Probabilities of conditionals: updating Adams*

Ivano Ciardelli and Adrian Ommundsen
Munich Center for Mathematical Philosophy
LMU Munich

Preprint. Final version published in *Noûs* (link)

Abstract

The problem of probabilities of conditionals is one of the long-standing puzzles in philosophy of language. We defend and update Adams' solution to the puzzle: the probability of an epistemic conditional is not the probability of a proposition, but a probability under a supposition.

Close inspection of how a triviality result unfolds in a concrete scenario does not provide counterexamples to the view that probabilities of conditionals are conditional probabilities: instead, it supports the conclusion that probabilities of conditionals violate standard probability theory.

This does not call into question probability theory per se; rather, it calls for a more careful understanding of its role: probability theory is a theory of probabilities of propositions; but as conditionals do not express propositions, their probabilities are not subject to the standard laws.

We argue that both conditional probabilities and probabilities of conditionals are best understood in terms of the dynamics of supposing, modeled as a restriction operation on a probability space. This version of the suppositionalist view allows us to connect Adams' Thesis to the widely held restrictor view of the semantics of conditionals.

We address two common objections to Adams' view: that the relevant probabilities are 'probabilities only in name', and that giving up conditional propositions puts us at a disadvantage when it comes to interpreting compounds.

Finally, we argue that some putative counterexamples to Adams' Thesis can be diagnosed as fallacies of probabilistic reasoning: they arise from applying to conditionals laws of standard probability theory which are invalid for them.

*We are grateful to Matt Mandelkern and Paolo Santorio for valuable discussion of the present ideas. This project was supported by the German Research Foundation (DFG)—Project number 446711878. Gefördert durch die Deutsche Forschungsgemeinschaft (DFG)—Projektnummer 446711878.

1 Introduction

Indicative conditionals and probabilities are two key items in our conceptual and linguistic toolbox to deal with uncertainty. Yet, after forty years of intensive study, the problem of how to best construe the probabilities of conditionals remains open. The crux of the problem, as brought out by the triviality results of Lewis (1976) and many since,[1] is a tension existing between, on the one hand, standard assumptions about probabilities, and on the other hand, an important desideratum, namely, Adams' Thesis that the probability of an indicative conditional is a conditional probability: the probability of the consequent on the supposition of the antecedent.[2]

In response to this, the attitude that initially prevailed was the one of Lewis, who gave up on the desideratum and concluded that probabilities of conditionals simply cannot match conditional probabilities. Fifty years later, however, this response looks more and more on the wrong track, for two sorts of reasons.

First, there is much to be said in favor of Adams' Thesis. There is empirical evidence: to use the words of McGee (1989), the thesis predicts "with uncanny accuracy" the way conditionals are assessed in situations of uncertainty by competent speakers.[3,4] Perhaps speakers are systematically mistaken in their judgments about conditionals; but the relevant intuitions are very basic, consistent, and robust under reflection; and the error theories that have been put forward to explain these judgments have been found unconvincing.[5] There are also conceptual reasons why the thesis should hold. The most solid intuition we have about the interpretation of conditionals is arguably captured by the Ramsey test idea: when assessing a conditional "if A then B", one supposes A, and then assesses B in the resulting hypothetical state. In particular, when assessing the conditional for probability, one supposes A and then assesses the probability of B in the resulting hypothetical state. In those cases in which the relevant kind of supposition is well-modeled by conditionalization, the result of this process is just the conditional probability of the consequent given the antecedent. So, the fact that our intuitions conform to Adams' Thesis is not a surprising finding; rather, it is an empirical validation of an independently plausible hypothesis. This suggests that theories that invalidate Adams' Thesis are missing something fundamental about the way we interpret and assess conditionals.

Second, it turns out that one can get to triviality results without appealing to Adams' Thesis. One can replicate Lewis's proof using much weaker assumptions that concern only *certainty* about conditionals (Bradley, 2000), and which seem

[1]See Khoo and Santorio (2018) for an overview of the literature.

[2]A note on this terminological choice. We call this Adams' Thesis, since it is the central tenet of Adams' theory ("The fundamental assumption of this work is: the probability of an indicative conditional is a conditional probability" Adams (1975), p. 3). Many authors refer to it as Stalnaker's Thesis (after Stalnaker, 1970), though Stalnaker only held the view briefly. Many of these authors assume that 'Adams' Thesis' is a different claim, equating some other quantity about conditionals with conditional probability. This is motivated by the view that what Adams calls probabilities of conditionals are not *real* probabilities, but something else. We will argue against this view in Section 5, and thus we disagree that two different theses are at stake. (Adams himself talked about 'assertibility' in early work (Adams, 1965), but later on he consistently and deliberately used 'probability'.)

[3]For empirical studies, see among many others, Over and Evans (2003); Evans *et al.* (2003). A more comprehensive list of references can be found in Douven and Verbrugge (2013), p. 712.

[4]At least, it seems to describe the rule. There are exceptions, that we will come back to in Section 6.

[5]See Edgington (1995); DeRose (2010); Khoo and Santorio (2018) for discussion of the error theories proposed by Lewis (1976) and Jackson (1979); see von Fintel and Gillies (2015); Khoo and Santorio (2018); Mandelkern (2018); Ciardelli (2021) for discussion of an error theory based on Kratzer's restrictor view.

overwhelmingly plausible.[6] One can also obtain triviality results from plausible assumptions which do not concern conditionals at all, but which constrain the probabilities of sentences involving, for instance, the epistemic modals *might* and *probably* (Russell and Hawthorne, 2016; Goldstein, 2019b). These findings suggest that the problem discovered by Lewis is a deep and general one. Adams' Thesis is just one of many different and independently plausible desiderata, each of which is at odds with the standard set of assumptions about probabilities. Giving up each of these desiderata whenever one encounters a triviality result seems wrong-headed: there seems to be an important lesson to be learned here about the limitations of the standard theoretical framework. Our aim in this paper is to contribute to understanding what the lesson is.

We will start by arguing, in Section 2, that if we observe a triviality result unfold in a concrete scenario where we have sharp intuitions about probabilities, what we find is not a counterexample to Adams' Thesis, but rather a counterexample to standard probability theory: probabilities of conditionals can increase dramatically upon conditionalization, in a way that violates the standard laws of probability. In Section 3 we will then ask *why* probabilities of conditionals have this special behavior. We will argue that probability theory is best understood as providing a model of probabilities of *propositions*; substantial assumptions about how language works are then involved in going from probabilities of propositions to probabilities of sentences in context, which is what our judgments are about. Following Adams (1975), we argue that certain conditionals, which we call *epistemic*, do not express propositions, and that their probabilities are, therefore, not probabilities of propositions; on the way, we raise some issues for an alternative diagnosis of the data which appeals to context dependence. In Section 4 we propose that both conditionals and conditional probabilities are best understood in terms of the notion of supposition. The view is in the spirit of Adams, but it improves on his original theory in two ways: (i) it allows for a more general understanding of conditional probability in terms of a model-theoretic restriction operation rather than in terms of the ratio formula; and (ii) it allows us to connect Adams' Thesis with a plausible general story about the semantics of conditionals—namely, the idea that *if*-clauses are devices to restrict sets of possibilities.[7] In Section 5, we respond to two salient objections to this line of thought, both of which go back to Lewis: (i) if the relevant probabilities do not obey standard probability theory, why even call them probabilities? and (ii) if conditionals do not express propositions, how to interpret logical compounds of conditionals? In Section 6 we use our findings to diagnose why Adams' Thesis is thought to fail in certain contexts: these are contexts that encourage us to estimate probabilities using reasoning patterns which—while valid for factual sentences, and thus usually reliable—are invalid for epistemic conditionals. In Section 7, we discuss recent work by Goldstein and Santorio and explain how it relates to the conclusions of this paper. Section 8 concludes and mentions some open problems.

[6] Though they, too, have been challenged: see Mandelkern and Khoo (2019).

[7] On the connection between Adams' Thesis and the restrictor view of conditionals, see also Kratzer (1986); Egré and Cozic (2011); Rothschild (2012); these accounts, however, are concerned with overt probability statements of the form "It is x probable that if A then B" rather than with probability judgments about plain conditionals "If A, then B". The difference is essential, since on the views above, the truth of "It is x probable that if A then B" does not imply that the probability of "If A, then B" is x.

2 Probabilities of conditionals are odd

In this section we will take a classic triviality result and, rather than letting it unfold in a vacuum, we look closely at how it plays out in a concrete scenario. If A and B are sentential clauses, we denote by $\mathsf{A} \Rightarrow \mathsf{B}$ the indicative conditional with antecedent A and consequent B, and by $p(\mathsf{A})$ the probability of A in the context under consideration. The triviality result we will consider, discussed by Skyrms (1980), McGee (1989), and Fitelson (2015), brings out a tension between (i) standard probability theory (ii) Adams' Thesis and (iii) the probabilistic Import-Export principle:

$$p(\mathsf{A} \Rightarrow (\mathsf{B} \Rightarrow \mathsf{C})) = p(\mathsf{A} \wedge \mathsf{B} \Rightarrow \mathsf{C})$$

From these assumptions we can derive the conclusion that the probability function p is trivial: in particular, one cannot find three incompatible propositions each of which has positive probability. Contrapositively, the result implies that in any situation in which the relevant probability function is non-trivial, one of the three assumptions above must fail: standard probability theory, Adams' Thesis, or Import-Export. This suggests the following way of making progress: look at a specific scenario involving non-trivial probabilities, and see what happens. Which of the above assumptions do we get a counterexample to? The example we look at is inspired by the discussion in Khoo and Santorio (2018).

Triviality in action. Consider the following scenario: a fair die was just rolled. The outcome has not yet been revealed. Say the outcome is *low* if it is 1, 2, or 3, and *high* if it is 4, 5, or 6. The die is fair, so each outcome has the same probability, 1/6. We will take for granted that, for sentences not including conditionals, probabilities and conditional probabilities are the ones given by standard probability theory: so, the probability that the outcome was even is 1/2, the conditional probability that it was even given that it was low is 1/3, and so on.

Consider first the following conditional:

(1) If the outcome was low, it was a two. $\mathsf{low} \Rightarrow \mathsf{two}$

What is the probability of (1) in the given context? Adams' Thesis yields an answer:

$$p(\mathsf{low} \Rightarrow \mathsf{two}) = p(\mathsf{two}|\mathsf{low}) = 1/3 \tag{i}$$

Now consider the following nested conditional:

(2) If the outcome was even, then if it was low, it was a two. $\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})$

What is the probability of (2)? Adams' Thesis and Import-Export yield an answer:

$$p(\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})) = p(\mathsf{even} \wedge \mathsf{low} \Rightarrow \mathsf{two}) = p(\mathsf{two}|\mathsf{even} \wedge \mathsf{low}) = 1 \tag{ii}$$

However, Adams' Thesis and standard probability theory lead to a different conclusion. Standard probability theory validates the following principles for any A and B:

Ratio: $p(\mathsf{B}|\mathsf{A}) = p(\mathsf{A} \wedge \mathsf{B})/p(\mathsf{A})$, if $p(\mathsf{A}) > 0$;

Conjunction: $p(\mathsf{A} \wedge \mathsf{B}) \leq p(\mathsf{B})$.

Putting these principles together we obtain:

Upper Bound: $p(\mathsf{B}|\mathsf{A}) \leq p(\mathsf{B})/p(\mathsf{A})$.

Now using Adams' Thesis, Upper Bound, and the result of Equation (i), we get:

$$p(\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})) = p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) \leq \frac{p(\mathsf{low} \Rightarrow \mathsf{two})}{p(\mathsf{even})} = \frac{1/3}{1/2} = 2/3 \quad \text{(iii)}$$

We reached a contradiction: Equation (ii) says that the probability of the conditional in (2) is 1, while Equation (iii) says it is at most 2/3. Given the triviality result mentioned above and given that the probability space we have here is non-trivial, a contradiction is what we expected. Now, however, we can use intuitions about our specific sentences in context to diagnose what has gone wrong.

Diagnosis. Equations (ii) and (iii) yield different conclusions about the probability of the conditional $\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})$: one says that this probability is 1, the other says it is at most 2/3. Which is right? Well, consider again sentence (2): intuitively, this is clearly something we should be certain of in the given scenario. The correct probability value is thus 1, in accordance with Equation (ii).

This immediately allows us to absolve Import-Export: this principle is only used in Equation (ii), but that is not where things go wrong. This speaks against proposals such as those of Khoo and Mandelkern (2018) and Fitelson (2019), that aim to avoid triviality results by invalidating Import-Export. These proposals invalidate the equation that gives the correct result, and instead validate the one that gives the wrong result.

The problem, then, lies within Equation (iii). This in turn uses the result of Equation (i). Could the problem lie there? Equation (i) says that the probability of $\mathsf{low} \Rightarrow \mathsf{two}$ is 1/3. Intuitively, this is the correct result.

The problem must then lie in the first two steps of Equation (iii). The first step uses Adams' Thesis, the second Upper Bound. If Adams' Thesis is right, we have:

$$p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = p(\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})) = 1$$

On the other hand, if Upper Bound is right, we have :

$$p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) \leq \frac{p(\mathsf{low} \Rightarrow \mathsf{two})}{p(\mathsf{even})} = 2/3$$

These are different predictions about the conditional probability $p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even})$. Which one is correct? Given the information that the die landed even, what attitude is it rational to have towards the conditional $\mathsf{low} \Rightarrow \mathsf{two}$? Should we regard it as certain? Intuitively, we should: the only outcome which is both low and even is 2; thus, if we are given the information that the outcome was even, we can be certain that if it was low, it was a two. Thus, the intuitively correct result is $p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = 1$, in accordance with Adams' Thesis and in contravention to Upper Bound.

This speaks against proposals that invalidate Adams' Thesis for nested conditionals, while retaining Upper Bound (Khoo, 2022): these theories settle for the wrong conditional probability here.

Thus, our scenario provides no reason to deny Adams' Thesis for nested conditionals. Quite the opposite: Adams' Thesis performs perfectly in each of the three cases we considered. Instead, it supports the conclusion that probabilities of conditionals violate the Upper Bound condition stemming from standard probability theory. In our scenario, we have the following situation:

$$1 = p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) \not\leq \frac{p(\mathsf{low} \Rightarrow \mathsf{two})}{p(\mathsf{even})} = 2/3$$

Summing up: looking at a triviality result unfold in a concrete setting provides neither a counterexample to Adams' Thesis nor to Import-Export, but instead supports the conclusion that conditional probabilities of conditionals do not behave in accordance with standard probability theory.[8]

Looking at different triviality results leads to similar conclusions. E.g., looking at Lewis-style triviality proofs supports the conclusion that probabilities of conditionals do not validate the law of total probability (when the latter is formulated in terms of conditional probabilities).[9]

But *why* do probabilities of conditionals behave in this strange way?

3 Why are probabilities of conditionals odd?

Based on considerations similar to ours, Bradley (2006) argues that Adams' Thesis provides motivation from departing from standard probability theory in favour of a non-monotonic probability theory. This view, however, leaves something important unexplained: violations of probability theory occur specifically in connection with conditionals. We would like a theory that does not simply give up the laws of probability theory across the board, but rather accurately predicts within what boundaries they hold and where we might find counterexamples.

In order to develop such a theory, it is crucial to distinguish two different issues. First, there is the issue of how to model credence. Credence attaches primarily to propositions, not sentences.[10] This component of the theory has to do with the modeling of idealized cognitive states. Second, there is the issue of how a certain a credal state—a graded outlook on how things are—determines an assignment of probabilities to sentences in context. This is the linguistic component of the theory.

In the literature, assumptions about these two components are often lumped together into a single assumption: that the set of sentences ordered by entailment forms a Boolean algebra, and that the map $p(\cdot)$ assigning probabilities to sentences satisfies the algebraic version of the Kolmogorov axioms over this algebra. We think that, instead, a clear diagnosis of the problem requires carefully disentangling

[8]Recall that Upper Bound follows from Ratio and Conjunction. Since Upper Bound fails, at least one of Ratio and Conjunction must fail for conditionals. This highlights two ways to obtain a non-classical theory of the probabilities of conditionals: a more moderate line preserves the Kolmogorov axioms, and in particular Conjunction, and drops the Ratio principle for conditional probabilities; in this kind of approach, probabilities of conditionals are statically classical, and their non-classicality is revealed only by the dynamics of supposition, reflected by conditional probabilities; the account of Goldstein and Santorio (2021) that we will discuss in Section 7 falls in this class. A more radical line, by contrast, holds that probabilities of conditionals are non-classical already at the static level, i.e., do not satisfy the Kolmogorov axioms (in fact, perhaps the undelying algebra of contents is not even Boolean). The view we defend in this paper is compatible with both approaches, although in Section 5 we will point to some data that suggest that the Conjunction principle is indeed false for conditionals, and that the more radical kind of departure from classicality may therefore be warranted.

[9]An example: given that the outcome was odd, we can exclude that if the outcome was low it was a two. So intuitively we should have $p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{odd}) = 0$. Using the fact that $p(\mathsf{low} \Rightarrow \mathsf{two}) = 1/3$ and $p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = 1$ we have the following violation of total probability:

$$p(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even})p(\mathsf{even}) + p(\mathsf{low} \Rightarrow 2|\mathsf{odd})p(\mathsf{odd}) = 1 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2} = \frac{1}{2} \neq \frac{1}{3} = p(\mathsf{low} \Rightarrow \mathsf{two})$$

[10]One can be confident to different degrees in propositions that one cannot articulate linguistically: for instance, I can consider it likely that my friend will show up tonight with a certain expression on his face that I am familiar with, though I cannot describe that particular expression linguistically.

assumptions pertaining to the two components of the theory. Let us start with the first: the formal modeling of credence.

3.1 Modeling credence

Probability spaces as credal states. A central notion in probability theory is that of a *probability space*. We will zoom in on the special case of *discrete* probability spaces.[11] A discrete probability space is a pair $s = \langle W, m \rangle$ where W, the *sample space*, is a non-empty set and $m : W \to [0,1]$ is a map such that $\sum_{w \in W} m(w) = 1$.

When we use such an object to model a credal state, we think of the elements of W as representing different ways things might be, called *possible worlds*.[12] We think of the number $m(w)$ as quantifying the extent to which the agent believes w to correspond to the way things actually are.

Given a proposition X, the extent to which an agent with state s believes X to be true is obtained by summing over the ways in which X may be true:

$$\pi_s(X) = \sum_{\substack{w \in W \\ X \text{ true in } w}} m(w)$$

Since $\pi_s(X)$ depends only on the set of worlds where X is true, we can for our purposes identify X with this set of worlds and simplify the above definition:

$$\pi_s(X) = \sum_{w \in X} m(w)$$

In this way, probability theory provides a natural model of credal states and of how such states determine probabilities of propositions. Note that given this model, probabilities of propositions obviously satisfy the Kolmogorov axioms.

Conditionalization and supposition. Suppose $s = \langle W, m \rangle$ is a discrete probability space, and $X \subseteq W$ is such that $\pi_s(X) > 0$. We can naturally obtain a new probability space by restricting the sample space to X and rescaling the probabilities of the remaining worlds so that they sum up to 1 again. That is, we can define a new space $s_X = \langle X, m_X \rangle$, where m_X is given by:

$$m_X(w) = \frac{m(w)}{\pi_s(X)}$$

The transformation $s \mapsto s_X$ is called the *conditionalization* of s on X. The *conditional probability* of a proposition Y given X in s is just the probability that Y is true, after conditionalization on X. More formally, since after conditionalization the only Y-worlds that remain are those in $Y \cap X$, we define:

$$\pi_s(Y|X) := \pi_{s_X}(Y \cap X)$$

[11] The points we make carry over to the general case, but since they are already clearly visible in the much simpler setting of discrete spaces, we stick to this setting for the sake of clarity.

[12] The "possible worlds" at stake here need not be maximally specific; they just need to be complete with respect to those aspects that are relevant for the reasoning situation at hand. For instance, in our die roll scenario, each outcome could correspond to a possible world. Alternatively, we could think of the elements in W as cells of a partition of the space of maximally specific possible worlds, where two worlds are in the same partition cell if they agree with respect to the features relevant in the context.

Simply spelling out the definitions, we obtain that the ratio formula holds:

$$\pi_s(Y|X) = \pi_{s_X}(Y \cap X) = \sum_{w \in Y \cap X} m_X(w) = \sum_{w \in Y \cap X} \frac{m(w)}{\pi_s(X)} = \frac{\pi_s(X \cap Y)}{\pi_s(X)}$$

Notice that at the heart of the notion of conditional probability is the conditionalization operation, which is connected to the natural idea of *restricting* a space. The ratio formula has a derivative significance. We will come back to this point in Section 4.

When probability spaces are used to model credence, the operation of conditionalization has a very natural interpretation: it can be used to model the process of making a supposition.[13] When we suppose that X is true, we enter a hypothetical state where X is treated as certain. All worlds in which X is false are dropped from consideration. The probabilities of the remaining worlds are normalized so that the probability of X becomes 1. The relative probabilities of the X worlds remain the same, since merely supposing that X is true gives us no information on which X worlds are more or less likely. This means that, if our original credal state was modeled by a space s, then our hypothetical state after the supposition of X is modeled by s_X. Thus, as long as $\pi_s(X) > 0$, the process of supposing X is naturally modeled by the conditionalization operation.[14] A conditional probability $\pi_s(Y|X)$ is therefore naturally interpreted as a measure of one's conditional credence in Y under the supposition that X.

3.2 Attaching probabilities to sentences

We saw how an agent's credal state s, modeled as a probability space, determines a certain assignment π_s of probabilities to propositions. But what about *sentences*? In order to make predictions about that, we also need a semantic theory that pairs sentences in context with semantic values, and a bridge principle that determines how a credal state and the semantic value of a sentence jointly determine the probability of the sentence in the given context.

A straightforward theory is often presupposed here: sentences in context express propositions, and (conditional) probabilities of sentences are (conditional) probabilities of the corresponding propositions. Let us give a name to these assumptions.

Factualism. Relative to a context, a sentence A expresses a proposition $\boldsymbol{\mathsf{A}}$.[15]

Factual Bridge. If A and B express propositions $\boldsymbol{\mathsf{A}}$ and $\boldsymbol{\mathsf{B}}$ in the relevant context, the probability of $\boldsymbol{\mathsf{A}}$ relative to a credal state s is $p_s(\mathsf{A}) = \pi_s(\boldsymbol{\mathsf{A}})$, and the conditional probability of A given B is $p_s(\mathsf{A}|\mathsf{B}) = \pi_s(\boldsymbol{\mathsf{A}}|\boldsymbol{\mathsf{B}})$.

These assumptions imply that probabilities and conditional probabilities of sentences obey standard probability theory: for they boil down to probabilities and

[13] It is also notably connected to the process of learning new information, but we focus on supposing since it is directly relevant for conditionals.

[14] Conditionalization models one specific mode of supposition—call it the *epistemic* mode. Different supposition processes, usually modeled by imaging (Lewis, 1976) or by interventions in a causal network (Pearl, 2000), are relevant for different kinds of hypothetical reasoning and, typically, for the assessment of subjunctive conditionals.

[15] Notice that Factualism only becomes a well-defined claim once we specify the fragment of language to which it is supposed to apply. For instance, even in the standard view, Factualism is supposed to be restricted to declarative sentences. What we will question below is not Factualism *per se*, but its application to a certain class of sentences.

conditional probabilities of certain propositions, and we saw that under the probabilistic model of credence these obey the Kolmogorov axioms as well as the ratio formula for conditional probability.

However, we argued in the previous section that probabilities of conditionals do *not* satisfy standard probability theory. Thus, something must be wrong in the story we just described.

3.3 What to give up?

The puzzle we faced above seems linguistic in nature: it arises specifically in connection with a certain class of sentences—conditionals.[16] This suggests that the problem does not stem from the probabilistic model of credence and supposition, but rather from our assumptions about how probabilities attach to conditionals. If this is right then either (i) conditionals do not express propositions, or (ii) they do express propositions but their probabilities are not derived from these propositions in the way described by Factual Bridge. We will first consider a way to pursue option (ii), raise a problem for it, and then turn to option (i).

Rejecting Factual Bridge. Some authors (Van Fraassen, 1976; Douven and Verbrugge, 2013; Bacon, 2015) have defended the view that while conditionals express propositions, the proposition expressed by the conditional depends in part on a contextually relevant credal state. It would be natural to argue that this complicates the bridge principle for conditional probability. To see why, suppose we are in a credal state s and we ask about the probability of a conditional C: then *relative to* s, C will express a proposition $\mathbf{C}_s$, and the probability of C in s is the probability of this proposition:

$$p_s(\mathsf{C}) = \pi_s(\mathbf{C}_s)$$

But now suppose we ask about the probability of C *given* B: we can understand this as asking about the probability of C relative to a new credal state $s_{\mathbf{B}}$ resulting from the supposition of B. That is the probability of the proposition expressed by C, not in s, but in $s_{\mathbf{B}}$:

$$p_s(\mathsf{C}|\mathsf{B}) = p_{s_{\mathbf{B}}}(\mathsf{C}) = \pi_{s_{\mathbf{B}}}(\mathbf{C}_{s_{\mathbf{B}}})$$

Absent further constraints relating $\mathbf{C}_s$ and $\mathbf{C}_{s_{\mathbf{B}}}$, there are no systematic relations between $p_s(\mathsf{C})$ and $p_s(\mathsf{C}|\mathsf{B})$. In particular, the upper bound principle $p_s(\mathsf{C}|\mathsf{B}) \leq p_s(\mathsf{C})/p_s(\mathsf{B})$ may well fail.

This approach may seem to provide a way to reconcile the observations we made with the idea that conditionals express propositions after all. We will argue, however, that a variant of our observations is still problematic for this strategy.

Recall that on the contextualist view, when a speaker asserts A, the context of utterance fixes the corresponding proposition $\mathbf{A}$ that was expressed: $\mathbf{A}$ is then the content put forward by the speaker. Different interlocutors, with different credal states, may then assess this proposition, judging its probability or its conditional probability given some assumption. In so doing, however, they are not *changing* the proposition that the speaker expressed.

Now let us see what happens in the case of a conditional assertion. Consider again our die scenario. The outcome of the roll has not yet been revealed. We overhear our friend Alice make the following guess:

[16] As we mentioned in the introduction, related problems arise for sentences involving epistemic modals.

(3) If the outcome was low, it was a two. $\mathsf{C} := \mathsf{low} \Rightarrow \mathsf{two}$

We leave the room before the outcome is revealed. Now let us ask:

- What is the probability that Alice is right?
- What is the probability that Alice is right, given that the outcome was even?

The intuitive answers are, respectively, 1/3 and 1. Our judgments are about what Alice said. If by means of her utterance Alice expressed proposition C, then our judgments are about C. And so, relative to our credal state s, we have:

- $\pi_s(\mathsf{C}) = 1/3$
- $\pi_s(\mathsf{C}|\mathsf{E}) = 1$

But this is impossible, since probabilities of propositions obey the upper bounding principle, so we should have $\pi_s(\mathsf{C}|\mathsf{E}) \leq \pi_s(\mathsf{C})/\pi_s(\mathsf{E}) \leq 2/3$. So the explanation of the observations in terms of context dependence has a problem.

The above strategy relies on the idea that when we assess conditional probabilities we make an assumption, and that assumption is supposed to change the proposition expressed by the conditional. But for this to work here, we need to assume that the mere fact that *we* later make an assumption can retroactively affect the proposition that Alice expressed by her utterance.

This possibility, if admitted, requires a radical change in the standard picture of the semantics/pragmatics interface: if different hearers, making different assumptions, can each privately change the proposition expressed in a different way, it no longer makes sense to think that there is one proposition that has actually been expressed. There seem to be many private propositions, one for each hearer at each time.[17] While we have no conclusive objection to this sort of relativist view, a different diagnosis seems to us more plausible.

Rejecting Factualism. The more natural diagnosis, in our view, is the one given by Adams (1975) and others in the "suppositionalist tradition" (Gibbard, 1980; Edgington, 1986, 1995; Bennett, 2003): epistemic conditionals of the sort we are concerned with here do not express propositions, and their probabilities are, therefore, not probabilities of propositions.

Note that to claim that conditionals do not express propositions is not to claim that they express nothing at all. It is just to say that what they express is not the same sort of object that run-of-the-mill sentences express; crucially, it is not the sort of object on which the probability function π_s is defined.[18]

Independent evidence against factualism for conditionals. The idea that conditionals of the sort we are considering do not express propositions is supported by observations independent of probability judgments. We will mention two.

[17] A view along these lines has in fact been advocated by Weatherson (2009), precisely in connection with conditionals, although that particular account would not predict the above judgments.

[18] The non-factualist position is often presented (including by its proponents, see Edgington, 1986; Bennett, 2003) as the view that conditionals lack truth conditions. But one must be careful: to say that conditionals do not express propositions is to say that they do not have bivalent truth conditions relative to possible worlds, since that would determine a corresponding proposition. The view *is* compatible with conditionals having truth conditions of some other kind; in particular, it is compatible with the idea that they have partial or trivalent truth conditions, in line with a tradition that goes back to de Finetti (1936) (see also McDermott, 1996; Milne, 1997; Cantwell, 2008; Rothschild, 2014; Lassiter, 2020; Égré *et al.*, 2020).

While we do not claim these are knock-down arguments, they provide at least a *prima facie* case against factualism for conditionals.

Argument from acceptance. One argument against factualism for conditionals comes from considering qualitative intuitions about *full acceptance* of conditionals, rather than quantitative intuitions about credence (for an argument along similar lines, see Edgington, 1986).

Consider an agent, Alice, whose qualitative doxastic state is represented by s—the set of worlds that she thinks might be actual. Suppose that A and B are factual, so they express propositions $\mathbf{A}$ and $\mathbf{B}$, and suppose $s \cap \mathbf{A} \neq \emptyset$, so Alice considers it possible that A is true.

If $s \cap \mathbf{A} \subseteq \mathbf{B}$, that means that according to Alice, if A is true, B is also true. In this situation, Alice is intuitively in a position to accept the conditional $\mathsf{A} \Rightarrow \mathsf{B}$. On the other hand, if $s \cap \mathbf{A} \not\subseteq \mathbf{B}$, then according to Alice it is possible that A is true and B false. In this case, Alice is not in a position to accept $\mathsf{A} \Rightarrow \mathsf{B}$. In sum, then, Alice is in a position to accept $\mathsf{A} \Rightarrow \mathsf{B}$ just in case $s \cap \mathbf{A} \subseteq \mathbf{B}$.

Now suppose that accepting $\mathsf{A} \Rightarrow \mathsf{B}$ amounts to accepting the truth of a proposition $\mathbf{C}$. Then we also have that Alice accepts $\mathsf{A} \Rightarrow \mathsf{B}$ iff $s \subseteq \mathbf{C}$. Putting things together, we have that, for all s which intersect $\mathbf{A}$:

$$s \subseteq \mathbf{C} \quad \text{iff} \quad s \cap \mathbf{A} \subseteq \mathbf{B}$$

Provided $\mathbf{A}$ and $\mathbf{B}$ are mutually consistent, this can hold only if $\mathbf{C} = \overline{\mathbf{A}} \cup \mathbf{B}$. This means that, if $\mathsf{A} \Rightarrow \mathsf{B}$ expresses a proposition, this must be the one expressed by the material conditional. Since there are powerful arguments against the material account (for instance, it yields terrible predictions about probabilities and about compounds), we should reject the hypothesis that $\mathsf{A} \Rightarrow \mathsf{B}$ expresses a proposition.

Argument from disagreement. Another argument comes from an example from Gibbard (1980). Here is a simpler variant. Imagine a context where there are three marbles in an urn: one is red, one blue, and one yellow. Alice, Bea, and Carla each draw one marble. Alice got red, Bea blue, and Carla yellow. Each of them cannot see what the others drew. Now, suppose an external observer utters the following conditional:

(4) If Carla did not draw yellow, she drew blue.

Now, Alice is clearly in a position to fully accept (4): she knows she has red, so Carla has either yellow or blue; thus if she does not have yellow, she has blue. Bea, on the other hand, can fully reject (4): she knows *she* has blue, so she can conclude that if Carla did not draw yellow, she drew red—and not blue.

Now suppose that (4), as uttered in the given context, expresses a proposition $\mathbf{C}$. Then, presumably, to accept or reject (4) would be to take $\mathbf{C}$ to be true or false, respectively. Then Alice's information state s_a should consist only of worlds where $\mathbf{C}$ is true, and Bea's state s_b only of worlds where $\mathbf{C}$ is false. So s_a and s_b should be disjoint. But these states are not disjoint, since neither Alice nor Bea is ruling out the actual world, and thus the actual world lies in the intersection $s_a \cap s_b$. Hence, what Alice and Bea are doing when they accept or reject (4) is not to accept or reject a proposition.[19]

[19] A standard response to Gibbard (Kratzer, 1986) invokes the idea that the same conditional can express different propositions when uttered by different speakers. This response does not help here, since we are imagining a single utterance that Alice and Bea are reacting to.

Looking beyond indicative conditionals. Besides being independently motivated, this diagnosis of the problem has another merit: it generalizes in the right way. If the non-standard behavior of probabilities of conditionals is due to the fact that conditionals do not express propositions, we should expect that analogous problems might arise with other classes of sentences which do not express propositions, and that, conversely, such problems will *not* arise with factual sentences. This prediction seems to be borne out. In the recent literature, it has been argued, independently of triviality results, that sentences involving epistemic modals like *might* and *probably* are non-factual (see, a.o., Yalcin, 2011; Swanson, 2011; Willer, 2013): and indeed, as expected, in the literature we find triviality results about the probabilities of such sentences (Russell and Hawthorne, 2016; Goldstein, 2019b). Similarly, though more controversially, Edgington (2008) argued that counterfactual conditionals are also non-factual, and indeed we find corresponding triviality results for counterfactuals (Williams, 2012; Leitgeb, 2012; Santorio, 2022a; Schultheis, 2022). Conversely, to our knowledge, there are no triviality results involving sentences that have not been independently argued to be non-factual. This suggests that a story that identifies non-factualism as the source of the problem is on the right track.

4 Explaining the observations

Having rejected a component of the classical view, we now need to say how to replace it. Our aim here is not to develop a full-fledged theory, but to say enough to be able to explain our judgments on probabilities of conditionals and to bring certain conceptual points in focus. Our explanation will rely on three key assumptions; as we will indicate at the end of this section, these assumptions can be vindicated by different specific theories of the compositional semantics of conditionals and of the way probabilities attach to sentences. While these theories differ in some important respects, at a suitable level of abstraction—the one at which our explanation is formulated—they can thus be seen as agreeing on the solution to the triviality problem.

Assumptions. Our first assumption is that the *factual bridge* is correct in connection with factual sentences: if a sentence expresses a proposition, then its probability is the probability of that proposition.

- **Assumption 1: factual bridge**
  If A expresses a proposition $\mathbf{A}$, then $p_s(\mathsf{A}) = \pi_s(\mathbf{A})$.

Our second assumption is that conditional probabilities of sentences, just like conditional probabilities of propositions, are to be understood as probabilities that result from conditionalization.

- **Assumption 2: conditional probabilities**
  If A expresses a proposition $\mathbf{A}$, then $p_s(\mathsf{B}|\mathsf{A}) = p_{s_{\mathbf{A}}}(\mathsf{B})$.

Notice that we do *not* define conditional probabilities in terms of the ratio formula. The more fundamental understanding of conditional probability, in our view, is in terms of the restriction operation. It is this characterization that directly connects with the construal of conditional probabilities as probabilities under a supposition. The ratio formula is just a way to *calculate* conditional probabilities; this way can be shown to be adequate for probabilities of propositions, but there is no reason to expect that it will give the right results for non-factual sentences, given that their probabilities are not probabilities of corresponding propositions.[20]

Our third assumption concerns probabilities of conditionals. How do probabilities attach to conditionals, if not via a proposition expressed? Our answer is: via the Ramsey test procedure. To judge the probability of a conditional is to judge the probability of the consequent under the supposition of the antecedent.

- **Assumption 3: probabilities of conditionals**
  If A expresses a proposition $\mathbf{A}$, then $p_s(\mathsf{A} \Rightarrow \mathsf{B}) = p_{s_{\mathbf{A}}}(\mathsf{B})$.

Given these assumptions, the validity of Adams' Thesis has a simple explanation: when thinking and communicating in situations of uncertainty, it is often very useful to zoom in on a restricted set of possibilities; a clause of the form "*if* A" is an object-language device to achieve such a restriction; the locution "*given* A" is a meta-language device to achieve the same result.

- **Fact 1: Adams' Thesis.**
  For any factual A and any C: $p_s(\mathsf{A} \Rightarrow \mathsf{C}) = p_s(\mathsf{C}|\mathsf{A})$.

We also predict the validity of probabilistic import-export, at least for factual antecedents: for any propositions $\mathbf{A}$ and $\mathbf{B}$, the states $(s_{\mathbf{A}})_{\mathbf{B}}$ and $s_{\mathbf{A} \cap \mathbf{B}}$ are defined in the same circumstances, and whenever defined, they are equal.

- **Fact 2: Probabilistic Import-Export.**
  For any factual A, B and any C: $p_s(\mathsf{A} \Rightarrow (\mathsf{B} \Rightarrow \mathsf{C})) = p_s(\mathsf{A} \wedge \mathsf{B} \Rightarrow \mathsf{C})$.

Finally, from Adams' Thesis and Import-Export we also get a version of Adams' Thesis for conditional probabilities.

- **Fact 3: Strong Adams' Thesis.**
  For any factual A, B and any C: $p_s(\mathsf{B} \Rightarrow \mathsf{C}|\mathsf{A}) = p_s(\mathsf{C}|\mathsf{A} \wedge \mathsf{B})$.

Back to the dice. These assumptions are sufficient to diagnose what is going on in our initial example. First, the predictions: given that Adams' Thesis and Import-Export are both valid, we get the intuitively correct results for the two conditionals we considered:

- $p_s(\mathsf{low} \Rightarrow \mathsf{two}) = p_s(\mathsf{two}|\mathsf{low}) = 1/3$
- $p_s(\mathsf{even} \Rightarrow (\mathsf{low} \Rightarrow \mathsf{two})) = p_s(\mathsf{even} \wedge \mathsf{low} \Rightarrow \mathsf{two}) = p_s(\mathsf{two}|\mathsf{even} \wedge \mathsf{low}) = 1$

We also correctly predict that, *given that the outcome was even*, it is certain that if it was low, it was a two. Indeed, by Strong Adams' Thesis, we have:

- $p_s(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = p_s(\mathsf{two}|\mathsf{even} \wedge \mathsf{low}) = 1$

As a consequence, the observed violation of Upper Bound is predicted:

- $p_s(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = 1 \not\leq 2/3 = p_s(\mathsf{low} \Rightarrow \mathsf{two})/p_s(\mathsf{even})$

Moreover, we can now see clearly *why* this violation occurs. What happens here is that the probability of $\mathsf{low} \Rightarrow \mathsf{two}$, which is initially low, jumps up to 1 upon conditionalizing on even, which has probability 1/2. The Upper Bound principle tells us that probabilities of propositions cannot increase so dramatically upon conditionalization. By contrast, *conditional* probabilities of propositions can increase unboundedly upon conditionalization. To put it more precisely: for any value $\varepsilon > 0$, it is possible to find instances of probability spaces s and propositions X, Y, Z such

[20] We are not the first to argue that the ratio formula should not be viewed as a definition of conditional probabilities, but as an analysis: see also Hájek (2003) and Cantwell (2021).

that $\pi_s(X|Y) < \varepsilon$ and $\pi_s(Z) > 1 - \varepsilon$ and yet such that $\pi_{s_Z}(X|Y) = 1$.[21] If probabilities of conditionals are not probabilities of propositions, but *conditional* probabilities of propositions, then it is not surprising that they turn out to violate the Upper Bound principle. On the contrary, it is expected.[22]

Similarly, we can now see why, unlike probabilities of factual sentences, probabilities of conditionals violate the law of total probability: this is because their probabilities do not track the size of a certain portion of the logical space, which can be measured by splitting the space into pieces and taking a weighted average.[23]

Updating Adams. The view that we just advocated is very much in the spirit of Adams (1975): conditionals do not express propositions, yet they have probabilities, and these are conditional probabilities. However, we think that the view improves on the original theory of Adams in some important respects.

In Adams' theory, the basic semantic object used to interpret conditionals is an assignment p of probabilities to factual sentences; the notion of conditional probability is understood in terms of the ratio formula, $p(\mathsf{B}|\mathsf{A}) := p(\mathsf{A} \wedge \mathsf{B})/p(\mathsf{A})$; this notion is then used to extend p to simple conditionals by setting $p(\mathsf{A} \Rightarrow \mathsf{B}) = p(\mathsf{B}|\mathsf{A})$ when A, B are factual.

In the view that we just sketched, by constrast, the basic semantic object is a discrete probability space—a set of possible worlds with weights assigned to them. Such a space can be taken to represent the credal state of an idealized agent, in a way which is independent of language altogether. In this setting, conditionalization can be understood in terms of a transformation $s \mapsto s_X$ that restricts a credal state to the X-worlds and rescales the weights accordingly. It is this fundamental semantic operation that underlies both the notion of conditional probability, and the semantics of epistemic conditionals of the sort we have been looking at.

In our view, this model-theoretic perspective is conceptually more natural: it reflects the fact that credence attaches primarily to propositions, and only derivatively to sentences; moreover, it has two other advantages.

Advantage 1: conditional probabilities of conditionals. Suppose Alice and Bob make the following claims:

(5)	Alice: If the outcome was low, it was a two.	$\mathsf{low} \Rightarrow \mathsf{two}$
	Bob: If the outcome was high, it was a six.	$\mathsf{high} \Rightarrow \mathsf{six}$

Now consider the following questions:

[21] Proof: take a natural number n such that $1/n < \varepsilon$. Let s be a uniform distribution on a space of n^2 worlds $\{w_{ij} \mid i, j \in \{1, \dots, n\}\}$. Let $X = \{w_{11}\}$, $Y = \{w_{11}, \dots, w_{1n}\}$, $Z = X \cup \overline{Y}$. Then X, Y, Z satisfy the above description.

[22] Since we construe conditional probabilities in terms of restriction, and not in terms of the ratio formula, we can predict the violation of Upper Bound without making any assumptions about probabilities of Boolean compounds involving conditionals. Nevertheless, it is worth noting that since the Upper Bound principle follows from the principles we called Ratio and Conjunction, any extension of our view that assigns probabilities to such compounds is bound to invalidate at least one of these principles.

[23] In particular, we predict the particular counterexample which we claimed to be empirically supported in Footnote 9, since Strong Adams' Thesis yields:

- $p_s(\mathsf{low} \Rightarrow \mathsf{two}) = p_s(\mathsf{two}|\mathsf{low}) = 1/3$
- $p_s(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = p_s(\mathsf{two}|\mathsf{even} \wedge \mathsf{low}) = 1$
- $p_s(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{odd}) = p_s(\mathsf{two}|\mathsf{odd} \wedge \mathsf{low}) = 0$.

(6) Given that the outcome was even, what is the probability that $\left\{\begin{array}{l}\text{Alice is right?}\\ \text{Bob is right?}\end{array}\right.$

Intuitively, the right answers are, respectively, 1 and 1/2.

Adams' theory does not account for these intuitions. Conditionals are assigned probabilities, but not *conditional* probabilities. This is because conditional probabilities are defined by the ratio formula. Thus, a conditional probability $p(\mathsf{B} \Rightarrow \mathsf{C}|\mathsf{A})$ would amount to the ratio $p(\mathsf{A} \wedge (\mathsf{B} \Rightarrow \mathsf{C}))/p(\mathsf{A})$, but the theory does not say how to assign probabilities to compounds such as $\mathsf{A} \wedge (\mathsf{B} \Rightarrow \mathsf{C})$.[24]

By constrast, in our view, we can make perfectly good sense of conditional probabilities of conditionals: these are not derived from probabilities of compounds via the ratio formula; instead, they are simply probabilities of conditionals under a supposition. In formulas, we have:

$$p_s(\mathsf{B} \Rightarrow \mathsf{C}|\mathsf{A}) = p_{s_{\mathsf{A}}}(\mathsf{B} \Rightarrow \mathsf{C}) = p_{s_{\mathsf{A}\cap\mathsf{B}}}(\mathsf{C}) = p_s(\mathsf{C}|\mathsf{A} \wedge \mathsf{B})$$

This immediately yields the desired predictions for the above example:

- $p_s(\mathsf{low} \Rightarrow \mathsf{two}|\mathsf{even}) = p_s(\mathsf{two}|\mathsf{low} \wedge \mathsf{even}) = 1$
- $p_s(\mathsf{high} \Rightarrow \mathsf{six}|\mathsf{even}) = p_s(\mathsf{six}|\mathsf{high} \wedge \mathsf{even}) = 1/2$

Advantage 2: semantics of conditionals. Another advantage becomes evident when we ask what is the *semantics* of an indicative conditional. In Adams' theory, the claim that the probability of a conditional is a conditional probability constitutes the semantics of the indicative conditional—there is no more fundamental semantic level from which this result is deduced. Other suppositionalists have followed this idea (e.g., Bennett, 2003, pp. 58, 104). This, however, is simply not plausible as a general story about the semantics of conditionals, even of the indicative kind. Consider the following two examples:

(7) a. If you don't give a presentation you have to submit a paper.
    b. If it is sunny, we usually have breakfast in the garden.

A speaker uttering either of these sentences is not expressing a high conditional probability. Under their most salient reading, these are factual claims. The first sentence, (7-a), states the existence of a certain conditional obligation; if it is intended as a specification of how to get credit for a course, it can be paraphrased as: all the worlds where you get credit but do not give a presentation are worlds where you submit a paper. The second sentence, (7-b), states a fact about breakfasts on sunny days—roughly, that most of them are taken in the garden.

Examples like these point to the following generalization: an *if*-clause is a restricting device. In (7-a), the *if*-clause restricts the set of possibilities the modal 'have to' ranges over to those where you don't give a presentation. In (7-b), it restricts the set of breakfast occasions that 'usually' ranges over to those where it is sunny. This general idea is known as the *restrictor view* of conditionals.[25]

[24] McGee (1989), describes how to extend Adams' theory so as to make the probabilities of compounds defined. This would allow us to define conditional probabilities of conditionals via the ratio formula, but the results are not in accordance with intuition. For instance, for the questions in (6) we would get the results 5/9 and 4/9 instead of 1 and 1/2.

[25] The view originated from Lewis's work on adverbs of quantification (Lewis, 1975), and was put forward as a general view of conditionals by Kratzer (1981, 1986) (see Rothschild, 2012, for recent discussion). It is, however, important to distinguish the restrictor idea as such from Kratzer's specific theory. Indeed, we will propose below to reject a key assumption of Kratzer's theory.

Our proposal has the merit of showing that Adams' Thesis need not be taken as providing the *semantics* of indicative conditionals, but can be viewed as following from the more basic, and more general, idea that *if*-clauses are restrictors. One way to achieve this is the following: suppose that sentences are assessed not (just) relative to a world parameter, but (also) relative to an information state parameter s. The idea is familiar from many previous accounts of conditionals (Veltman, 1985; Gillies, 2004, 2009; Yalcin, 2007; Bledin, 2014; Ciardelli, 2020; Punčochář and Gauker, 2020). Suppose we model s not as a simple set of worlds, but as a probability space. We can then explain why Adams' Thesis holds if we reject Kratzer's assumption that an *if*-clause always restricts the range of some, possibly covert, operator occurring in the syntax of the conditional. Instead, it is natural to assume that one thing an *if*-clause can be used to restrict is the information state relative to which the conditional is evaluated. We characterize a conditional $\mathsf{A} \Rightarrow \mathsf{B}$ as being an *epistemic* conditional if the restriction targets the information state parameter (as opposed to, e.g., the modal base for a deontic modal, as in (7-a)). Thus, if $\mathsf{A} \Rightarrow \mathsf{B}$ is an epistemic conditional, the semantics of the *if*-clause makes sure that evaluating $\mathsf{A} \Rightarrow \mathsf{B}$ relative to a state s amounts to evaluating B in the restricted state s_{A}. In particular, if we assess the conditional for probability, we get precisely the content of Assumption 3: the probability of $\mathsf{A} \Rightarrow \mathsf{B}$ relative to s is the probability of B relative to s_{A}.

This is only a sketch of how our Assumption 3 can be derived from a restrictor semantics, but hopefully sufficient to convey the idea. The task of developing a precise compositional semantics for a language involving *if*-clauses that can restrict, among other semantic parameters, an underlying information state, is to some extent independent of the present issues, and is taken up in a separate paper (Ciardelli, 2022).[26]

5 Two objections

In this section we defend our view from two commonly made objections—both of which go back to Lewis (1976).

Objection 1: probabilities or "probabilities"? The first objection goes as follows: if the quantities we assign to conditionals do not conform to the laws of probability, they are not probabilities; they must be something else, and then they had better be called something else. Lewis (1976) famously put it as follows:

[26] This is not the only account compatible with the assumptions we laid out above. There are at least two other salient options, both of which connect in a natural way to the restrictor idea. The first option, explored in slightly different versions by Cantwell (2021) and Goldstein and Santorio (2021), is to (i) interpret sentences with respect to sequences of worlds; (ii) take *if*-clauses to restrict the set of worlds appearing in the sequence; and (iii) show how to lift a given probability distribution on worlds to one on sequences. The second option (de Finetti, 1936; McDermott, 1996; Milne, 1997; Cantwell, 2008; Rothschild, 2014; Lassiter, 2020; Égré *et al.*, 2020) is to (i) interpret sentences as having partial (or trivalent) truth conditions with respect to worlds; (ii) take an *if*-clauses to be a partializer, which makes a conditional neither true nor false when the antecedent is not true; and (iii) define the probability of a sentence as the probability that the sentence is true, given that it has a truth value. (The resulting truth functions can also be used to get the restrictor behavior exemplified in (7) in the way described by Belnap (1970), see also Rothschild (2012).) A detailed comparison of the relative merits of these options must wait for another occasion. Our main aim here is not to settle on a specific account, but to argue for a certain solution to the triviality problem which can be vindicated by more than one account.

> But if it be granted that the "probabilities" of conditionals do not obey the standard laws, I do not see what is to be gained by instisting on calling them "probabilities".

To formulate our response to this objection, it is helpful to think of an analogy. Consider the notion of *logical consequence.* This could be taken to be a term of art, which is intrinsically defined by the laws of classical logic or by its construal in terms of truth-preservation. Yet, many theories—dynamic semantics, for example—come with notions of consequence which violate classical logic, and which are not understood in terms of truth-preservation. One could raise a complaint like the above one about these notions: if these relations do not obey the standard laws, why even call them "consequence"? The reason is that "consequence" is not just a term of art of classical logic, but also a pre-theoretical notion: we have some pre-theoretical understanding about what follows from what, and about what constitutes a sound piece of reasoning, and we use this understanding as a guide in our theorizing about consequence. We then come up with theories that attempt to formally capture or "explicate" the notion of consequence. Thus, there is a clear sense in which revisionary notions of consequence are still notions of consequence, insofar as they are intended as alternative theories of the same pre-theoretical notion.

We think that the situation is exactly the same with probability. Like "consequence", "probability" is not just a term of art, but also a pre-theoretical notion. We think and talk all the time about things being certain, likely, unlikely, or more likely than something else, and we ascribe high or low probability to statements made by other people, including conditionals. Just like our pre-theoretical intuitions about consequence might in some domains diverge from the predictions of classical logic, and be better vindicated by a different approach, so also our pre-theoretical intuitions about probability might in some domains diverge from the predictions of standard probability theory (in fact, the two things are likely to be related, since standard probability theory builds on classical logic). Nothing should prevent us, then, from developing an alternative account of (sentential) probabilities; doing so does not amount to changing the subject.

Objection 2: compounds. Lewis had another "non-conclusive objection" to a non-factualist approach to conditionals: once we deny that conditionals express propositions, we can no longer avail ourselves of the standard account of logical connectives and quantifiers to interpret compounds involving conditionals. We then face the challenge of developing a new theory of these items—at least to the extent that conditionals can embed under them.

First, let us note that this objection is non-conclusive indeed. There is near-universal consensus, for instance, that interrogatives do not express propositions. But interrogatives can occur embedded under connectives and quantifiers, as illustrated by the sentences in (8); so there are other phenomena in language that require a more general theory of these logical items.

(8) a. Where are you going, and when will you come back?
b. Who has a car we can borrow, or where can we rent one?
c. What present did every guest bring?

Moreover, in the case of conditionals, a departure from the standard theory is independently motivated. The standard treatment of connectives simply does not yield the right predictions for compounds of conditionals. In fact, once we look at such compounds, we find further violations of standard probability theory. For instance, take again our die scenario and consider the following sentence:

(9) If the outcome was even it was a two, and if it was odd it was either one or three.
$(\mathsf{even} \Rightarrow \mathsf{two}) \wedge (\mathsf{odd} \Rightarrow \mathsf{one} \vee \mathsf{three})$

Intuitively, this sentence simply claims that the outcome was either one, two, or three—i.e., that the outcome was low. Its probability should accordingly be 1/2. But the probability of the first conjunct, $\mathsf{even} \Rightarrow \mathsf{two}$, is 1/3. Thus, it seems that the probability of a conjunction of conditionals can be higher than the probability of one conjunct.[27] This might at first strike us as exceedingly strange, but there is a perfectly reasonable diagnosis: the two conjuncts are each making a restricted claim about a subset of the possibilities; conjunction "glues" these restricted claims together into an unrestricted claim about the whole set of possibilities. The restricted claims have different probabilities, which balance out when they are glued together. So the probability of a conjunction of this kind ends up being intermediate between the probabilities of the conjuncts, rather than below both of them.

In view of this, the putative advantage of a factualist account in interpreting compounds turns into a disadvantage. If conditionals express propositions, to which standard logical operators can be applied, observations such as the one we just made are really puzzling. But if conditionals express something other than ordinary propositions, then it is unsurprising that logical operations on them will also exhibit some unfamiliar features; and it is also clear why a more general story about such compounds is needed.[28]

6 Failures of Adams' Thesis?

If our view is right, Adams' Thesis is not just true, but in a sense, trivially true: conditional probabilities and probabilities of epistemic conditionals are nothing but probabilities under a supposition. We think this is a good prediction: when we ask a theoretically unbiased speaker about the probability of an epistemic conditional, they seem to just interpret the question as asking about a conditional probability.

Yet, the literature contains several putative counterexamples to Adams' Thesis (McGee, 2000; Kaufmann, 2004; Rothschild, 2013; Moss, 2018; Khoo and Santorio, 2018; Magidor, 2019). Aren't these, then, counterexamples to our theory? We will argue that they are not.

Some purported counterexamples (Kaufmann, 2004; Rothschild, 2013; Moss, 2018) arise, in our view, from the fact that indicative conditionals involving *will* have, alongside the epistemic reading to which Adams' Thesis applies, a second "on-

[27] The same point can be made by considering the sentence:

(i) If the outcome was low it was even, and if it was high it was even. $(\mathsf{low} \Rightarrow \mathsf{even}) \wedge (\mathsf{high} \Rightarrow \mathsf{even})$

Intuitively, in our context (i) is just a roundabout way to claim that the outcome was even. Its probability should be 1/2. But the probability of the first conjunct is 1/3.

[28] Lewis's challenge has in fact been taken up by proponents of non-propositional approaches. In the setting of a trivalent approach, McDermott (1996) gives an account of connectives that seems to make reasonable predictions about the probabilities of propositional compounds of conditionals. Accounts of compounds have also been given by proponents of information-based approaches (Veltman, 1985; Dekker, 1993; Gillies, 2004, 2009; Yalcin, 2007; Punčochář and Gauker, 2020). However, these accounts are either silent about probabilities, or they yield unintuitive predictions. To our knowledge, the question of how to account for the probabilities of compounds like (9) within a information-based theory is open.

tic" reading. The existence of this reading is not a matter of stipulation, but follows from the possibility to interpret the *if*-clause as restricting the modal *will*. Conditionals including a modal are generally known to be ambiguous in this way (this is the ambiguity between C-readings and O-readings discussed in Geurts, 2004), and it would be really surprising if conditionals involving *will* were an exception. The probability of a conditional under an ontic reading need not equal the corresponding conditional probability, but this is not a counterexample to Adams' Thesis, provided the latter is understood as a generalization about conditionals under the epistemic reading (or C-reading, in the terminology of Geurts (2004)).

Other putative counterexamples, however, do not involve the modal *will*. We would like to argue that, in these cases, the judgments stem from subtle fallacies: they are not based just on linguistic intuition, but on complex probabilistic reasoning; and this reasoning uses probabilistic principles which, while valid for factual sentences, and thus usually reliable, are invalid when applied to conditionals.

Example 1: Boxes and tickets. This example is a variation by Khoo and Santorio (2018) on examples by Kaufmann (2004) (see also Khoo, 2016):[29]

> You drew a ticket from one of two boxes. The box you drew from was selected randomly, and you don't know which was selected. In Box 1, there were 100 tickets, of which 90 were red and 81 red with a dot. In Box 2, there were 100 tickets, of which 10 were red and 1 red with a dot. How likely is it that:
>
> (10) If you drew a red ticket, it had a dot.

In these examples there are two attractive methods of reasoning which reach different answers:

- *1/2*: In Box 1, 9/10 of the red tickets have a dot and in Box 2, 1/10 of the red tickets have a dot. So supposing you drew from Box 1, there was a 9/10 chance that you drew a dotted ticket if you drew red, and supposing you drew from Box 2, there was a 1/10 chance that you drew a dotted ticket if you drew red. You had the same probability of drawing from Box 1 as from Box 2, so all-in-all the probability of (10) is the average: 1/2.
- *High*. If you drew a red ticket, then most likely it was drawn from Box 1. Therefore, since almost all of the red tickets in Box 1 have a dot, the probability of (10) is high.

Only the second answer is in accordance which Adams' Thesis. But, which answer gives the right probability of the conditional in this context? Khoo (2016) says *both*: he draws the conclusion that (10) has two readings, and then goes on to offer theories of how and why these two different readings arise.

Our view allows for a simpler account of these observations: (10) only has the high probability reading, in accordance with Adams' Thesis. The reasoning in favour of the answer 1/2 is indeed natural, but fallacious.[30] Notice that the reasoning above relies on the law of total probability and then to an instance of

[29] Kaufmann's original example uses a conditional involving *will*, and we think should be analyzed as ambiguous between an ontic and an epistemic reading as sketched above.

[30] Rothschild (2013) considers a related idea. He suggests that some putative violations of Adams' Thesis might involve a cognitive error, in particular they might be due to a version the base rate fallacy. As we will see, our account of why the reasoning is fallacious is of a different kind.

Strong Adams' Thesis (using the obvious abbreviations):

$$\begin{aligned} p_s(\mathsf{red} \Rightarrow \mathsf{dot}) &= p_s(\mathsf{red} \Rightarrow \mathsf{dot} \mid \mathsf{one}) \cdot p_s(\mathsf{one}) + p_s(\mathsf{red} \Rightarrow \mathsf{dot} \mid \mathsf{two}) \cdot p_s(\mathsf{two}) \\ &= p_s(\mathsf{dot} \mid \mathsf{red} \wedge \mathsf{one}) \cdot p_s(\mathsf{one}) + p_s(\mathsf{dot} \mid \mathsf{red} \wedge \mathsf{two}) \cdot p_s(\mathsf{two}) \\ &= 9/10 \cdot 1/2 + 1/10 \cdot 1/2 = 1/2 \end{aligned}$$

But, as we argued, the first step is fallacious: the law of total probability is invalid for conditionals like (10). However, the law *does* hold for factual sentences. Our first line of reasoning seems natural while nevertheless being mistaken because it involves over-extending a usually correct and reliable principle.

Our diagnosis is supported by the observation that the judgment 1/2 can be shown to be wrong by the following reasoning, which seems compelling:

> Look, we have here a situation with 200 tickets, each of which is drawn with the same probability. There are in total 100 red tickets, and of these, 82 have a dot. Since each red ticket was equally likely to be picked, it is very likely (82%) that if you drew a red ticket, it had a dot.

The relevant judgment is, in our view, not very different from the mistaken judgment that people give in the Monthy Hall problem: it is indeed "natural", but it can be shown to be wrong, and the mistake can be given a plausible explanation.

Example 2: Sherlock Holmes. Perhaps the best known putative counterexample to Adams' Thesis is due to McGee (2000). Here is the scenario:

> Murdoch drowned in the lake in his garden. There's no evidence that his death was not accidental. A person whom you believe to be Sherlock Holmes—but who may also be a fraud—tells you that Murdoch was murdered, almost certainly by Brown, but if not by Brown then by someone else.

What credence should we assign to the following conditional?

(11) If Brown didn't kill Murdoch, someone else did.

McGee reasons as follows: it is very likely that the informant is the infallible Sherlock Holmes, and *given* that he is Sherlock Holmes and that he is certain of (11), (11) is certainly right. Hence we should consider (11) very likely.

However, the conditional probability is not necessarily high: in fact, it can be very low provided the probability of the person being a fraud is higher than the probability that he *is* Holmes but the culprit is not Brown, as shown in Figure 1.

We argue McGee's reasoning commits a fallacy akin to that of the previous example. It relies on the following pattern of reasoning: it is likely that A, and given A, it is certain that B; therefore, it is likely that B:

$$p_s(\mathsf{A}) \text{ is high and } p_s(\mathsf{B}|\mathsf{A}) = 1, \text{ therefore } p_s(\mathsf{B}) \text{ is high.}$$

While valid when A and B are factual sentences, this reasoning is invalid when B is a conditional, for precisely the reasons brought out in this example.

For a much simpler illustration, in the context of a die roll, consider:

(12) A The outcome was not 1.
B If the outcome was 1 or 3, it was 3.

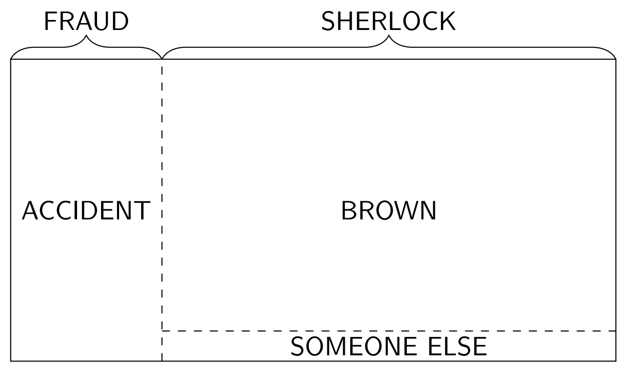

Figure 1: A visual representation of probabilities in McGee's example

Here, $p_s(\mathsf{A}) = 5/6$ and $p_s(\mathsf{B}|\mathsf{A}) = 1$, and yet $p_s(\mathsf{B}) = 1/2$. In fact, using the idea described in Footnote 21, for every $\varepsilon > 0$ we can give examples where $p_s(\mathsf{A}) > 1 - \varepsilon$ and $p(\mathsf{B}|\mathsf{A}) = 1$, and where nevertheless $p_s(\mathsf{B}) < \varepsilon$. This shows that McGee's reasoning, while valid when applied to factual sentences, can lead us astray when applied to conditionals.

7 Comparison with Goldstein and Santorio

After completing a first draft of this paper, we became aware of recent work by Goldstein and Santorio—unpublished at the time of writing, but available online—which makes some points closely related to the ones we make here. We think it is helpful to briefly discuss how the views in the two papers relate.

Among the many points of convergence, both papers locate the source of the triviality problem in the fact that epistemic conditionals differ semantically from factual sentences in that they do not express standard propositions. Moreover, Goldstein and Santorio (henceforth, G&S) agree with us that the probability of an epistemic conditional equals the probability of the consequent on the supposition of the antecedent, which need not coincide with the value given by the ratio formula.

While we have no substantial point of disagreement with the theory of G&S, we advocate a somewhat different perspective than the one G&S take in their paper. For G&S, the main lesson of the triviality problem is that conditionals call for a new theory of how credences change upon supposition: standard conditionalization should be abandoned in favor of a new update rule, 'hyperconditionalization'.

The lesson that we would like to draw is somewhat different. We think we need to distinguish carefully two components of a theory of credences: on the one hand, the modeling of credal states and of supposition (the lower part of Figure 2); and on the other hand, the account of how sentences are assigned probabilities based on a credal state (the upper part of Figure 2). Insofar as the first component is concerned, we see no reason to deviate from standard probability theory: a credal state can be modeled as a probability space s, and the process of supposing a proposition X can be modeled by conditionalizing s to X—i.e., restricting s to X and rescaling the probabilities of the remaining worlds. However, what we need to reconsider is how probabilities attach to sentences in context, i.e., the vertical arrows in Figure 2. Once we grant that some sentences do not express standard propositions, their

$$\begin{tikzcd} p_s & & p_{s_X} \\ s \arrow[u] \arrow[rr] & & s_X \arrow[u] \end{tikzcd}$$

Figure 2: Credal states, conditionalization, and induced probabilities

probabilities may well behave differently from probabilities of propositions. In particular, there is no reason to expect that the way these probabilities change upon conditionalization can be reliably estimated by the ratio formula.[31]

We should stress that our perspective is compatible with G&S's theory. In this theory, sentences are associated with certain fine-grained semantic contents (formally, these are sets of paths, where paths are sequences of possible worlds). Some of these contents correspond to standard propositions, while others—in particular, the ones expressed by epistemic conditionals—do not. G&S then explain how a credal state (which they call a 'proto-epistemic space') determines an assignment of probabilities to these fine-grained contents (they call this derived assignment an 'epistemic space') and, thereby, to sentences, including epistemic conditionals. Now, G&S themselves think of the dynamics of supposition as unfolding at the higher level of epistemic spaces (i.e., in the upper part of Figure 2), and this motivates the shift from conditionalization to hyperconditionalization. However, one may take an alternative perspective according to which the dynamics of supposition unfolds by standard conditionalization at the level of the underlying credal states (their proto-epistemic spaces), as in Figure 2, while probabilities attach to sentences via the induced epistemic space. On this alternative perspective, supposition still goes by conditionalization, while hyperconditionalization can be seen as a rule to calculate the effects of supposition on the induced probabilities of sentences.

One last difference is that, on our diagnosis, one may in principle expect even plain (i.e., non-conditional) probabilities of non-factual sentences to diverge from the usual laws of probability theory. By contrast, in G&S's account, probabilities of sentences—factual and non-factual alike—obey the standard Kolmogorov axioms. As we saw in Section 5, there is some reason to think compounds of conditionals in fact violate probability theory, even with respect to non-conditional probabilities.

[31] A confound in the comparison is that G&S take conditional probabilities to be defined by the ratio formula, whereas we understand them in terms of the operation of restricting a probability space to a sub-domain of possibilities. As a result, G&S will claim, for instance, that probabilities of conditionals do not generally align with conditional probabilities, whereas we claim they do. The difference on this point is ultimately terminological and insubstantial. Nevertheless, we think the choice of what to call 'conditional probability' is not entirely arbitrary. In our view, the operation of restricting a probability space to a subset of possibilities is eminently natural from a mathematical point of view, and it is conceptually significant due to its direct link to supposition. The ratio formula, by contrast, is significant only derivatively, as a way of *calculating* conditional probabilities. So, we think when generalizing to a context where the results of restriction and those the ratio formula come apart—as they do in the case of conditionals—it is the former notion that still captures the central idea of conditional probability.

8 Conclusion

Let us sum up the main points we have been making. We started out by observing that our linguistic intuitions support the conclusion that probabilities of epistemic conditionals violate standard probability theory—a point that was made before by Bradley (2006), but which seems to have received little attention in the literature. This gives us a non *ad-hoc* way to block triviality results—since such results assume that probabilities of conditionals can be calculated in accordance to standard laws.

We then considered what explains this observation. After rejecting a diagnosis in terms of context dependency, we argued for an explanation along suppositionalist lines: standard probability theory is best seen as a theory of probabilities of propositions; but there is independent reason to think that epistemic conditionals do not express propositions; if so, their probabilities will naturally not be probabilities of propositions, and therefore, they are not expected to conform to the laws that apply to such probabilities.

We proposed to understand both the semantics of epistemic conditionals and the notion of conditional probabilities in terms of *supposition*, modeled as a restriction operation on a probability space. We argued that this preserves the spirit of Adams' theory, but it improves on it in two ways: it allows us to assign conditional probabilities to conditionals in a natural way, and it allows us to connect the view to a more general, and independently motivated, idea about the semantics of conditionals: that *if*-clauses are restricting devices.

We argued against the common complaint that the resulting probabilities are probabilities 'in name only': *probability*, just like *consequence*, *supposition*, *belief*, etc., is not a mere term of art, but one of a family of concepts about thought and communication that our theories are meant to elucidate. In our view, it is helpful to distinguish two notions, namely, propositional probability (which attaches to propositions) and sentential probability (which attaches to sentences in context). A sentential probability sometimes, but not always, amounts to a propositional probability. But just because probabilities of conditionals are not propositional, that does not mean that they should be viewed as degrees of assertibility rather than degrees of belief.

Finally, we saw that realizing that probabilities of conditionals violate standard probability theory does not just allow us to diagnose what goes wrong in triviality proofs. It also allows us to account for some putative counterexamples to Adams' Thesis: the relevant judgments are not mere linguistic intuitions, but result from complex reasoning which involves applying standard probability theory to the probabilities of conditionals, treating them as if they were propositional probabilities.

Looking ahead, we view this paper as part of a large collective enterprise: understanding how, and to what extent, the theoretical picture developed specifically for factual sentences needs to be revised in order to analyze non-factual sentences, in particular those including epistemic modals and conditionals. This enterprise involves not only assigning suitable semantic values to such sentences, but also formulating more general accounts of logical notions like consequence, speech acts like assertion, attitudes like acceptance, credence, and desire, and mental acts like supposition. For some recent contributions see, among others: Gillies (2004); Yalcin (2007, 2015); Swanson (2011, 2016); Bledin (2014, 2020); Starr (2014); Moss (2015); Goldstein (2019a,b); Ciardelli (2020); Punčochář and Gauker (2020); Hawke and Steinert-Threlkeld (2021); Cantwell (2021); Santorio (2022b); Goldstein and Santorio (2021).

Especially important for the topics discussed here is the issue of supposition. A brand of triviality results that we have not surveyed involves looking at the effects of supposing a conditional, in particular in combination with the negation of its consequent (Russell and Hawthorne, 2016).[32] We think that an appropriate response to these arguments requires recognizing that supposing epistemic sentences is a complicated issue. When we suppose such a sentence, we are not just restricting attention to a set of possibilities; rather, it seems, we are making sure our hypothetical state has a certain property (being compatible with a proposition, assigning it high probability, or entailing a conclusion under a supposition). The dynamics of such suppositions is bound to be unfamiliar. Indeed, we know from Yalcin (2007) that there are pairs of sentences, like $\neg\mathsf{A}$ and $\mathsf{might\ A}$, which are in a sense compatible ($\mathsf{might\ A}$ does not rule out $\neg\mathsf{A}$) and yet cannot be jointly supposed in a consistent way. This suggests that we should be suspicious of arguments that assume that epistemic suppositions behave similarly to factual suppositions. A detailed investigation must be left for future work.

References

Adams, E. (1965). The logic of conditionals. *Inquiry*, **8**(1-4), 166–197.

Adams, E. (1975). *The logic of conditionals: An application of probability to deductive logic*, volume 86. Springer Science & Business Media.

Bacon, A. (2015). Stalnaker's thesis in context. *The Review of Symbolic Logic*, **8**(1), 131–163.

Belnap, N. (1970). Conditional assertion and restricted quantification. *Noûs*, **4**(1), 1–12.

Bennett, J. F. (2003). *A philosophical guide to conditionals*. Oxford University Press.

Bledin, J. (2014). Logic informed. *Mind*, **123**(490), 277–316.

Bledin, J. (2020). Fatalism and the logic of unconditionals. *Noûs*, **54**(1), 126–161.

Bradley, R. (2000). A preservation condition for conditionals. *Analysis*, **60**(3), 219–222.

Bradley, R. (2006). Adams conditionals and non-monotonic probabilities. *Journal of Logic, Language and Information*, **15**(1-2), 65–81.

Cantwell, J. (2008). The logic of conditional negation. *Notre Dame J. Formal Logic*, **49**(3), 245–260.

Cantwell, J. (2021). An expressivist analysis of the indicative conditional with a restrictor semantics. *Review of Symbolic Logic*, **14**(2), 487–530.

Ciardelli, I. (2020). Indicative conditionals and graded information. *Journal of Philosophical Logic*, **49**(3), 509–549.

[32] The triviality results of Stalnaker (1976) and Mandelkern (2021), which employ left-nested conditionals, can also be seen as related, since on a Ramsey test view the interpretation of such a nested conditional involves the supposition of a conditional.

Ciardelli, I. (2021). Restriction without quantification: embedding and probability for indicative conditionals. *Ergo*, **8**(21), 598–636.

Ciardelli, I. (2022). The restrictor view, without covert modals. *Linguistics and Philosophy*, **45**, 1–28.

de Finetti, B. (1936). La logique de la probabilité. In *Actes du congrès international de philosophie scientifique*, volume 4, pages 1–9. Hermann Paris.

Dekker, P. (1993). *Transsentential Meditations. Ups and Downs in Dynamic Semantics.* Ph.D. thesis, ILLC, University of Amsterdam.

DeRose, K. (2010). The conditionals of deliberation. *Mind*, **119**(473), 1–42.

Douven, I. and Verbrugge, S. (2013). The probabilities of conditionals revisited. *Cognitive Science*, **37**(4), 711–730.

Edgington, D. (1986). Do conditionals have truth conditions? *Crítica: Revista Hispanoamericana de Filosofía*, pages 3–39.

Edgington, D. (1995). On conditionals. *Mind*, **104**, 235–329.

Edgington, D. (2008). Counterfactuals. In *Proceedings of the Aristotelian Society (Hardback)*, volume 108, pages 1–21. Wiley Online Library.

Egré, P. and Cozic, M. (2011). If-clauses and probability operators. *Topoi*, **30**.

Égré, P., Rossi, L., and Sprenger, J. (2020). De Finettian logics of indicative conditionals. part I: Trivalent semantics and validity. *Journal of Philosophical Logic*, pages 1–27.

Evans, J. S. B., Handley, S. J., and Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **29**(2), 321.

von Fintel, K. and Gillies, A. (2015). Hedging your *if*s and vice versa. Manuscript.

Fitelson, B. (2015). The strongest possible lewisian triviality result. *Thought: A Journal of Philosophy*, **4**(2), 69–74.

Fitelson, B. (2019). How to model the epistemic probabilities of conditionals. Handout from a talk at MCMP, LMU Munich.

Geurts, B. (2004). On an ambiguity in quantified conditionals. Manuscript, University of Nijmegen. Available at http://ncs.ruhosting.nl/bart/papers/conditionals.pdf.

Gibbard, A. (1980). Two recent theories of conditionals. In *Ifs: Conditionals, beliefs, decision, chance, time*, pages 211–247. Reidel, Dordrecht.

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

Gillies, A. (2009). On truth-conditions for if (but not quite only if). *Philosophical Review*, **118**(3), 325–349.

Goldstein, S. (2019a). Generalized Update Semantics. *Mind*, **128**(511), 795–835.

Goldstein, S. (2019b). Triviality results for probabilistic modals. *Philosophy and Phenomenological Research*, **99**(1), 188–222.

Goldstein, S. and Santorio, P. (2021). Probability for epistemic modalities. *Philosophers' Imprint*, **21**(33), 1–37.

Hájek, A. (2003). What conditional probability could not be. *Synthese*, **137**(3), 273–323.

Hawke, P. and Steinert-Threlkeld, S. (2021). Semantic expressivism for epistemic modals. *Linguistics and Philosophy*, **44**, 475–511.

Jackson, F. (1979). On assertion and indicative conditionals. *The Philosophical Review*, **88**(4), 565–589.

Kaufmann, S. (2004). Conditioning against the grain. *Journal of Philosophical Logic*, **33**(6), 583–606.

Khoo, J. (2016). Probabilities of conditionals in context. *Linguistics and Philosophy*, **39**(1), 1–43.

Khoo, J. (2022). *The meaning of* if. Oxford University Press.

Khoo, J. and Mandelkern, M. (2018). Triviality Results and the Relationship between Logical and Natural Languages. *Mind*, **128**(510), 485–526.

Khoo, J. and Santorio, P. (2018). Probability of conditionals in modal semantics. Lecture Notes for a course at NASSLLI 2018, available at https://justindkhoo.files.wordpress.com/2020/06/ks-nasslli2018.pdf.

Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. Berlin: Walter de Gruyter.

Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society*, **22**(2), 1–15.

Lassiter, D. (2020). What we can learn from how trivalent conditionals avoid triviality. *Inquiry*, **63**(9-10), 1087–1114.

Leitgeb, H. (2012). A probabilistic semantics for counterfactuals. part A. *The Review of Symbolic Logic*, **5**(1), 26–84.

Lewis, D. (1975). Adverbs of quantification. In E. Keenan, editor, *Semantics of Natural Language*, pages 3–15. Cambridge University Press.

Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. In W. Harper, R. Stalnaker, and G. Pearce, editors, *Ifs*, pages 129–147. Reidel.

Magidor, O. (2019). Conditional acceptance. In E. Lepore and D. Sosa, editors, *Oxford Studies in Philosophy of Language Volume 1*, pages 99–121. Oxford University Press.

Mandelkern, M. (2018). Talking about worlds. *Philosophical Perspectives*, **32**(1), 298–325.

Mandelkern, M. (2021). If p, then p! *The Journal of Philosophy*, **118**(12).

Mandelkern, M. and Khoo, J. (2019). Against preservation. *Analysis*, **79**(3), 424–436.

McDermott, M. (1996). On the truth conditions of certain 'if'-sentences. *Philosophical Review*, **105**(1), 1–37.

McGee, V. (1989). Conditional probabilities and compounds of conditionals. *The Philosophical Review*, **98**(4), 485–541.

McGee, V. (2000). To tell the truth about conditionals. *Analysis*, **60**(265), 107–111.

Milne, P. (1997). Bruno de Finetti and the logic of conditional events. *The British Journal for the Philosophy of Science*, **48**(2), 195–232.

Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, **8**(5), 1–81.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford University Press.

Over, D. E. and Evans, J. S. B. (2003). The probability of conditionals: The psychological evidence. *Mind & language*, **18**(4), 340–358.

Pearl, J. (2000). *Causality: models, reasoning and inference*, volume 29. Springer.

Punčochář, V. and Gauker, C. (2020). Indicative conditionals in objective contexts. *Theoria*, **86**(5), 651–687.

Rothschild, D. (2012). A note on conditionals and restrictors. In J. Hawthorne and L. Walters, editors, *Conditionals, Probability, and Paradox: Themes from the Philosophy of Dorothy Edgington*. OUP.

Rothschild, D. (2013). Do indicative conditionals express propositions? *Noûs*, **47**(1), 49–68.

Rothschild, D. (2014). Capturing the relationship between conditionals and conditional probability with a trivalent semantics. *Journal of Applied Non-Classical Logics*, **24**(1-2), 144–152.

Russell, J. S. and Hawthorne, J. (2016). General dynamic triviality theorems. *Philosophical Review*, **125**(3), 307–339.

Santorio, P. (2022a). General triviality for counterfactuals. **82**(2), 277–289.

Santorio, P. (2022b). Path semantics for indicative conditionals. *Mind*, **131**(521), 59–98.

Schultheis, G. (2022). Counterfactual probability. Forthcoming in *Journal of Philosophy*. Available at https://philpapers.org/rec/SCHCAP-36.

Skyrms, B. (1980). *Causal Necessity. A pragmatic investigation of the necessity of laws.* Yale University Press.

Stalnaker, R. C. (1970). Probability and conditionals. *Philosophy of science*, **37**(1), 64–80.

Stalnaker, R. C. (1976). Possible worlds. *Noûs*, **10**(1), 65–75.

Starr, W. (2014). A uniform theory of conditionals. *Journal of Philosophical Logic*, **43**(6), 1019–1064.

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

Swanson, E. (2016). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, **45**(2), 121–146.

Van Fraassen, B. C. (1976). Probabilities of conditionals. In *Foundations of probability theory, statistical inference, and statistical theories of science*, pages 261–308. Springer.

Veltman, F. (1985). *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.

Weatherson, B. (2009). Conditionals and indexical relativism. *Synthese*, **166**(2), 333–357.

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

Williams, R. (2012). Counterfactual triviality: A Lewis-impossibility argument for counterfactuals. *Philosophy and Phenomenological Research*, **85**(3), 648–670.

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

Yalcin, S. (2011). Nonfactualism about epistemic modality. In A. Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press.

Yalcin, S. (2015). Epistemic modality de re. *Ergo, an Open Access Journal of Philosophy*, **2**.