shows up in both red and teal, the agent is twice as confident in w as she is in u or v.

Our second example illustrates a failure of Lukasiewicz's Principle to preserve probability. Here, $Pr(\Box\mathsf{A}) = Pr(\{\langle \mathsf{teal}, w\rangle, \langle \mathsf{teal}, u\rangle\}) = Pr(\langle \mathsf{teal}, w\rangle) + Pr(\langle \mathsf{teal}, u\rangle) = \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$. By contrast, $Pr(\mathsf{A}) = Pr(\{\langle \mathsf{teal}, w\rangle, \langle \mathsf{teal}, u\rangle, \langle \mathsf{red}, w\rangle\}) = Pr(\langle \mathsf{teal}, w\rangle) + Pr(\langle \mathsf{teal}, u\rangle) + Pr(\langle \mathsf{red}, w\rangle) = \frac{1}{4} + \frac{1}{4} + \frac{1}{4} = \frac{3}{4}$. Similar examples generate similar failures of the Direct Argument and Right Nested Modus Ponens to preserve probability.

We have now generated an epistemic space where our three inferences of interest fail to preserve probability. Our next task is to give a systematic theory of when various epistemic spaces allow various notions of entailment to preserve probability or certainty.

5. Credence and consequence

In this section, we'll build on Santorio (2018a) to establish three main connections between credence and consequence. First, the framework above guarantees that classical consequence preserves probability. Second, there is a well behaved class of epistemic spaces on which proper consequences preserves probability. Third, there is another well behaved class of epistemic spaces where informational consequence preserves certainty; but this last class is not closed under conditionalization.

The first immediate result of the framework above is that classical consequence preserves probability. Suppose that A classically implies B, and consider some epistemic space $E = \langle W, F, Pr\rangle$. Then $\mathbf{A} \subseteq \mathbf{B}$. So every point in $\mathbf{A}$ assigned probability by Pr is also in $\mathbf{B}$. So $Pr(\mathsf{A}) \leq Pr(\mathsf{B})$. The same observation extends to the case of multiple premises. Here, let the uncertainty associated with A be 1 minus the probability of A. Then an inference preserves probability iff the uncertainty of the conclusion never exceeds the sum of the uncertainty of the premises.[8] In the system above, classical consequence preserves probability.

Definition 7. Classical consequence preserves probability in E iff $\mathsf{A}_1, \ldots, \mathsf{A}_n \models_c \mathsf{B}$ implies that $1 - Pr_E(\mathsf{B}) \not> \sum_1^n 1 - Pr_E(\mathsf{A}_i)$.

Observation 2. Classical consequence preserves probability in every epistemic space.

Proof. Suppose $\mathsf{A}_1, \ldots, \mathsf{A}_n \models_c \mathsf{B}$. Then the conjunction $\mathsf{A}_1 \cdots \wedge \ldots \mathsf{A}_n \models_c \mathsf{B}$, and so B is true at every point at which the conjunction $\mathsf{A}_1 \cdots \wedge \ldots \mathsf{A}_n$ is true, and so its probability must be at least as high as $\mathsf{A}_1 \cdots \wedge \ldots \mathsf{A}_n$. But the probability of $\mathsf{A}_1 \cdots \wedge \ldots \mathsf{A}_n$ must itself be no lower than the sum of the distance from 1 of each conjunct. □

The next natural question is whether proper consequence preserves probability.

Definition 8. Proper consequence preserves probability in E iff $\mathsf{A}_1, \ldots, \mathsf{A}_n \models_p \mathsf{B}$ implies that $1 - Pr_E(\mathsf{B}) \not> \sum_1^n 1 - Pr_E(\mathsf{A}_i)$.

There are epistemic spaces in which this fails. But only under one condition. Say that $\langle s, w\rangle$ is proper just in case $w \in s$. Say that an epistemic space is proper just in case every live point is proper.

Definition 9. E is proper iff $\langle s, w\rangle \in \mathsf{Live}(E)$ only if $w \in s$.

8. See Adams (1998) and Field (2015).

Proper consequence corresponds to an epistemic space being proper.

Observation 3. Proper consequence preserves probability in E iff E is proper.

Proof. Suppose that E is not proper. Then there is some live $\langle s, w\rangle$ with $w \notin s$. Now let A be true throughout s but false at w. Then since $\Box\mathsf{A} \supset \mathsf{A}$ is false at $\langle s, w\rangle$, $Pr(\neg(\Box\mathsf{A} \supset \mathsf{A})) > 0$. But $\neg(\Box\mathsf{A} \supset \mathsf{A}) \models_p \bot$, and yet $Pr_E(\bot) = 0$, and so proper consequence doesn't preserve probability in E.

Conversely, suppose that E is proper, and that $\mathsf{A}_1, \ldots, \mathsf{A}_n \models_p \mathsf{B}$. Since E is proper, Pr_E only assigns positive probability to $\langle s, w\rangle$ where $w \in s$. At all such points where the premises are true, the conclusion is true; so proper consequence preserves probability in E. $\Box$

For illustration, consider the inference from $\Box\mathsf{A}$ to A, which is classically invalid but properly valid. Suppose the agent's epistemic space is as follows:

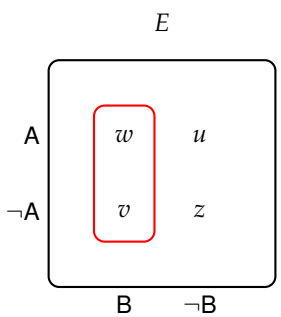

state	world	Pr_E
red	w	1/4
red	v	1/4
red	u	1/2

Here, the agent assigns positive probability to the pair $\langle \mathsf{red}, u\rangle$, even though u is not included within red. In this model, the agent is certain of $\Box\mathsf{B}$, since she only assigns positive credence to the state red. Nonetheless, her credence in B is only $\frac{1}{2}$. This is because she assigns $\frac{1}{2}$ credence to the point $\langle \mathsf{red}, u\rangle$. So this agent does not treat proper consequence as probability preserving.

If we think that proper consequence does preserve the credence of rational agents, then we can make propriety a constraint on rational credence functions. On this proposal, an agent is rational only if her epistemic space is proper. (In an appendix, we explore a method of encoding this requirement into the semantics for any expression, in a way that provides further motivation for requiring that any epistemic space be proper.)

For comparison with what follows, it's worth noting that if we require the class of rational epistemic spaces to be proper, we can still accept Closure Under Conditionalization. The class of proper epistemic spaces is itself closed under conditionalization on new pieces of evidence, so that whenever Pr is proper, $Pr(\cdot \mid \mathsf{A})$ is guaranteed to also be proper. We'll see in a moment that other constraints are not compatible with Closure Under Conditionalization.

Now let's turn to informational consequence. Our goal here is to develop a constraint on epistemic spaces which ensures that information consequence preserves certainty, so that whenever an agent is certain of the premises, she is certain of the conclusion.

Definition 10. Informational consequence preserves certainty in E iff $\mathsf{A}_1, \ldots, \mathsf{A}_n \models_i \mathsf{B}$ implies that if $Pr_E(\mathsf{A}_i) = 1$ for every $i \in n$, then $Pr_E(\mathsf{B}) = 1$.

This principle does not hold in every epistemic space. Instead, it requires that the epistemic space be plentiful, so that whenever a state is live, every world in the state is assigned positive credence when paired with the state.

Definition 11. E is plentiful just in case if $s \in \mathsf{LiveStates}(E)$, then $\forall v \in s : Pr(\langle s, v\rangle) > 0$.

Suppose that the agent has eliminated $\langle s, w\rangle$ as a possible state-world pair, with $w \in s$. Then at best she considers $s - w$ as a live state; not s itself. We can think of this constraint as a level bridging principle, connecting an agent's attitude towards states with her attitude towards state-world pairs.

More precisely, preservation of information consequence turns out

to be equivalent to the conjunction of being proper and plentiful.[9]

Observation 4. Informational consequence preserves certainty in E iff E is proper and plentiful.

Proof. Suppose E is proper and plentiful. Suppose $\mathsf{A}_1, \ldots, \mathsf{A}_n \underset{i}{\models} \mathsf{B}$. Suppose that $Pr_E(\mathsf{A}_i) = 1$ for every $i \in n$. Then every live $\langle s, w\rangle$ is in $[\![\mathsf{A}_i]\!]$ for every $i \in n$. Since E is plentiful, this implies that A_i is true throughout every live s (since whenever s is live, $\langle s, w\rangle$ is live for all $w \in s$). So B is true throughout every live s. Since E is proper, this implies that every live $\langle s, w\rangle$ is in $[\![\mathsf{B}]\!]$ (since $\langle s, w\rangle$ is live only if $w \in s$.)

Suppose E is either not proper or not plentiful. We must show that in either case informational consequence fails to preserve certainty. Now suppose E is improper. Then there is some live $\langle s, w\rangle$ with $w \notin s$. Let $[\![\mathsf{q}]\!] = \{\langle s', w'\rangle \mid s' \in \mathsf{LiveStates}(E), w' \in s'\}$, and let $[\![\mathsf{p}]\!] = [\![\mathsf{q}]\!] \cup \{s, w\}$. p informationally entails q, since whenever p is true throughout s, q is too. Furthermore, $Pr(\mathsf{p}) = 1$. But $Pr(\mathsf{q}) = Pr(\mathsf{p}) + Pr(\langle s, w\rangle)$, and so $Pr(\mathsf{q}) \neq 1$.

Now suppose E is not plentiful. Then there is a live s^* where $w \in s^*$ and $\langle s^*, w\rangle$ is not live. Let $X = \{\langle s, w\rangle \in \mathsf{Live}(E) \mid s \neq s^*\}$. Let $Y = \{\langle s^*, w\rangle \mid \langle s^*, w\rangle \in \mathsf{Live}(E)\}$. Let $[\![\mathsf{p}]\!] = X \cup Y$, and let $[\![\mathsf{q}]\!] = X$. p informationally implies q, since they are true throughout the same states ($\mathsf{LiveStates}(E) - s^*$). $Pr(\mathsf{p}) = 1$, since $[\![\mathsf{p}]\!] = \mathsf{Live}(E)$. But $Pr(\mathsf{q}) < 1$, since $Pr(Y) > 0$. □

Santorio (2018a) identifies the property of plenitude and introduces a class of models where informational consequence preserves certainty. But the relevant models all assign a probability of 1 to a single information state (a principle we will critically explore and ultimately defend in the final section), and for this reason Santorio (2018a) does not establish a necessary condition for informational consequence to preserve certainty.

We've now said enough to explain our general strategy for solving our initial puzzle. We propose that every rational epistemic space is proper and plentiful. For this reason, rational agents should treat inferences like the Direct Argument, Lukasiewicz's Principle, and Right Nested Modus Ponens as preserving certainty but not probability.

To implement this strategy, however, we need to say something about the triviality results from earlier. Here, we reject the principle of Closure Under Conditionalization. The crucial observation is that this principle is incompatible with our requirement that states be plentiful.

Observation 5. The set of plentiful epistemic spaces is not closed under conditionalization.

To see the problem, let's return to our first example.

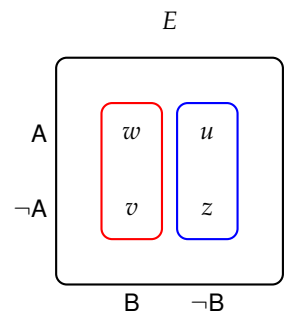

state	world	Pr_E
red	w	1/4
red	v	1/4
blue	u	1/4
blue	z	1/4

E is a proper epistemic space. This space models an agent who has some credence that A obtains, but has no credence that $\Box\mathsf{A}$ obtains, since her credence is concentrated on states that are agnostic with respect to A. Now imagine that our agent conditionalizes on A, eliminating the points $\langle \mathsf{red}, v\rangle$ and $\langle \mathsf{blue}, z\rangle$. She reaches the following space:

9. Strictly speaking, this observation and the last one hold not of epistemic spaces but of epistemic frames, which hold fixed an assignment of probability to state-world pairs but vary the interpretation of atomic sentences. For simplicity in what follows, I omit this complexity by assuming that every set of state-world pairs is the interpretation of some sentence.

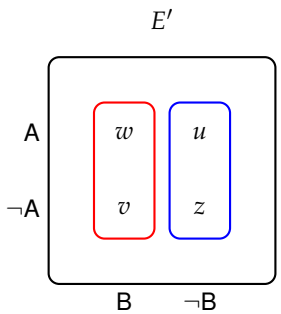

state	world	$Pr_{E'}$
red	w	1/2
blue	u	1/2

This new space E' is not plentiful, because $Pr_{E'}(\mathsf{A}) = 1$, but $Pr_{E'}(\Box\mathsf{A}) = Pr_E(\Box\mathsf{A}) = \frac{1}{2}$. Since A informationally implies $\Box$A by Lukasiewicz' Principle, this shows that E' does not preserve information consequence.

We can't require rational epistemic spaces to be plentiful while allowing that Closure Under Conditionalization is a sufficient condition for rationality. For take some rational agent whose space is plentiful. Given this space, we can find some new piece of evidence A where conditionalizing on A will produce a new epistemic space which is not plentiful. Closure Under Conditionalization then predicts that the new epistemic space is rational, contradicting the requirement to be plentiful.

Our theory rejects the principle that rational credence is closed under conditionalization. Since we reject conditionalization as a rationality-preserving update procedure, we need to supply an alternative. In the next section, we do exactly this. We enrich conditionalization with a further 'normalization' operation that updates the live states in an epistemic space.

6. Updating

We now develop an update rule which rational agents are permitted to use, so that if they are rational prior to updating, they remain rational afterwards. First, we review why the ordinary rule of conditionalization is incompatible with the requirement that informational credence preserves certainty. Then we develop an alternative to ordinary conditionalization: conditionalization with normalization.

Let's begin by considering conditionalization in more detail. On the standard proposal, the credence function of an agent with epistemic space E after learning A would be $Pr_E(\cdot \mid \mathsf{A})$, which assigns any claim B the ratio of $Pr_E(\mathsf{A} \wedge \mathsf{B})$ to $Pr_E(\mathsf{A})$.

Definition 12. $E +_{\mathsf{C}} \mathsf{A} = E +_{\mathsf{C}} \mathbf{A} = \langle W'_E, F'_E, Pr_E(\cdot \mid \mathbf{A})\rangle$

Unfortunately, this theory of updating conflicts with our initial data. An agent can begin in a plentiful space which preserves informational consequence, and then conditionalize on ordinary information to reach a new space which does not preserve informational consequence. To see the problem, consider our running example.

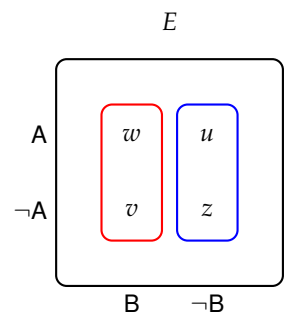

state	world	Pr_E
red	w	1/4
red	v	1/4
blue	u	1/4
blue	z	1/4

This epistemic space is plentiful, since every world in red is assigned some probability when paired with red, and likewise for blue. In this space, the agent has some credence that A obtains, but has no credence that $\Box$A obtains, since her credence is concentrated on states that are agnostic with respect to A. Now imagine that our agent conditionalizes on A. She reaches the following space:

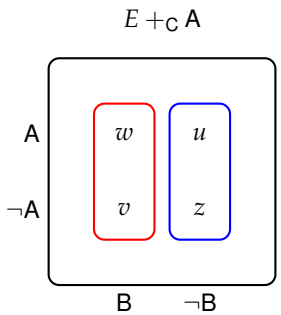

state	world	$Pr_{E+_{\mathsf{C}}\mathsf{A}}$
red	w	1/2
blue	u	1/2

$E +_{\mathsf{C}} \mathsf{A}$ is not plentiful. While red is a live state, the point $\langle \mathsf{red}, v\rangle$ is not live. So this space treats some states as live while excluding some ways that state could be realized. For this reason, $E +_{\mathsf{C}} \mathsf{A}$ does not preserve informational consequence. While $Pr_{E+_{\mathsf{C}}\mathsf{A}}(\mathsf{A}) = 1$, $Pr_{E+_{\mathsf{C}}\mathsf{A}}(\Box\mathsf{A}) = 0$.

To solve this problem, we say that while E is a rational epistemic space, $E +_{\mathsf{C}} \mathsf{A}$ is not. The class of rational epistemic spaces is not closed under conditionalization. Instead, it is closed under a new update rule: conditionalization with normalization. This update rule factorizes updating into parts: first conditionalization, and then normalization. The problem with the example above is that after the agent learns A, her live states stay exactly the same. Before and after she learns A, she assigns credence to states that are agnostic with respect to A. In her posterior state, the agent's credence is no longer plentiful. She treats some states as live which contain worlds that she does not treat as live.

To solve this problem, we model updating with A so that the agent first conditionalizes her state on A, and then normalizes the resulting state. With normalization, we allow the agent to replace defective states with plentiful counterparts. More precisely, let $s + E$ represent the idea of updating or normalizing s with the information in E, by narrowing down the state s to the set $s + \mathsf{Live}(E)$ of worlds w where $\langle s, w\rangle$ is live in E (recall that $s + \mathbf{A}$ is the set of worlds w in s where $\langle s, w\rangle \in \mathbf{A}$). To normalize a space E, we systematically replace every live state s in E with its counterpart $s + E$, by distributing any credence the agent assigned to arbitrary $\langle s, w\rangle$ to its counterpart $\langle s + E, w\rangle$. We then enrich conditionalization with this second step of normalization:

Definition 13. Where $E = \langle W', F', Pr\rangle$:

1. $E +_{\mathsf{CN}} \mathsf{A} = E +_{\mathsf{CN}} \mathbf{A} = \mathsf{Norm}(E +_{\mathsf{C}} \mathbf{A})$, where:
   (a) $\mathsf{Norm}(E) = \langle W', F', \mathsf{Norm}(Pr)\rangle$
   (b) $s + E = s + \mathsf{Live}(E) = \{w \in s \mid \langle s, w\rangle \in \mathsf{Live}(E)\}$
   (c) $\mathsf{Norm}(Pr)(\langle s, w\rangle) = \sum_{s'} Pr(\langle s', w\rangle)$ $\quad s = s' + E$
   (d) $\mathsf{Norm}(Pr)(\mathbf{A}) = \sum_{\langle s,w\rangle} \mathsf{Norm}(Pr)(\langle s, w\rangle)$ $\quad \langle s, w\rangle \in \mathbf{A}$

To see this theory in action, start with our first example:

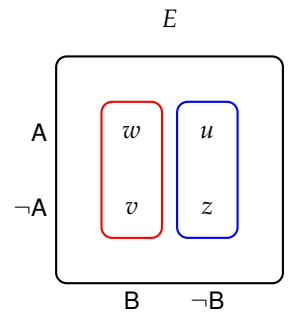

state	world	Pr_E
red	w	1/4
red	v	1/4
blue	u	1/4
blue	z	1/4

We saw above that when we conditionalize E on A, we reach a space that is no longer plentiful. For example, while $\langle \mathsf{red}, v\rangle$ is assigned no probability, the space still treats red as live:

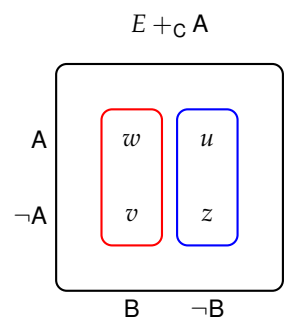

state	world	$Pr_{E+_{\mathsf{C}}\mathsf{A}}$
red	w	1/2
blue	u	1/2

To avoid this problem, we normalize this space $E +_{\mathsf{C}} \mathsf{A}$ to produce $\mathsf{Norm}(E +_{\mathsf{C}} \mathsf{A}) = E +_{\mathsf{CN}} \mathsf{A}$. When we normalize the space, we shrink each of red and blue down to $\{w\}$ and $\{u\}$, because for example w is the only world in s that receives positive probability when paired with s.

$E +_{\mathsf{CN}} \mathsf{A}$

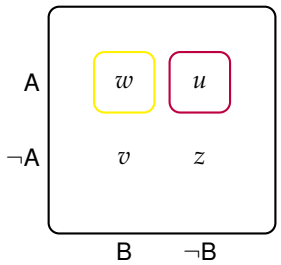

state	world	$Pr_{E+_{\mathsf{CN}}\mathsf{A}}$
yellow	w	1/2
purple	u	1/2

In the resulting state, both A and □A receive a probability of 1.

In our first example, the live states are a partition. Now consider a more complex example:

E

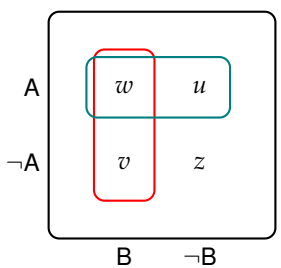

state	world	Pr_E
red	w	1/4
red	v	1/4
teal	w	1/4
teal	u	1/4

When we conditionalize E on $\mathsf{A} \wedge \mathsf{B}$, we produce the following state:

$E +_{\mathsf{C}} \mathsf{A} \wedge \mathsf{B}$

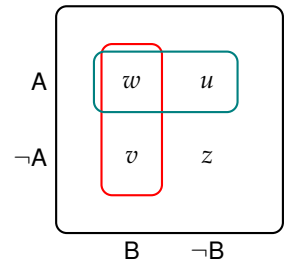

state	world	$Pr_{E+_{\mathsf{C}}\mathsf{A}\wedge\mathsf{B}}$
teal	w	1/2
red	w	1/2

When normalizing, we map both red and teal to $\{w\}$, and move the probability assigned to those previous states over to $\{w\}$.

$E +_{\mathsf{CN}} \mathsf{A} \wedge \mathsf{B}$

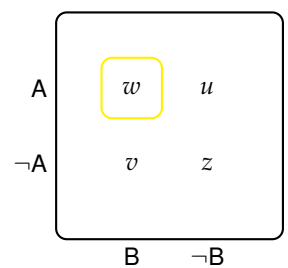

state	world	$Pr_{E+_{\mathsf{CN}}\mathsf{A}\wedge\mathsf{B}}$
yellow	w	1

We've now explored how conditionalization with normalization treats a variety of examples. Let's now turn to some of its structural properties. First, we've seen above that this update procedure respects informational consequence. This is because for any epistemic space E, the result of applying conditionalization with normalization is plentiful.

Observation 6. $E +_{\mathsf{CN}} \mathsf{A}$ is plentiful.

Conditionalization with normalization is a simple recipe for taking any state and producing a plentiful state, which lets informational consequence preserve certainty. On our official theory, the epistemic space of any rational agent is plentiful. We now enrich this theory with

the assumption that the class of rational epistemic spaces is closed under conditionalization with normalization.

(13) **Closure Under Conditionalization With Normalization.** If E is a rational epistemic space, then $E +_{\mathsf{CN}} \mathsf{A}$ is a rational epistemic space.

Observation 6 more than ensures that the two constraints of closure and being plentiful are compatible.

Now let's compare conditionalization with normalization with the ordinary rule of conditionalization. The first thing to see is that the two rules deliver the same verdicts regarding the probability of any non-modal sentence. After all, conditionalization with normalization simply adds a second step of updating to ordinary conditionalization: probabilities must be systematically shifted from certain information states to others. But note that any non-modal sentence has its probability controlled by the world component of any point. So when the probability of $\langle s, w\rangle$ is shifted to $\langle s + E, w\rangle$, the probability of each non-modal sentence stays the same. In this sense, our update rule is a conservative revision of ordinary conditionalization.

There are several structural differences between conditionalization with normalization and ordinary conditionalization. The first difference is that conditionalization with normalization does not automatically preserve certainty. Consider our first example again:

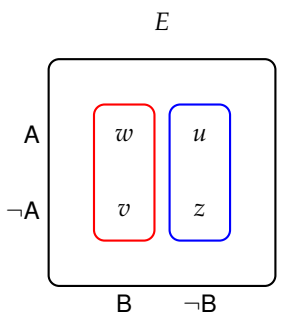

state	world	Pr_E
red	w	1/4
red	v	1/4
blue	u	1/4
blue	z	1/4

Now consider the result of updating with A.

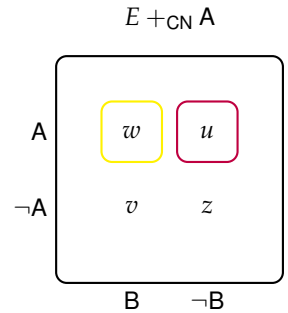

state	world	$Pr_{E+_{\mathsf{CN}}\mathsf{A}}$
yellow	w	1/2
purple	u	1/2

In E, the agent is certain of $\Diamond\neg\mathsf{A}$, since $\Diamond\neg\mathsf{A}$ is true through both of her live states red and blue. In addition, E treats A as possible. But in $E +_{\mathsf{CN}} \mathsf{A}$, $\Diamond\neg\mathsf{A}$ is no longer certain.[10]

Definition 14. $+$ is persistent iff $Pr_{E+_{\mathsf{CN}}\mathsf{A}}(\mathsf{B}) = 1$ whenever $Pr_E(\mathsf{B}) = 1$ and $Pr_E(\mathsf{A}) > 0$.

Observation 7. $+_{\mathsf{CN}}$ is not persistent.

The failure of persistence is a natural consequence of the nonmonotonicity of updating. The live states of the posterior space are not a subset of the live states of the prior space, but rather are the resulting of intersecting each live state from the prior with the worlds where A is true.

In addition, $+_{\mathsf{CN}}$ is not always successful: sometimes, there are claims that are treated as possible, where learning them does not produce a state that is certain of them.

Definition 15. $+$ is successful iff $Pr_{E+\mathsf{A}}(\mathsf{A}) = 1$.

Observation 8. $+_{\mathsf{CN}}$ is not successful.

10. In this way our theory interacts with the impossibility results in Fuhrmann (1989) concerning epistemic possibility.

To see this structural feature, consider the result of learning an epistemic contradiction of the form $\mathsf{A} \wedge \Diamond\neg\mathsf{A}$. Consider the simple example of red and blue above. This epistemic contradiction, is true at $\langle \mathsf{red}, w\rangle$ and $\langle \mathsf{blue}, u\rangle$, since red and blue both contain A and $\neg\mathsf{A}$ worlds. So ordinary conditionalization, $E +_{\mathsf{C}} \mathsf{A} \wedge \Diamond\neg\mathsf{A}$, produces the same state as $E +_{\mathsf{C}} \mathsf{A}$. This is another failure to respect informational consequence, since $Pr_{E+_{\mathsf{C}}\mathsf{A}\wedge\Diamond\neg\mathsf{A}}(\mathsf{A} \wedge \Diamond\neg\mathsf{A}) = 1$ and $Pr_{E+_{\mathsf{C}}\mathsf{A}\wedge\Diamond\neg\mathsf{A}}(\emptyset) = 0$, even though $\mathsf{A} \wedge \Diamond\neg\mathsf{A} \models_i \bot$.

By contrast, when we normalize the result of conditionalizing on $\mathsf{A} \wedge \Diamond\neg\mathsf{A}$, we avoid this prediction. The posterior probability of $\mathsf{A} \wedge \Diamond\neg\mathsf{A}$ is 0. This is because when we normalize $E +_{\mathsf{C}} \mathsf{A} \wedge \Diamond\neg\mathsf{A}$, we do not reach an absurd state. Rather, we simply shrink our two states red and blue to remove worlds where A is false. After updating with $\mathsf{A} \wedge \Diamond\neg\mathsf{A}$, the posterior probability assigned to this claim is 0, even though the resulting space is perfectly coherent. Summing up, in our system epistemic contradictions can have a probability greater than 0 (but never a probability of 1). But updating with an epistemic contradiction is guaranteed to produce a state in which they have a probability of 0. In this way, the system's treatment of epistemic contradictions is analogous to the treatment in Veltman (1996).[11]

7. Avoiding triviality

We've now enriched our theory of modal credence with an account of the norms governing changes in learning. Our official theory is that the epistemic space of any rational agent is proper and plentiful, and that the set of rational epistemic spaces is closed under conditionalization with normalization. This theory predicts that proper consequence preserves probability and informational consequence preserves certainty.

To summarize the ideas in this paper, let's walk through how our theory deals with the triviality results with which we began. First, we can see that each of our initial inferences preserves certainty.

(14) a. **The Direct Argument.** $\neg\mathsf{A} \vee \mathsf{B} \models \mathsf{A} \rightarrow \mathsf{B}$
b. **Lukasiewicz's Principle.** $\mathsf{A} \models \Box\mathsf{A}$
c. **Right Nested Modus Ponens.** $\mathsf{A}, \mathsf{A} \rightarrow (\mathsf{B} \rightarrow \mathsf{C}) \models \mathsf{B} \rightarrow \mathsf{C}$

For illustration, consider the Direct Argument. Suppose an agent is certain that either the gardener or the butler did it. Then at every live world in her epistemic space, one of these suspects committed the crime. But then plenitude ensures that at every world in every live state, either the gardener or the butler did it. So when we restrict our attention to the worlds in a live state where the gardener didn't do it, we know that the butler did it throughout the resulting state. So the agent is certain that the butler did it if the gardener didn't.

Second, none of these inferences preserves probability. Consider an agent where every live state contains worlds with three suspects: butler, gardener, and cook. The agent's credence that the butler did it if the gardener didn't is 0. But she can have a quite high credence that either the butler or the gardener did it, as long as she has some credence that the cook did.

Now let's turn to the resulting threat of triviality. For simplicity, let's continue focusing on the Direct Argument.

(15) a. $Pr(\mathsf{A} \rightarrow \mathsf{B}) = Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) \times Pr(\neg\mathsf{A} \vee \mathsf{B}) + Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg(\neg\mathsf{A} \vee \mathsf{B})) \times Pr(\neg(\neg\mathsf{A} \vee \mathsf{B}))$
b. $Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) = 1$
c. $Pr(\mathsf{A} \rightarrow \mathsf{B}) = Pr(\neg\mathsf{A} \vee \mathsf{B}) + Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg(\neg\mathsf{A} \vee \mathsf{B})) \times Pr(\neg(\neg\mathsf{A} \vee \mathsf{B}))$
d. So $Pr(\mathsf{A} \rightarrow \mathsf{B}) \geq Pr(\neg\mathsf{A} \vee \mathsf{B})$

Our theory accepts the first step of this argument. Since we're assigning probability to sets of state-world pairs, we know that our probability functions satisfy the ordinary rules of probability, even for modal claims.

To respond to this argument, we reject the second premise. We have no guarantee that $Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) = 1$. This is because not every

11. On the other hand, we could also introduce a more dynamic meaning for conjunction, on which $\mathsf{A} \wedge \Diamond\neg\mathsf{A}$ is false at every $\langle s, w\rangle$. In that case, we would predict that every agent's credence in that claim is 0.

epistemic space is rationally permitted. Only plentiful ones are allowed. But we can start with a probability function Pr that is plentiful, and then conditionalize on $\neg\mathsf{A} \vee \mathsf{B}$ to reach a state that is not plentiful. In this new state, informational consequence doesn't preserve certainty, and so $Pr(\mathsf{A} \rightarrow \mathsf{B} \mid \neg\mathsf{A} \vee \mathsf{B}) \neq 1$. On the other hand, we do accept an analogue of the second premise for our preferred rule of updating, conditionalization with normalization. We know that if Pr_E is a rational credence function, then $Pr_{E+_{\mathsf{CN}}\neg\mathsf{A}\vee\mathsf{B}}(\mathsf{A} \rightarrow \mathsf{B}) = 1$. But an analogue of our first premise which replaces conditional probabilities with posterior probabilities after normalization fails.

Summarizing, our theory of modal credence explains the puzzle with which we began. The first key idea is that various constraints on epistemic spaces allow information consequence to preserve certainty while proper consequence preserves probability. The second key idea is that the consistency of these ideas requires us to give up the principle that rationality is preserved by conditionalization on new evidence. Instead, the proper rationality preserving rule of updating is conditionalization with normalization.

8. Transparency

So far, we've offered a series of constraints on credences in epistemic modals. But we have not yet provided an interpretation of what it is for an agent to have such credences. In order to provide such an interpretation, we turn to a new question: can an agent have uncertainty about what information state they occupy? We'll see in the next section that if we do not allow this kind of uncertainty, we can reduce credence in epistemic modal claims to ordinary credence in possible worlds propositions.

Let's call an epistemic space 'transparent' just in case it only treats one information state as live.

Definition 16. E is transparent iff $\exists s \forall s'$: if $s' \in \mathsf{LiveStates}(E)$, then $s = s'$.

When an agent's credence is transparent, she is maximally opinionated about modal matters. Each modal claim receives a probability of 1 or 0.

Observation 9. If E is transparent, then $Pr_E(\Box\mathsf{A}) \in \{0, 1\}$ and $Pr_E(\mathsf{A} \rightarrow \mathsf{B}) \in \{0, 1\}$.

Proof. Suppose E is transparent. Then there is some s which is the unique live state in E. Either $\Box\mathsf{A}$ (alternatively, $\mathsf{A} \rightarrow \mathsf{B}$) is true throughout s, or false throughout s. In the former case, $Pr_E(\Box\mathsf{A}) = 1$; in the latter case, $Pr_E(\Box\mathsf{A}) = 0$. □

Santorio (2018a) develops a model of modal credence on which it is transparent. Much of this paper's task has been to extend that model to a more general setting in which transparency is not assumed. In this section, we'll critically explore the prospects for transparency. We'll see that there are a variety of reasons to accept transparency, although none are conclusive.

Let's begin with an argument for modal opinionation. Suppose that I am about to flip a fair coin. Now consider the following conditional:

(16) If I flip the coin, it will land heads.

When considering your credence in (16), there are two natural reactions. One answer is 1/2, the conditional probability of heads given that the coin is flipped (see for example Stalnaker (1970) and Adams (1998)). But another answer is that the probability of (16) is 0. After all, in this scenario we assign a credence of 1 to the following *might* conditional:

(17) If I flip the coin, it might land tails.

But it is natural to think that (16) and (17) are incompatible. Indeed, the theory of modal credence defended above has the consequence that $\mathsf{A} \rightarrow \Diamond\mathsf{B}$ and $\mathsf{A} \rightarrow \neg\mathsf{B}$ are probabilistically incompatible, since exactly one of these claims will be true at each state.[12]

12. But see Santorio (2018b), Santorio (Forthcoming) for an alternative theory of conditionals on which these claims are not probabilistically incompatible, although they are informationally inconsistent.

Let's think more systematically now about whether our credence in conditionals is always opinionated. To reach that result, let's continue to assume that the Direct Argument preserves certainty, so that:

(18) If $Pr(\mathsf{B} \mid \mathsf{A}) = 1$, then $Pr(\mathsf{A} \to \mathsf{B}) = 1$.

Now, building on the observation above, let's add the further premise that whenever an agent is not conditionally certain of B given A, she rejects the conditional $\mathsf{A} \to \mathsf{B}$.

(19) If $Pr(\mathsf{B} \mid \mathsf{A}) < 1$, then $Pr(\mathsf{A} \to \mathsf{B}) = 0$.

Again, our argument for this relies on two more general principles about *might* conditionals:

(20) a. If $Pr(\mathsf{B} \mid \mathsf{A}) > 0$, then $Pr(\mathsf{A} \to \Diamond\mathsf{B}) = 1$
b. $Pr(\mathsf{A} \to \Diamond\neg\mathsf{B}) + Pr(\mathsf{A} \to \mathsf{B}) = 1$

These two principles entail (19). For suppose $Pr(\mathsf{B} \mid \mathsf{A})$ is less than 1. Then $Pr(\neg\mathsf{B} \mid \mathsf{A})$ is positive, and so the agent is certain of $\mathsf{A} \to \Diamond\neg\mathsf{B}$, and thereby rejects $\mathsf{A} \to \mathsf{B}$ entirely.

When we put together our two constraints on credence in conditionals, we reach an opinionated theory. Our credence in a conditional $\mathsf{A} \to \mathsf{B}$ is always 1 or 0, depending on whether or not our conditional credence in B given A is 1.

Now let's turn from conditionals to modals. Again, we hold fixed that Lukasiewicz's Principle preserves certainty:

(21) If $Pr(\mathsf{A}) = 1$, then $Pr(\Box\mathsf{A}) = 1$.

But now consider a further principle governing epistemic possibility: whenever the agent has some credence in A, she is certain of $\Diamond\mathsf{A}$:

(22) If $Pr(\mathsf{A}) > 0$, then $Pr(\Diamond\mathsf{A}) = 1$.

Principles of this kind have been studied for some time in connection with impossibility results in Fuhrmann (1989):

> Your epistemic state commits you to *It might be that p* iff it does not commit you to $\neg p$. (Gillies (2006))

Finally, we assume as is standard that $\Diamond\mathsf{A}$ and $\Box\mathsf{A}$ are probabilistically incompatible, so that any agent who is certain in $\Diamond\neg\mathsf{A}$ will completely reject $\Box\mathsf{A}$. This gives us the result that whenever $Pr(\mathsf{A})$ is less than 1, $Pr(\Box\mathsf{A})$ is 0.

The principles we've explored above imply that agents are fully opinionated on modal matters. It turns out that on the theory above, any proper, plentiful, and transparent theory validates just these bridge principles:

Observation 10. If E is proper, plentiful, and transparent, then:

1. $Pr(\mathsf{A} \to \mathsf{B}) = \begin{cases} 1 & \text{if } Pr(\mathsf{B} \mid \mathsf{A}) = 1 \\ 0 & \text{otherwise} \end{cases}$

2. $Pr(\Box\mathsf{A}) = \begin{cases} 1 & \text{if } Pr(\mathsf{A}) = 1 \\ 0 & \text{otherwise} \end{cases}$

3. $Pr(\Diamond\mathsf{A}) = \begin{cases} 1 & \text{if } Pr(\mathsf{A}) > 0 \\ 0 & \text{otherwise} \end{cases}$

The principles above imply that an agent's credence in modal claims supervenes on their credence in non-modal claims. Once we know an agent's credence in A, we know her credence in $\Box\mathsf{A}$ and $\Diamond\mathsf{A}$. Once we know her credence in B given A, we know her credence in $\mathsf{A} \to \mathsf{B}$.[13]

13. On the other hand, transparent theories do not imply that every sentence containing an epistemic modal has a probability of 1 or 0. For example, when $\Box\mathsf{A}$ has a probability of 1, the probability of the conjunction $\Box\mathsf{A} \wedge \mathsf{B}$ is identical to the probability of B, which may lie strictly between 0 and 1.

9. Reduction

Transparent epistemic spaces allow us to validate substantive bridge principles reducing modal credence to non-modal credence. A natural question at this point is whether there is anything systematic to say about how transparent epistemic spaces come about. What information state in particular should an agent be certain of? Here, we can draw on previous work in Yalcin (2012), Beddor and Goldstein (2018), and Goldstein (2019) to offer an answer. We now introduce a construction $\uparrow$ which maps any ordinary probability function over worlds into a corresponding transparent epistemic space, which assigns probability to state-world pairs.

In particular, let's introduce the notion of an agent's proto epistemic space, consisting of a set of possible worlds and a probability function defined over that set. We might understand this probability function, for example, in terms of an agent's betting dispositions on non-modal claims.

Definition 17. A proto epistemic space $U = \langle W', Cr \rangle$, where:

1. $W' \subseteq W$ is a set of worlds.
2. (a) Cr is a probability measure over W'.
   (b) $Cr(w)$ abbreviates $Cr(\{w\})$.
3. $\mathsf{LiveWorlds}(U) = \{w \mid Cr_U(w) > 0\}$ is the set of live worlds in U.

Now we can introduce an operation $\uparrow$ which maps each proto epistemic space to a unique epistemic space. Our construction starts from the observation that there is a special information state consistent with the live worlds of U: $\mathsf{LiveWorlds}(U)$ itself, the set of worlds assigned positive probability by U. Then the probability Pr_E of any claim $\mathbf{A}$ is the probability that Cr_U assigns to the result of updating this minimal information state $\mathsf{LiveWorlds}(U)$ with $\mathbf{A}$. (Recall that $s + \mathbf{A}$, the result of updating state s with meaning $\mathbf{A}$, is the set of worlds w in s where $\langle s, w \rangle \in \mathbf{A}$.)

Definition 18. $E = \uparrow U$ iff:

1. $W_E = W_U$
2. $F_E = W_U \times \mathcal{P}(W_U)$
3. $Pr_E(\mathbf{A}) = Cr_U(\mathsf{LiveWorlds}(U) + \mathbf{A})$

In other words, the construction above guarantees that $Pr_E(\mathbf{A}) = Cr_U(\{w \mid \langle \mathsf{LiveWorlds}(U), w \rangle \in \mathbf{A}\})$.

Here's an example: consider the proto epistemic space $U = \langle W, Cr \rangle$, where W is our earlier set $\{w, v, u, z\}$ and Cr is indifferent over W.

world	Cr_U
w	1/4
v	1/4
u	1/4
z	1/4

When we lift U with $\uparrow$, we reach a transparent space that assigns all probability to the single weakest state consistent with the live worlds.

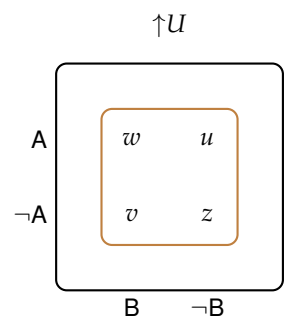

state	world	$Pr_{\uparrow U}$
brown	w	1/4
brown	v	1/4
brown	u	1/4
brown	z	1/4

Now that we've illustrated our construction with an example, let's turn to its structural properties. Let's begin with our main topic, the connection between credence and consequence. The first observation is that our theory guarantees that any constructed epistemic state is proper and plentiful. This means that our construction is sufficient to solve our initial puzzle, so that proper consequence preserves probability and

informational consequence preserves certainty. In addition, we know that any constructed epistemic space is transparent: $\mathsf{LiveWorlds}(U)$ is the unique state assigned positive probability.

Observation 11. If $E = \uparrow U$, then E is proper, plentiful, and transparent.

Proof. Suppose that $E = \uparrow U$. Then there is one live state in E: the set of live worlds in U. To show that E is proper, take some arbitrary $\langle s, w\rangle$, live in E. Since $\langle s, w\rangle$ is live in E, we know that w is a live world in U. This guarantees that $w \in s$. To show that E is plentiful, again consider the unique live s in E. For every world w in s, we know that $\langle s, w\rangle$ is live in E, since w is live in U. □

To fully explain the triviality results with which we began, we introduced a new update rule: conditionalization with normalization. To see our update rule at work in a transparent setting, consider the following example, where the agent is certain of red:

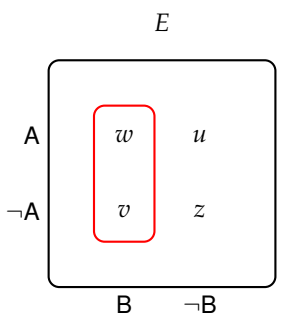

state	world	Pr_E
red	w	1/2
red	v	1/2

Now imagine that our agent updates on A. To do so, she first rules out the point $\langle \mathsf{red}, v\rangle$, assigning $\langle \mathsf{red}, w\rangle$ a probability of 1. But now her state is not plentiful. To resolve this tension, the agent must shift her credence entirely from $\langle \mathsf{red}, w\rangle$ to $\langle \{w\}, w\rangle$.

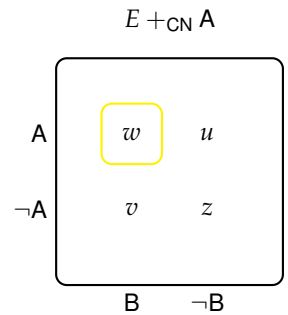

state	world	$Pr_{E+_{\mathsf{CN}}\mathsf{A}\wedge\mathsf{B}}$
yellow	w	1

At every time, a transparent agent treats a single state as live. But which state is live depends on what she has learned. When she learns new information, she flips from fully accepting one state to fully accepting another.

We can understand the same doxastic evolution again from the perspective of our construction procedure. Now imagine that our agent begins with the following proto epistemic space:

world	Cr_U
w	1/2
v	1/2

When we apply $\uparrow$ to this proto epistemic space, we reach exactly the epistemic space above, with an agent who is certain of the state red, but equally unsure about w and v.

Likewise, consider the second epistemic space from above, $E +_{\mathsf{CN}} \mathsf{A} \wedge \mathsf{B}$. We can think of this epistemic space as derived from $\uparrow$ by yet another proto epistemic space:

world	$Cr_{U'}$
w	1

U and U' are related in a familiar way: we reach U' by conditionalizing

Cr_U on w, the unique world at which $\mathsf{A} \wedge \mathsf{B}$ are true.

The point generalizes: suppose an epistemic space E is constructed from a proto epistemic space U. Suppose we then update the epistemic space E with some claim A using Conditionalization with Normalization. The resulting epistemic space is just the one we reach by simply conditionalizing the probability function in U on $\mathsf{LiveWorlds}(U) + \mathsf{A}$. (Again, $\mathsf{LiveWorlds}(U) + \mathsf{A}$ is the set of live worlds in U where A is true relative to the live worlds.)

Observation 12. Suppose $E = \uparrow U$. Then $E +_{\mathsf{CN}} \mathsf{A} = \uparrow\langle W_U, Cr_U(\cdot \mid \mathsf{LiveWorlds}(U) + \mathsf{A})\rangle$.

Proof. Let $E^* = \uparrow\langle W_U, Cr_U(\cdot \mid \mathsf{LiveWorlds}(U) + \mathsf{A})\rangle$. We show that $Pr_{E+_{\mathsf{CN}}A}(\langle s, w\rangle) = Pr_{E*}(\langle s, w\rangle)$ for arbitrary $\langle s, w\rangle$. Either $s = \mathsf{LiveWorlds}(U) + \mathsf{A}$ or not. If not, $Pr_{E*}(\langle s, w\rangle) = 0 = Pr_{E+_{\mathsf{CN}}A}(\langle s, w\rangle)$. Otherwise, we have $Pr_{E+_{\mathsf{CN}}A}(\langle \mathsf{LiveWorlds}(U) + \mathsf{A}, w\rangle) = \mathsf{Norm}(Pr_E(\langle \mathsf{LiveWorlds}(U) + \mathsf{A}, w\rangle \mid \mathsf{A}) = Pr_E(\langle \mathsf{LiveWorlds}(U), w\rangle \mid \mathsf{A}) = Cr_U(w \mid \mathsf{LiveWorlds}(U) + \mathsf{A}) = Pr_{E*}(\langle \mathsf{LiveWorlds}(U) + \mathsf{A}, w\rangle)$. $\square$

This last result provides a powerful vindication of conditionalization with normalization. Ultimately, we can think of this update rule as emerging naturally out of more basic ingredients. Fundamentally, an agent is endowed with a credence function over worlds. This credence function changes over time by the familiar process of ordinary conditionalization. But at any time, an agent's doxastic perspective extends beyond her credence in worlds. Her credence in worlds determines her broader outlook on modal matters via our construction procedure. From the modal perspective, the agent doesn't count as updating via the rule of conditionalization alone. But from the worldly perspective, she does.

10. Transparency without opinionation

In the previous sections, we've explored the prospects for a transparent theory of modal credence. The cost of such a theory is readily apparent: it requires agents to be maximally opinionated about modal matters. But the benefits of the theory are also powerful. First, the theory validates a variety of plausible bridge principles connecting modal and non-modal credence. Second, the theory allows us to reduce modal credence to non-modal credence, via our construction procedure $\uparrow$. Finally, this construction procedure vindicates our update rule of conditionalization with normalization: this rule is just what one would expect if non-modal credence is updated via ordinary conditionalization.

A transparent theory of modal credence offers a simple philosophical reduction of credence in state-world pairs: this state supervenes on our ordinary credence over worlds, itself potentially reducible to behavioral dispositions of some kind of other. If we reject transparency, further work remains to provide a philosophical interpretation of credence in sets of pairs of information states and possible worlds.

If we accept a transparent theory of modal credence, we need to explain why agents do not always appear maximally opinionated about modal matters. One option is to retain the theory of modal credence above, but offer a new semantics for epistemic modals. Following Goldstein (2020), we can let the meaning of an epistemic modal claim be sensitive to both the world and information state parameter, so that modal claims can be true at some worlds and false at others. Then even an agent who is certain of the information state could still be agnostic about modal matters, as long as some modal claim is true at some live worlds and false at others. Such a theory would invalidate the specific bridge principles between modal and non-modal credence in §8, which encode opinionation. But the theory could still embrace the construction procedure above, so that modal credence is still reducible to and supervenient on non-modal credence. To solve our initial puzzles, any such theory would still need to guarantee that inferences like the Direct Argument are still informationally valid (see Goldstein (2020) for characterizations of the informational validity of these inferences within a world and information sensitive semantics).

For example, one option explored in Goldstein (2020) is that epistemic modalities quantify over worlds determined by a classical accessibility relation that is appropriately restricted to the current information

state. So let Rw represent the epistemic possibilities at w according to relation R. Rw might for example be the worlds consistent with what is known by a contextually salient group. Then:

Definition 19.

1. $[\![\Diamond \mathsf{A}]\!]^{s,w} = 1$ iff $\exists v \in Rw \cap s : [\![\mathsf{A}]\!]^{Rw \cap s, v} = 1$
2. $[\![\Box \mathsf{A}]\!]^{s,w} = 1$ iff $[\![\mathsf{A}]\!]^{Rw \cap s} = 1$
3. $[\![\mathsf{A} \rightarrow \mathsf{B}]\!]^{s,w} = 1$ iff $[\![\mathsf{B}]\!]^{(Rw \cap s) + \mathsf{A}} = 1$

To deal with our initial data, this semantics can be combined with an epistemology on which rational agents are proper, plentiful, and transparent. On the resulting theory, the agent always assigns all of her credence to a single information state, and the agent's epistemic modal credence supervenes on her credences in ordinary claims. Yet the agent may not be maximally opinionated about epistemic modal claims, because a claim can be true at an information state and one world, while false at that information state when paired with a different world. This strategy for dealing with epistemic modals differs from Moss (2015) and Charlow (2020), who instead model epistemic modal ignorance by assigning epistemic modal claims semantic values at a higher type than sets of information state-world pairs.[14]

11. Conclusion

We close with one open question for our theory: the probability of conditionals. The framework above successfully predicts that when the conditional probability of B given A is 1, the probability of the conditional is 1. This follows from the Direct Argument preserving certainty. However, the theory above does not validate Stalnaker's Thesis when the probability of the conditional is non-extreme. That is, we do not have the general result that $Pr(\mathsf{A} \rightarrow \mathsf{B}) = Pr(\mathsf{B} \mid \mathsf{A})$.

14. Finally, see Beddor and Goldstein (2021) for an approach to epistemic modals on which epistemic modal belief may be opinionated, but epistemic modal knowledge is not.

This is not an artifact of the system above.[15] In the semantics above, the conditional is strict. Relative to all our notions of consequence, we have that $\mathsf{A} \rightarrow \mathsf{C}$ implies $(\mathsf{A} \wedge \mathsf{B}) \rightarrow \mathsf{C}$. But no such conditional can satisfy Stalnaker's Thesis while allowing classical or proper consequence to preserve probability. The problem is that conditional probability itself is non-monotonic. We can have $Pr(\mathsf{C} \mid \mathsf{A}) > Pr(\mathsf{C} \mid (\mathsf{A} \wedge \mathsf{B}))$, even though the preservation of classical consequence for the strict conditional implies that $Pr(\mathsf{A} \rightarrow \mathsf{C}) \leq Pr((\mathsf{A} \wedge \mathsf{B}) \rightarrow \mathsf{C})$.

The upshot is that if we want to extend our theory to validate Stalnaker's Thesis, we need to give up the strict semantics for the conditional. In this way, the original triviality results in Lewis (1976) are more powerful than the triviality results we have focused on in this paper. A sequel to this paper, Goldstein and Santorio (Forthcoming) pursues exactly this strategy. The sequel develops an enrichment of information sensitive semantics on which meanings are sets of paths, and paths are linear orderings of possible worlds (as in Santorio (2018b)). All of the ideas from this paper can be transposed to that richer setting in a way that preserves the results above while validating Stalnaker's Thesis with a variably strict conditional semantics.[16]

15. Thanks to Paolo Santorio for help here.
16. Thanks to Paolo Santorio, and the fourth Taiwan Metaphysics Colloquium.

Appendix: propriety

This appendix explores another way to motivate propriety as a constraint on epistemic spaces. In particular, we will modify our semantics so that the meaning of every sentence in L is itself proper, in that $\langle s, w\rangle \in [\![\mathsf{A}]\!]$ only if $w \in s$. This proposal is not arbitrary. Consider the update semantics in Veltman (1996), where $[\cdot]$ is a function from information states to information states, so that $s[\mathsf{A}]$ represents the result of updating state s with sentence A. To update s with an atomic sentence like *it's raining*, we remove any world from s where it isn't raining. To update s with $\neg\mathsf{A}$, we remove from s any world that survives update with A. To update s with $\mathsf{A} \wedge \mathsf{B}$ and $\mathsf{A} \vee \mathsf{B}$, we explore the result of updating with A and updating with B, either intersecting or unioning the two states. Then the meaning of $\Diamond\mathsf{A}$, $\Box\mathsf{A}$, and $\mathsf{A} \rightarrow \mathsf{B}$ are tests, exploring whether the initial information state is consistent with A, implies A, or implies B after updating with A. In each case, failure of the test produces the absurd state, while passing the test produces the initial state.

Definition 20.

1. $s[\mathsf{p}] = \{w \in s \mid w(\mathsf{p}) = 1\}$
2. $s[\neg\mathsf{A}] = \{w \in s \mid w \notin s[\mathsf{A}]\}$
3. $s[\mathsf{A} \wedge \mathsf{B}] = s[\mathsf{A}] \cap s[\mathsf{B}]$
4. $s[\mathsf{A} \vee \mathsf{B}] = s[\mathsf{A}] \cup s[\mathsf{B}]$
5. $s[\Diamond\mathsf{A}] = \{w \in s \mid s[\mathsf{A}] \neq \emptyset\}$
6. $s[\Box\mathsf{A}] = \{w \in s \mid s[\mathsf{A}] = s\}$
7. $s[\mathsf{A} \rightarrow \mathsf{B}] = \{w \in s \mid s[\mathsf{A}][\mathsf{B}] = s[\mathsf{A}]\}$

At first glance, the semantics above seems to encode similar ideas about epistemic modals to the information sensitive semantics earlier, but in a different framework. As a matter of fact, the two systems are closer than they first appear. Given any function $[\mathsf{A}]$ from states to states $[\mathsf{A}]$, we can construct an 'uncurried' set of state-world pairs $[\![\mathsf{A}]\!]$, where $\langle s, w\rangle$ is in $[\![\mathsf{A}]\!]$ if and only if w survives an update of s with $[\mathsf{A}]$.[17]

Definition 21. Where $\mathbf{A}$ is a function of type $\langle\langle s, t\rangle, \langle s, t\rangle\rangle$, let $\circlearrowright\mathbf{A} = \{\langle s, w\rangle \mid w \in \mathbf{A}(s)\}$.

When we systematically uncurry Definition 20, we reach a variant of the information sensitive semantics from earlier:

Definition 22.

1. $[\![\mathsf{p}]\!]^{s,w} = 1$ iff $w \in s$ & $w(\mathsf{p}) = 1$
2. $[\![\neg\mathsf{A}]\!]^{s,w} = 1$ iff $w \in s$ & $[\![\mathsf{A}]\!]^{s,w} = 0$
3. $[\![\mathsf{A} \wedge \mathsf{B}]\!]^{s,w} = 1$ iff $w \in s$ & $[\![\mathsf{A}]\!]^{s,w} = 1$ & $[\![\mathsf{B}]\!]^{s,w} = 1$
4. $[\![\mathsf{A} \vee \mathsf{B}]\!]^{s,w} = 1$ iff $w \in s$ & $([\![\mathsf{A}]\!]^{s,w} = 1$ or $[\![\mathsf{B}]\!]^{s,w} = 1$
5. $[\![\Diamond\mathsf{A}]\!]^{s,w} = 1$ iff $w \in s$ & $\exists v \in s : [\![\mathsf{A}]\!]^{s,v} = 1$
6. $[\![\Box\mathsf{A}]\!]^{s,w} = 1$ iff $w \in s$ & $[\![\mathsf{A}]\!]^{s} = 1$
7. $[\![\mathsf{A} \rightarrow \mathsf{B}]\!]^{s,w} = 1$ iff $w \in s$ & $[\![\mathsf{B}]\!]^{s+\mathsf{A}} = 1$

Observation 13. $\forall \mathsf{A} \in \mathsf{L} : \circlearrowright[\mathsf{A}] = [\![\mathsf{A}]\!]$

This semantics is just like the earlier one, except that it systematically strengthens the meaning of each sentence to include a requirement that w is in s. Imagine enriching the object language with a single proposition Proper, true at $\langle s, w\rangle$ iff $w \in s$. Then any sentence A in this semantics has the same meaning as the sentence $\mathsf{A} \wedge \mathsf{Proper}$ in our earlier semantics.

Observation 14. If $\langle s, w\rangle \in [\![\mathsf{A}]\!]$, then $w \in s$.

To see the difference between this semantics and our earlier one, consider again the inference from $\Box\mathsf{A}$ to A. We saw that in our earlier semantics, this inference was classically invalid, because $\Box\mathsf{A}$ could be true at an improper point $\langle\{w\}, v\rangle$ with A true at w and false at v. Our new semantics makes $\Box\mathsf{A}$ false at such improper points, and so predicts that $\Box\mathsf{A}$ classically implies A. More generally, an inference is classically and properly valid on our new semantics if and only if it was properly

17. See Schroeder (2015) and Rothschild (2017) for similar thoughts.

valid on our old semantics.

This new semantics provides a stronger motivation for imposing propriety as a constraint on epistemic spaces. For suppose we do not. Then $Pr(\mathsf{A} \vee \neg\mathsf{A})$ may be less than 1. After all, $\mathsf{A} \vee \neg\mathsf{A}$ is only true at proper points; but if E is improper than Pr_E will assign positive probability to some improper points, where $\mathsf{A} \vee \neg\mathsf{A}$ is false.

References

Ernest Adams. *A Primer of Probability Logic*. CSLI Publications, 1998.

David Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, 2001.

Bob Beddor and Simon Goldstein. Believing epistemic contradictions. *Review of Symbolic Logic*, 11(1), March 2018.

Bob Beddor and Simon Goldstein. Mighty knowledge. *Journal of Philosophy*, 2021.

Jonathan Bennett. *A Philosophical Guide to Conditionals*. Clarendon Press, 2003.

Fabrizio Cariani. On stalnaker's 'indicative conditionals'. In Louise McNally, Yael Sharvit, and Zoltan Szabo, editors, *Studies in Linguistics and Philosophy*, volume 100. Forthcoming.

Nate Charlow. Grading modal judgment. *Mind*, 129:769–807, 2020.

Dorothy Edgington. On conditionals. *Mind*, 104(414):235–329, 1995.

Hartry Field. What is logical validity? In Colin R. Caret and Ole T. Hjortland, editors, *Foundations of Logical Consequence*. Oxford University Press, 2015.

André Fuhrmann. Reflective modalities and theory change. *Synthese*, 81:115–134, 1989.

Anthony S. Gillies. Epistemic conditionals and conditional epistemics. *Noûs*, 38(4):585–616, 2004.

Anthony S. Gillies. What might be the case after a change in view. *Journal of Philosophical Logic*, 35(2):117–145, 2006.

Simon Goldstein. Triviality results for probabilistic modals. *Philosophy and Phenomenological Research*, 99:188–222, 2019.

Simon Goldstein. Free choice impossibility results. *Journal of Philosophical Logic*, 49:249–282, 2020.

Simon Goldstein and Paolo Santorio. Probability for epistemic modalities. *Philosophers' Imprint*, Forthcoming.

Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and modality. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179–216. Blackwell, 1996.

Alan Hájek. Staying regular? `https://www.youtube.com/watch?v=2qGAjxImLs0&ab_channel=DiaphoraEtn`, 2017.

Frank Jackson. *Conditionals*. Blackwell, 1987.

Nathan Klinedinst and Daniel Rothschild. Connectives without truth tables. *Natural Language Semantics*, 20(2):137–175, June 2012.

David Lewis. Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85:297–315, 1976.

Matt Mandelkern. Bounded modality. *Philosophical Review*, 128:1–61, 2019.

Sarah Moss. On the semantics and pragmatics of epistemic modals. *Semantics and Pragmatics*, 8:1–81, 2015.

Dilip Ninan. Relational semantics and domain semantics for epistemic modals. *Journal of Philosophical Logic*, 47:1–16, 2018.

Daniel Rothschild. Veltman-Yalcin. http://danielrothschild.com/dyncon/vy/, 2017.

Paolo Santorio. Credence for epistemic discourse. https://philpapers.org/archive/SANCFE.pdf, 2018a.

Paolo Santorio. Conditional excluded middle in informational semantics. In *Proceedings of the Amsterdam Colloquium*, volume 21, 2018b.

Paolo Santorio. Trivializing informational consequence. *Philosophy and Phenomenological Research*, 2021.

Paolo Santorio. Path semantics for indicative conditionals. *Mind*, Forthcoming.

Mark Schroeder. Is semantics formal? In *Expressing Our Attitudes*. Oxford University Press, 2015.

Moritz Schulz. Epistemic modals and informational consequence. *Synthese*, 174(3):385–395, 2010.

Moritz Schulz. *Counterfactuals and probability*. Oxford University Press,

2017.

Robert Stalnaker. Probability and conditionals. *Philosophy of Science*, 37 (1):64–80, March 1970.

Robert Stalnaker. Indicative conditionals. *Philosophia*, 5:269–86, 1975.

Eric Swanson. How not to theorize about the language of subjective uncertainty. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2011.

Eric Swanson. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, pages 1–28, 2012.

Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.

Kai von Fintel and Anthony Gillies. An opinionated guide to epistemic modality. In Tamar Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology*, volume 2, pages 32–62. Oxford University Press, Oxford, 2007.

Kai von Fintel and Anthony S. Gillies. *Must* ... stay ... strong! *Natural Language Semantics*, 18(4):351–383, 2010.

Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122 (1):45–92, 2013.

Malte Willer. Simplifying counterfactuals. In *Proceedings of the 20th Amsterdam Colloquium*, pages 428–437, 2015.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.

Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012.

Seth Yalcin. Epistemic modality *de re*. *Ergo*, 2(19), 2015.