Modal language and Bayesian reasoning I: The expression of uncertainty*

Daniel Lassiter
Stanford University

Abstract Probabilistic models have become increasingly widespread in theoretical linguistics in recent years, including formal semantics and pragmatics. This article begins by outlining the semantics of probability and Bayesian update, explaining how they follow as a simple upgrade of the familiar possible-worlds model of meaning and discourse dynamics. It then turns to survey the large but scattered literature on formal models and experimental investigations of epistemic vocabulary and conditionals that have been informed by probabilistic thinking, with attention to contributions from linguistics, logic, philosophy of language, epistemology, computer science, and experimental psychology.

Keywords: Uncertainty, probability, Bayesian inference, epistemic modality, conditionals, lexical semantics, dynamics

> There may be circumstances in which it is not unwise to cling to illusions, but in science we need a very different attitude, the *inductive attitude*. This attitude aims at adapting our beliefs to our experience as efficiently as possible. ... It requires a ready ascent from observations to generalizations, and a ready descent from the highest generalizations to the most concrete observations. It requires saying "maybe" and "perhaps" in a thousand different shades.
>
> — Pólya (1954: 7)

1 Introduction

Research in formal semantics and pragmatics has historically been oriented toward deterministic logics, but in recent years there has been a surge of interest in integrating quantitative models of uncertainty into the traditional toolkit. This trend has been due to many factors, including but not limited to: the rise of probabilistic models in all areas of computational linguistics; a surge of interest among researchers interested in meaning in psycholinguistic methods and theory, which also make extensive use of probabilistic models; and increased theoretical engagement with quantitative models of uncertainty and decision-making as they are used in other areas of cognitive science, including language processing, reasoning, learning, attention, memory, perception, motor control, and many other domains.

* Thanks to Luis Alonso Ovalle, three anonymous reviewers, and Boram Kim for helpful comments and discussion.

This action been enabled to a considerable extent by the rise of Bayesian methods of analysis. Bayesianism has been around for a long time, with a hard core of advocates in statistics, decision theory, philosophy of science, and epistemology. However, its widespread adoption in psychology and linguistics was hindered for many decades by the computational complexity of Bayesian inference, which made it difficult to determine the predictions even of many simple models. With fast modern computers and new methods of simulation and approximation, this issue is less problematic, and Bayesian models have been widely employed in recent cognitive science. One insight that has emerged clearly from this work is that the traditional opposition between structured/logical vs. statistical/probabilistic models is misguided (e.g., Tenenbaum, Kemp, Griffiths & Goodman 2011). Logical and probabilistic methods are not in competition: they can be combined productively in many ways, and indeed probability is best viewed as an *extension* of classical logic rather than a *competitor*.

For research in formal semantics and pragmatics, there are several points of interface with Bayesian models in other areas. For example, probabilistic models of syntactic and semantic processing drawn from computational psycholinguistics (e.g., Jurafsky 2003; Chater & Manning 2006; Crocker 2010) and from cognitive architectures (Anderson 2009) have considerable relevance to semantic and pragmatic theory (e.g., Brasoveanu & Dotlacil To appear). Bayesian theories of human learning (Perfors, Tenenbaum, Griffiths & Xu 2011; Perfors 2012), including language learning (Chater & Manning 2006) provide another important point of theoretical contact. A third connection involves the use of computational models that integrate Bayesian reasoning into a game-theoretic perspective on formal pragmatics (Frank & Goodman 2012; Goodman & Lassiter 2015; Franke & Jäger 2016).

This article focuses on a different interface: the potential for Bayesian models of knowledge representation, reasoning, and decision-making to inform the *representation* of information, the *lexical semantics* of modal language, and *reasoning* processes that involve modal language. I will survey some of the ideas about knowledge representation and modal semantics/reasoning that have recently emerged from this perspective, starting with a brief primer on Bayesian reasoning. To begin, we'll review the semantics of probability and see how it might be used to illuminate what Pólya (1954) called "'maybe' and 'perhaps' in a thousand different shades".

2 The "inductive attitude": Probability and Bayesian inference

The most common way of representing information in semantic and pragmatic research is the "sets-of-worlds" conception, derived from work in modal logic and semantics. Sets-of-worlds models essentially treat the distinction between things that have been excluded from consideration and things that are still relevant to reasoning. For example, the set of "epistemically possible" (or "accessible") worlds $\mathcal{E}$ is a set of possibilities that are not excluded by what is currently known. If some proposition A is known, then $\mathcal{E}$ entails A — $\mathcal{E} \subseteq A$ — and $\mathcal{E}$ excludes A's complement — $\mathcal{E} \cap \bar{A} = \varnothing$. Any set of worlds (proposition) that is entailed by $\mathcal{E}$ is a *necessity* relative to $\mathcal{E}$. A set of worlds B that is not excluded by $\mathcal{E}$—where $\mathcal{E} \cap B \neq \varnothing$—is a *possibility* relative to $\mathcal{E}$. (These notions can equally be used for non-epistemic concepts, such as teleological or deontic possibility and necessity, but this is not our focus at the moment.)

Generalizing a bit, it is often useful to reason about "sets-of-X" for more complicated objects X. If our uncertainty involves something with a bit more structure—say, worlds w paired with a center c, as in centered-worlds theories of the *de se* (Lewis 1979)—we can model uncertainty using a set of $\langle w, c\rangle$ pairs, each consisting of a world and a center. This idea is easily generalized to uncertainty about structured objects of greater complexity—e.g., those used in QUD theory (Roberts 2004) or Inquisitive Semantics (Ciardelli, Groenendijk & Roelofsen 2013).

Sets-of-X models are naturally associated with a simple treatment of learning: learning involves modifying the set to exclude things that are incompatible with what has been learned, without changing anything else. (This is sometimes called "Stalnakerian update", after Stalnaker (1978).) Formally, this means that update is an intersection operation. For example, suppose that the epistemically possible worlds at time t_1 are $\mathcal{E}_1$, and then at t_2 we learn that it is snowing (**snow**), and nothing else. The proposition that it is snowing is represented by the set of possible worlds where it is snowing, **snow**. So, at t_2 our information is represented by $\mathcal{E}_2 = \mathcal{E}_1 \cap \textbf{snow}$ (Figure 1, left column). This is the most conservative update of our information: from the set of worlds representing what we knew at t_1, we exclude any world where it is not snowing, and do not change anything else.

This kind of model encodes some useful properties of states of information and their dynamics. For instance, **snow** is possible both before and after update, as is the proposition A which intersects both $\mathcal{E}_1$ and **snow**. In addition, $\mathcal{E}_2$ is a subset of **snow**, and so **snow** is a necessity relative to $\mathcal{E}_2$. However, note that we have not learned anything about A—at least, nothing that can be stated in the language of possibility and necessity that we are using. A was possible before update with **snow** and remains possible afterwards. For some choices of A this seems sensible—for instance, if A were a proposition about the diameter of the earth, learning about the state of weather has no relevance to it. But for many choices of A we humans seem to have subtler intuitions than this model can capture. For example, if A were *The temperature is between 20 and 30 degrees Fahrenheit*, learning that it is snowing would in many situations render A more likely, without yet making it a necessity. To model such non-categorical effects, we need a model of information and its dynamics that can represent the way that learning about one proposition can be informative about another—without having to do anything as drastic as changing it from a possibility to a necessity, or from a possibility to an impossibility.

The probabilistic model of this situation is strictly richer than the sets-of-worlds model. The key components remain—possible worlds, propositions as sets of worlds, and update as exclusion—but we also tag propositions with numerical values that serve to measure their probabilities. For finite sets of possible worlds W—which we'll focus on, just to keep the math simpler—the additional constraints are:

i. $\mathcal{P}$ is a function from propositions (subsets of W) to numbers between 0 and 1.

ii. Every proposition has a probability greater than or equal to 0.

iii. A tautology has the maximum possible probability — i.e., $\mathcal{P}(W) = 1$.

iv. If propositions A and B are disjoint ($A \cap B = \varnothing$), then the probability of their disjunction ($\mathcal{P}(A \cup B)$) is equal to the sum of their probabilities ($\mathcal{P}(A) + \mathcal{P}(B)$).

In the right column of Figure 1 we see the probabilistic upgrade of the sets-of-worlds model. The

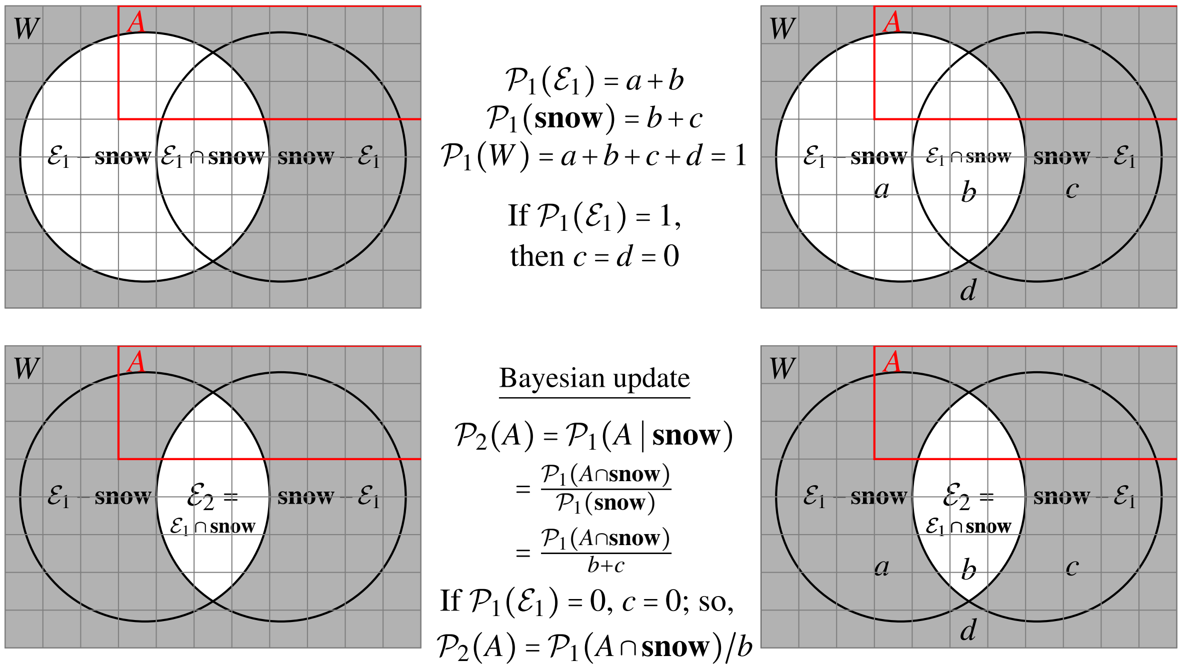

Figure 1 Update in the intersective/Stalnakerian (left column) vs. Bayesian (right column) models. Comparison from left to right reveals that the key features of the sets-of-worlds model have been preserved in the Bayesian model, with the addition of quantitative information. Comparison from top to bottom shows the closely related learning dynamics of the models.

top right panel shows the state at t_1, before we have learned that it's snowing. $\mathcal{E}_1$ is divided into two parts, one where it's not snowing (probability a) and one where it is (probability b). The portion of **snow** outside $\mathcal{E}_1$ has probability c, and the rest of logical space ($W - \mathcal{E}_1 - \mathbf{snow}$) has probability d. By rule (ii) of probability theory we know that $a, b, c, d \geq 0$. Since probabilities of disjoint sets add up by rule (iv), the union of the four sets $\mathcal{E}_1 - \mathbf{snow}$, $\mathcal{E}_1 \cap \mathbf{snow}$, $\mathbf{snow} - \mathcal{E}_1$, and $W - \mathcal{E}_1 - \mathbf{snow}$ has probability $a + b + c + d$. But since these sets cover logical space, their union is W. Since $\mathcal{P}$(W) is 1 by rule (iii), we can infer that $a + b + c + d = 1$. So, the rules of probability theory already place substantial constraints on the kinds of models we can entertain, even without knowing anything about the values a, b, c, and d.

So far we have added nothing to the sets-of-worlds conception beyond the notion of a probability measure and the associated constraints. One additional constraint that is helpful in bridging these models of information is to stipulate that, at any time i, the value of $\mathcal{P}_i(\mathcal{E}_i)$ is 1: the set of epistemically accessible worlds always has probability 1. This implies that any B that is not possible (does not overlap with $\mathcal{E}_i$) has probability 0, and any B that is necessary (is entailed by/a superset of $\mathcal{E}_i$) has probability 1. This additional constraint ensures that we do not wind up with a strange model where, for example, *It is not snowing* is considered impossible (excluded by the relevant $\mathcal{E}_i$) but has a high probability of being true.

The basic approach to update in a Bayesian model is to take in some new observation—here, *It's snowing*—and incorporate it into the model in the most conservative way possible. Here again, "conservative" means that we ensure that we don't change anything except as demanded by our evidence, as Polya's *inductive attitude* demands ("adapting our beliefs to our experience as efficiently as possible"). The most common update rule for probabilistic models is *conditioning*, also known under the moniker *Bayesian update*. Conditioning has a close connection to Stalnakerian update, which starts with a contraction of the set of worlds initially considered possible, $\mathcal{E}_1$, to the subset compatible with the observation **snow**—so, an update function that takes $\mathcal{E}_1$ and **snow** and outputs $\mathcal{E}_2 = \mathcal{E}_1 \cap \mathbf{snow}$. When we apply this operation to a probabilistic model, some probability will typically be lost: whatever was assigned under $\mathcal{E}_1$ to **no snow**. This portion of logical space now has probability 0, and so we must somehow adjust the probabilities of propositions that overlap $\mathcal{E}_2$ in order to ensure that the rule $\mathcal{P}(W) = 1$ is satisfied. It turns out that conditioning is the most conservative way to do this, in the sense that it is the only update operation that does not add any information beyond the fact that the observation **snow** is true (Williams 1980).

Update by conditioning requires that if we learn **snow** at time t_2, the updated probability measure $\mathcal{P}_2$ will assign any proposition A a probability equal to

$$\mathcal{P}_1(A \mid \mathbf{snow}) = \mathcal{P}_1(A \cap \mathbf{snow}) / \mathcal{P}_1(\mathbf{snow}).$$

To see why this is necessary, note that—in order to ensure that we don't redistribute probability in a way that is not justified by our observation—$\mathcal{P}_2(A)$ should be proportional to whatever value *It's snowing and A is true* had at t_1, for every A.[1] $\mathcal{P}_2(A)$ is proportional to $\mathcal{P}_1(A \cap \mathbf{snow})$, rather than equal, because the only way to ensure that the updated probability measure meets condition (iii) is to divide everything by the probability of **snow**. So, $\mathcal{P}_1$ and $\mathcal{P}_2$ are related by:

$$\forall A \subseteq W : \mathcal{P}_2(A) = \mathcal{P}_1(A \mid \mathbf{snow}) = \frac{\mathcal{P}_1(A \cap \mathbf{snow})}{\mathcal{P}_1(\mathbf{snow})} \quad \text{(if this ratio is defined)}$$

Since we are also requiring that $\mathcal{P}_i(\mathcal{E}_i) = 1$ at any time i, we can summarize Bayesian update on **snow** as a three-step process. First, perform Stalnakerian update, contracting $\mathcal{E}_1$ to $\mathcal{E}_2 = \mathcal{E}_1 \cap \mathbf{snow}$. Second, modify the probability measure by assigning probability 0 to any proposition excluded by $\mathcal{E}_2$. Third, adjust the probabilities of every other proposition as conservatively as possible while satisfying the requirement that $\mathcal{P}_2(W) = 1$. The result is that $\mathcal{P}_2(A)$ must be equal to $\mathcal{P}_1(A \mid \mathcal{E}_2)$.

The probabilistic upgrade of the sets-of-worlds model still allows us to use the language of possibility and necessity: things excluded by the relevant $\mathcal{E}_i$ are impossible, those entailed by $\mathcal{E}_i$ are necessary, and so on. But we also have access to much more information about the relative degrees of probability of various propositions. Importantly, the probabilistic model allows us to make sense of the intuition that learning **snow** can be informative about a proposition A even when A is compatible with both the original and the updated epistemic states. As we saw, the language of (im)possibility and necessity is not rich enough to treat our intuitions here: **snow** can be informative

1 In a bit more detail: conservative update on **snow** means that, for any two propositions B and C, the ratio of the probabilities-after-update $\mathcal{P}_2(B)/\mathcal{P}_2(C)$ is (if defined) equal to the ratio of the probabilities-before-update $\mathcal{P}_1(B \cap \mathbf{snow})/\mathcal{P}_1(C \cap \mathbf{snow})$. If you have only learned that it's snowing, adjust the relative probabilities of B and C to the extent that you previously thought they were probabilistically related to **snow**, and no more.

about A without making it necessary or impossible. Bayesian update gives us an account of this subtler kind of reasoning. Learning **snow** takes us from $\mathcal{P}_1(A)$ to

$$\mathcal{P}_2(A) = \mathcal{P}_1(A \mid \textbf{snow}) = \mathcal{P}_1(A \cap \textbf{snow})/\mathcal{P}_1(\textbf{snow})$$

For example, the relative probability of A and $\bar{A}$ changes from $\frac{\mathcal{P}_1(A)}{\mathcal{P}_1(\bar{A})}$ to $\frac{\mathcal{P}_1(A \cap \textbf{snow})}{\mathcal{P}_1(\bar{A} \cap \textbf{snow})}$. There is no reason why these ratios should be the same. If A is *The temperature is between 25 and 30 degrees Fahrenheit*, then learning **snow** will tend to increase the probability of A, even if it is not yet a necessity. This is because Bayesian update removes all of the probability mass from possible worlds the temperature would be too warm for snow and redistributes it among the worlds where the temperature is low enough. This is just one example of the usefulness of this more nuanced representation of information.

It's worth noting here that the rule for conditioning is often written in one of the following two forms, both of which may get the title "Bayes' rule": either

$$\mathcal{P}(A \mid B) = \frac{\mathcal{P}(B \mid A) \times \mathcal{P}(A)}{\mathcal{P}(B)}$$

or the even more complicated-looking

$$\mathcal{P}(A \mid B) = \frac{\mathcal{P}(B \mid A) \times \mathcal{P}(A)}{\sum_{A'} \mathcal{P}(B \mid A') \times \mathcal{P}(A')},$$

where the variable A' ranges over all of the relevant alternatives to A (for example, "not-A", or the members of any other set of mutually exclusive and logically exhaustive propositions that includes the target proposition A). These forms are frequently useful. For example, it's common in science and ordinary life to have information about the probability of an effect given some candidate causes—e.g., $\mathcal{P}(\textbf{symptom} \mid \textbf{disease})$ and $\mathcal{P}(\textbf{disease})$—when the value you want is the "inverse" probability that a particular disease is the cause of the symptoms—$\mathcal{P}(\textbf{disease} \mid \textbf{symptom})$. Under certain conditions, the formulas above can be used to work out what this value must be, using Bayesian inference.

However, these complicated and often more useful formulations of Bayesian update are equivalent to the simple, semantically motivated version of the basic conditioning rule that that we saw above—$\mathcal{P}(A \mid B) = \mathcal{P}(A \cap B)/\mathcal{P}(B)$. The numerator is the same because, by the definition of conditional probability and some trivial algebra,

$$\mathcal{P}(A \cap B) = \frac{\mathcal{P}(A \cap B)}{\mathcal{P}(A)} \times \mathcal{P}(A) = \mathcal{P}(B \mid A) \times \mathcal{P}(A)$$

Given this equality, the denominator of the most complicated version—$\sum_{A'} \mathcal{P}(B \mid A') \times \mathcal{P}(A')$— is simply a sum over $\mathcal{P}(B \mid A') \times \mathcal{P}(A') = \mathcal{P}(B \cap A')$ for all of the various values of A'. A straightforward consequence of the probability axioms (the "law of total probability") is that, whenever the A' are exhaustive and disjoint, this sum must be equal to $\mathcal{P}(B)$. So, the denominator is just a different way to package the value $\mathcal{P}(B)$, which was the denominator of the original conditioning rule.

3 "Saying 'maybe' and 'perhaps' in a thousand different shades"

An enriched representation of information is also useful because many languages—including English—have a rich vocabulary for talking about the informational status of propositions, with much gradation and nuance. Here I will briefly review some of the ways that probability and Bayesian conditioning have been put to use in analyzing the epistemic vocabulary of English. To keep the discussion manageable, I will not survey alternative approaches or engage in detailed theory comparison. Needless to say, none of the analyses discussed below are uncontroversial or obviously correct, and competing accounts should be studied carefully as well.

3.1 Theoretical perspectives

Epistemic expressions deal with information—aspects of knowledge, belief, and (un)certainty.[2] This description isn't totally precise, but it is enough to get a sense of the rich vocabulary that many languages, including English, have for expressing the nuance of our informational states. Some clear cases in English are *must, might, may, likely, probable, possible, certain, know, believe*, and *doubt*. Until recently most linguistically-oriented work on epistemic vocabulary concentrated on modals that were at the "extremes": auxiliaries like *must* and *might*, verbs like *know*, and adjectives like *possible*. These items provide a nice domain of application for the sets-of-worlds model of information, since they can plausibly be analyzed using "all" and "some"-quantifiers over sets of epistemically accessible worlds. For example, *It must/might be snowing* can plausibly be glossed as "In all/some epistemically accessible worlds, it is snowing".

However, there are other epistemic items that can't be treated in this way. Consider, for example,

(1) The Bills are more likely to win the Super Bowl than the Rams are.

It is not at all clear how to map the meaning of (1) onto a sets-of-worlds model, where our primary tools for analysis are quantification using "all" and "some", negation, and sub- or superset relations. But the Bayesian upgrade makes available a plausible analysis: (1) is true just in case the condition in (2) is satisfied.

(2) $\mathcal{P}(\textbf{Bills win}) > \mathcal{P}(\textbf{Rams win})$

One nice feature of this analysis is that *likely* is a gradable adjective, and probabilities are a kind of degrees. So, a compositional analysis of (1) is available that closely parallels a standard treatment of the gradable adjective *excited* in sentences like (3), where we are comparing degrees of excitement (von Stechow 1984; Kennedy 2007).

(3) The Bills are more excited than the Rams are.

Similarly we can analyze (1) as comparing degrees of probability provided by our probabilistic model.

2 It would perhaps be more accurate to distinguish "epistemic" modals for knowledge and "doxastic" modals for belief, but this distinction is not made consistently in the literature, and it is an open question whether and how it is relevant to the epistemic system of English and other languages.

On this analysis, the bare form *The Bills are likely to win* asserts that the Bills' probability of winning exceeds a contextually determined threshold θ—just as, in the degree semantics for gradable adjectives, *The Bills are excited* indicates that the Bills' degree of excitement exceeds a contextually determined threshold. See Yalcin 2007; Portner 2009; Lassiter 2010, 2017a; Klecha 2014 for more on the position of epistemic adjectives such as *likely* within the theory of gradable adjectives in general.

Many other epistemic items could be analyzed in a similar way. For example, a Bills victory could be *more probable* or *more certain* than a Rams victory. We might try to cash out the degrees implicit in these comparisons as degrees of probability. The natural analysis of (e.g.) *The Bills are certain to win* would then be that the Bills' probability of winning is 1 or close to 1. *Possible* is a controversial case, since English speakers seem to disagree about whether one thing can be *more possible* than another; in any case, once we have a Bayesian model of uncertainty in hand we could treat *It is possible that the Bills will win* as indicating that the probability of this event is greater than 0.

Belief and doubt are attitudes that comes in degrees. We can believe one thing more strongly than another, have more doubt regarding one matter than another, and so on. Given this, it is natural to ask whether we can make sense of belief and doubt from a probabilistic perspective: for instance, we might analyze *I believe (doubt) that it will snow* as saying that I assign a sufficiently high (resp. low) probability to the proposition **snow**. This suggestion is extremely plausible on face, but it also faces significant problems where belief is concerned. On the one hand, I can believe something without being certain of it. (An example from a photo caption found on the web: "I believe this was 7th street but I'm not certain.") So, if there is a probability threshold for belief it must be less than 1: belief is weak (Hawthorne, Rothschild & Spectre 2016). On the other hand, any particular choice of "weak" threshold runs up against Kyburg's (1961) famous Lottery paradox. If the threshold for belief is high but less than 1, then I am committed, in a sufficiently large lottery, to believing each member of a certain large set of propositions ({**Ticket 1 won't win, Ticket 2 won't win, ...**}) without believing the conjunction of those propositions (**No ticket will win**). This puzzle has generated a large literature and a variety of solutions: see, for instance, Williamson 2000; Lin & Kelly 2012; Leitgeb 2014. While there is obviously some connection between probability and belief, it has turned out to be surprisingly difficult to spell it out in a foolproof way.

Going Bayesian doesn't require us to analyze everything in terms of probabilities, of course. For instance, we could still follow the long tradition in modal semantics of treating items like *know* and *must* as universal quantifiers over the epistemically accessible worlds, and *might, perhaps*, and *maybe* as existential quantifiers over the same set. If $\mathcal{P}_i(\mathcal{E}_i)$ is required to be 1 at all times i, as we assumed above, this analysis would have the effect of enforcing probabilistic entailments for these items. Anything that I *know* at time t_i must have probability 1, since it is true throughout $\mathcal{E}_i$. Likewise, if my information verifies *It must be snowing*, I must assign probability 1 to snow; and if $\mathcal{P}_i(\textbf{snow}) > 0$, then it *might* snow according to $\mathcal{E}_i$. (However, see Lassiter 2016 for experimental and corpus evidence that *must* is weaker, and *might* stronger, than this analysis would predict.)

Alternatively, we might attempt to define the meanings of these items partly in terms of probabilities, in a direct way: for instance, it might be part of the meaning of *must* that *It must be snowing* indicates that **snow** has high or even maximal probability. It is important to be clear here that any probabilistic entailments of these items do not, on either style of analysis, exhaust the meanings of *know* and *must*. For example, knowledge seems to have entailments involving justification, and *must* may be connected to indirectness of evidence. For discussion of some of the thorny linguistic and philosophical issues at stake here, see for instance Williamson 2000; Hawthorne 2004; von Fintel & Gillies 2010; Lassiter 2016.

There are many other items—for instance *clear, evident, plausible, confident*, and so-called epistemic *ought* and *should*—whose meaning is clearly related to probability, but which seem to have a richer semantic texture. If you ask me whether Mary made it home yet and I reply that *She should be there by now*, I presumably judge that she is probably home. However, this does not seem to be an entailment per se. There is nothing odd about the reply *She should be there by now, but she isn't*, yet the matched statement with *probably*—*She is probably there by now, but she isn't*—is bizarre (Copley 2004; Swanson 2015; Yalcin 2016). Similarly, I might judge that a witness' account of the events surrounding a crime is *plausible*, and yet know that it is false. (I was actually the perpetrator, but covered my tracks well.) This would not make sense if *The account is plausible* implied that the probability of the account is moderately high: this would contradict my inside knowledge that the account's probability of truth is 0. Assertions of clarity are interesting in a different way. If I think it is *clear* that the Bills will win the Super Bowl, this certainly suggests that *I* find a Bills victory highly probable. However, assertions of clarity have additional implications involving the publicity of the information used to reach a conclusion: I cannot assert that *It is clear that I had eggs for breakfast* simply because I happen to know that I did (Barker & Taranto 2003; Barker 2009; Crone 2016).

Several central items in the epistemic vocabulary of English are difficult to treat within a sets-of-worlds conception of information, but yield to a probabilistic analysis that fits well within our broader understanding of the compositional semantics of English. In addition, there are numerous other items whose relationship to probabilistic information is evident, but less direct.

3.2 Experimental and quantitative work

One useful feature of the Bayesian perspective on modal semantics is that it allows one to make precise quantitative predictions and test them experimentally. In addition, probabilistic models make available direct connections between modal reasoning and probabilistic models that are employed in much work in psychology and artificial intelligence. These features make Bayesian thinking very attractive by opening up new perspectives on modal semantics, new theoretical connections with other fields, and new ways of testing and refining theories.

There is a small body of experimental work on the relationship between qualitative and quantitative expressions of uncertainty, starting from a probabilistic perspective on the language of uncertainty. For instance, Wallsten, Budescu, Rapoport, Zwick & Forsyth (1986) asked participants to assign expressions such as *almost certain, probable, likely, good chance, possible*, and *doubtful* to areas on a spinner. They used psychometric methods to infer the ranges of probabilities that participants considered to be appropriately described by the expressions. This method is interesting, though somewhat limited for a number of reasons: it attempts to measure the meanings of expressions outside of sentential and conversational context, and it cannot tease apart semantic and pragmatic factors. For instance, most participants did not appear to think that an event with probability .9 was appropriately described as 'possible'. This is presumably not because such an event is not possible, but because it would be better described in another way—by *likely* or *almost certain*, say.

In another example, Ülkümen, Fox & Malle (2015) examined the use of English epistemic expressions in light of a theoretical distinction between different kinds of probability that is widely employed in statistics, economics, and philosophy (Romeijn 2017) and sometimes in psychology (Kahneman & Tversky 1982; Lagnado & Sloman 2004; Fox & Ülkümen 2011). They showed evidence from ordinary language (or, at least, from the *New York Times*) that people distinguish between epistemic probability—uncertainty as a psychological feature of people—and "aleatory" probability, i.e., objective chance as a feature of an indeterministic world. Their corpus investigation found a number of patterns suggesting that some items (*likely, probability, chance*) were generally used to express objective chance, while others (*certain, sure, confident*) were primarily used to express epistemic probability. Ülkümen et al.'s (2015) interpretation makes no explicit contact with formal semantics, but the questions that they are asking are a natural extension of the general probabilistic perspective on epistemic language described above: given that probability is implicated, what kind of probability is involved? Their conclusions, while broadly consonant with the Bayesian perspective that we have elaborated here, suggest that what we usually call "epistemic" language may also encompass a variety of non-epistemic uses related to worldly indeterminacy (objective chance; see Lassiter 2018 for discussion).

Combining elements of the theoretical orientations of the last two papers discussed, Løhre & Teigen (2015) compared Norwegian students' production and interpretation of numerical probabilities with the two kinds of probability expressions. In their third experiment, two items were associated with objective factors (*Det er _% sikkert* "It is *n*% certain" and *Det er _% sannsynlig* "There is an *n*% probability"), and one item was associated with subjective judgments (*Jeg er _% sikker* "I am *n*% certain"). They found that participants' assignments of numbers were about 10% higher in the subjective frame than in either of the objective frames, which were indistinguishable. This pattern mainly reveals features of the psychological differences between the two kinds of probability judgment (cf. Lagnado & Sloman 2004). However, it could in principle be useful for linguists: if the effect holds generally, purported differences in subjective vs. objective interpretations of epistemic language should be associated with differences in confidence and assignments of numerical probabilities.

The psychology of reasoning provides another example of productive theoretical exchange made available by the probabilistic perspective on epistemic language. Rotello & Heit (2009); Heit & Rotello (2010) considered the effect of the epistemic items *plausible* and *necessary* on experimental participants' willingness to endorse logically valid syllogisms like (6) and logically contingent syllogisms like (5).

(4) Logically valid:
   a. Horses have property X.
   b. Cats have property X.
   c. So, horses have property X.

(5) Logically contingent:
   a. Dogs have property X.
   b. Cats have property X.
   c. So, horses have property X.

For each kind of argument, they asked participants either to judge whether the conclusion was "necessary" given that the premises are true, or (in another group of participants) whether the conclusion was "plausible" if the premises are true. Heit and Rotello found that the group judging plausibility tended to endorse both kinds of arguments, with only a small advantage for logically valid arguments. However, the "necessary" group tended to endorse most logically valid arguments while rejecting most contingent arguments. They interpreted this pattern in light of a theoretical account of the experimental task that treats it as a signal detection problem (Macmillan & Creelman 2005). Their key finding was that a psychological theory of the task on which the strength of arguments varies only along one dimension—say, the conditional probability of the conclusion given the premises—-could not, from a signal detection perspective, account for the difference among groups. A two-dimensional signal detection model gave a better account, a result which they interpreted as support for dual-process theories of reasoning that distinguish fast associative reasoning from slow deliberative thinking (Sloman 1996; Evans 2008; Kahneman 2011).

In response, Lassiter & Goodman (2017) argued that careful attention to linguistic semantics and pragmatics is necessary to understand the way that syllogistic reasoning is affected by the presence of an epistemic expression in the conclusion. In their experiment, participants were asked to "accept" or "reject" arguments like (6) with epistemic items in the conclusion. Each participant saw arguments with a variety of epistemic items.

(6) a. Dogs have sesamoid bones.
b. Cats have sesamoid bones.
c. So, it is possible/likely/necessary/... that horses have sesamoid bones.

The properties used in the arguments were always obscure biological or pseudo-biological features. By varying the animals in the premises and their relation to the animal in the conclusion, Lassiter & Goodman produced argument skeletons with a wide range of intuitive plausibility, from very low (e.g, "Seals ... dolphins ... So, horses ...") to very high (e.g., "Mammals ... So, horses ...). They then examined how the acceptance rates of a fixed argument skeleton was affected by the presence and identity of an epistemic expression, chosen from {*possible, plausible, likely, probable, certain, necessary*}.

Comparing quantitative response patterns across modals and argument skeletons allowed Lassiter & Goodman to verify a number of qualitative and quantitative predictions of their "Probability Threshold Model", which incorporates a number of ideas from formal semantics and pragmatics while associating argument strength with a one-dimensional scale—the conditional probability of the conclusion given the premises (Heit 1998; Oaksford & Chater 2007). They argued that this model was able to account for the phenomena noted by Heit and Rotello without assuming that the choice of modal item induced participants to engage in different modes of reasoning. If so, the effects of the choice between "plausible" and "necessary" in reasoning experiments do not provide compelling support for a two-dimensional (dual-process) theory of reasoning over a simpler one-dimensional theory based on conditional probability.

The results were broadly consistent with their "Probability Threshold Model", which combines three main components from semantic, pragmatic, and psychological theory: (i) a probabilistic theory of epistemic language as sketched above, (ii) a probabilistic theory of vagueness (e.g., Edgington 1997; Frazee & Beaver 2010; Lassiter 2011; Lassiter & Goodman 2017; Sutton 2017), and (iii) a Bayesian perspective on reasoning that identifies argument strength with the conditional probability of the conclusion given the premises (Heit 1998; Oaksford & Chater 2007). In this account, logically valid arguments are an extreme case where the probability of the conclusion given the premises is 1 since the premises jointly entail the conclusion.

These are just a few illustrations of the potential for productive theoretical exchange between formal semantics/pragmatics and other areas of cognitive science facilitated by the adoption of probabilistic methods in the study of epistemic language.

4 Conditionals

4.1 Theoretical perspectives

Historically, the first serious attempts to analyze English modal vocabulary in terms of probability theory involved the item *if*. Starting with Adams' (1966; 1975) investigations into the connection between conditional probability and assertibility, this connection has been discussed and debated in a huge body of work (see the excellent surveys in Edgington 1995; Bennett 2003). The key problematic in this literature is to make sense of the intuitive connection between the probabilities of indicative conditionals and conditional probabilities, as defined in section 2 above. As van Fraassen (1976: 273) puts it:

> The English statement of a conditional probability sounds exactly like that of the probability of a conditional. What is the probability that I throw a six if I throw an even number, if not the probability that: if I throw an even number it will be a six?

This connection is sometimes called "Stalnaker's Thesis" after Stalnaker 1970, who first framed it as a semantic proposal:

$$\mathcal{P}(\text{if } A \text{ then } B) = \mathcal{P}(B \mid A).$$

According to this suggestion, whatever proposition an indicative conditional denotes, it should have a probability equal to the conditional probability that its consequent would have on the supposition that its antecedent is true. (Note that Stalnaker's thesis is logically independent of the pragmatic claim that *If A, B* is acceptable only if $\mathcal{P}(B \mid A)$ is high, as suggested by Adams (1975); Oaksford & Chater (2007) and others.)

Stalnaker's suggestion held up the hope that one could use this intuitively obvious connection, together with what is known about probability theory, to make progress on the difficult question of what kinds of propositions conditionals denote. However, Lewis (1976) tempered these hopes considerably by proving—on fairly light assumptions—that no conditional connective could validate Stalnaker's Thesis except in extremely trivial models.[3] So, one of the assumptions must be wrong. A number of candidate solutions have been discussed.

An option that would likely be popular among linguists is to deny that *if* is a sentential connective—a two-place operator on propositions, like *and* and *or*. The most popular theory of con-

3 See Egré & Cozic 2011 for an enlightening presentation of Lewis' proof, which reveals a surprising connection to limits on first-order definability of generalized quantifiers.

ditionals in linguistic semantics is the "restrictor" theory, on which *if* is not a connective, but instead combines with a proposition to restrict the domain of quantification of a modal operator (Kratzer 1991a,b). For instance, *If it's snowing, it's cold* would be analyzed roughly as *MUST[**snow**][**cold**]*, where **snow** restricts a silent epistemic *must*. However, the prospects for validating Stalnaker's thesis within the restrictor theory are dim (Charlow 2016). Advocates of the restrictor theory can, however, deny the data, arguing that the connection highlighted in van Fraassen's quote is a sort of linguistic illusion (cf. Kratzer 1991a). This move is quite plausible within the context of the restrictor theory, but see Edgington 2014 for arguments that it does not fully solve the problem.

A second option is to deny that conditionals denote propositions. Adams' (1966; 1975) influential theory is of this type: on his account, the '$\mathcal{P}$' on the left side of Stalnaker's Thesis should be interpreted not as a true probability but as a measure of assertibility. (See also Edgington 1995; Schulz 2017 for a "suppositional" account with connections both to Adams' approach and to the restrictor theory.) Adams develops a detailed account of how we can reason probabilistically with a mix of conditional and non-conditional premises and conclusions, validating many intuitive patterns of inference as characterizations of rational states of probabilistic belief—all while treating conditionals as imposing constraints on probabilistic belief states, rather than mere propositions. However, Adams' theory is limited in that there is no way to interpret nested conditionals like those in (7) compositionally.

(7) a. If Al agrees to the bet, then I will win a lot if the Bills win the Super Bowl. (right-nested)
  b. If Al will agree to the bet if Mary does, then I need to convince Mary to bet. (left-nested)

Simply put, "$\mathcal{P}$(**I win if the Bills win** | **Al agrees**)" makes no sense on Adams' theory. Conditional probability is a function that requires two propositions as arguments. It yields a value only when there are propositions on both sides of the "|". Since *I win if the Bills win* does not denote a proposition on Adams' theory, the conditional probability that we are being asked to compute is undefined. But nested conditionals are in many cases quite natural—witness (7)—and as a result this restriction severely limits the usefulness of Adams' approach for linguistic analysis.

A third possibility is to suppose that conditionals have context-dependent meanings that depend on the information in the probability measure itself. As van Fraassen (1976) showed, it is possible to vindicate Stalnaker's Thesis in this way (see also Douven & Verbrugge 2013 for discussion and relevant experimental evidence). The specific semantic proposal of van Fraassen has been developed in detail by Stalnaker & Jeffrey (1994); Kaufmann (2005). A nice feature of this approach is that it can be extended to allow for unlimited left- and right-nesting of conditionals (Kaufmann 2009). Some potential downsides are that the denotational semantics associated with this approach is quite complex, and that the theory requires us to allow sentences to take on an infinity of values in $[0,1]$ where the classical truth-values 0 and 1 are merely extreme cases, and only conditional sentences make use of the middle range $(0,1)$.

Much less attention has been paid to the probabilities of counterfactual conditionals, but the intuitive argument is just as clear. Modifying van Fraassen's quote along the following lines makes it no less compelling:

> What is the probability that I would have thrown a six if I had thrown an even number, if not the probability that: if I had thrown an even number it would have

been a six?

Shouldn't we then have a variant of Stalnaker's Thesis for counterfactuals?

$$\mathcal{P}(\text{if were } A, \text{ would } B) = \mathcal{P}_{\mathbf{CF}}(B \mid A).$$

I'm intending "$\mathcal{P}_{\mathbf{CF}}(B \mid A)$" as a kind of "counterfactual probability"—the probability that B would have had if A were the case. A possible reason why this connection has not been explored in anything like the depth of the indicative case is that is that there is no standard definition of counterfactual probability. One possibility is that we find the probability that I would have thrown a six, if I had thrown even, by considering a "retrospective" conditional probability: the value that $\mathcal{P}(\mathbf{six} \mid \mathbf{even})$ had at some time before the throw happened (e.g., Edgington 1995; Leitgeb 2012). This would give us the right result for the die-throwing example, and it is an ordinary conditional probability. However, there are some clear counter-examples to this analysis involving events that occurred *after* the antecedent time, but which remain fixed in the counterfactual scenario (Barker 1998; Edgington 2003). For example, suppose I bet "Heads" on the flip of a fair coin and then the coin is flipped and comes up tails. Since the coin flip happened *after* the bet, the retrospective theory does not take into account the outcome of the coin flip: thus it predicts that both sentences in (8) should have probability .5.

(8) a. If I'd bet "heads", I'd have won.
  b. If I'd bet "tails", I'd have won.

But this is clearly incorrect: since the coin came up tails, the probability that I'd have won if I'd bet heads is 0, and the probability that I'd have won if I'd bet tails is 1. (Examples of this type are sometimes called "Morgenbesser cases" after Sidney Morgenbesser, who Slote (1978) credits with an example that can be used to make the point.)

While the empirical picture is very complex, it appears that the choice of which facts to hold fixed and which to ignore in counterfactual reasoning is primarily sensitive to causal relevance rather than temporal order (Barker 1998; Kaufmann 2001; Edgington 2003, 2008). This means that a robust theory of causality may be a prerequisite for developing an adequate theory of counterfactual probability, and indeed counterfactual reasoning in general. Causal Bayesian Networks, as developed in a substantial body of work in computer science, philosophy, and psychology, are an obvious starting point for developing a formally explicit and empirically adequate theory of counterfactual probability (Meek & Glymour 1994; Pearl 2000; Sloman 2005; Lucas & Kemp 2015; see Schulz 2011; Kaufmann 2013; Lassiter 2017b for relevant discussion from a linguistic perspective). See also Schulz 2017 for a recent and highly recommended account of the probabilities of counterfactuals.

Another consideration relevant to the theoretical dialectic here is the interpretation of probabilistic language that occurs within conditionals. For instance, consider the following intuitively true statements about a scenario where we know that a fair die has been rolled, but do not know how it came up.

(9) a. There's a 1/6 chance that the die came up 6.
  b. If the die came up even, there's a 1/3 chance that it came up 6.

Obviously, we would not be licensed in concluding from these premises that

(10) The die didn't come up even.

But then, as Yalcin (2012b) points out, we have a counter-example to the logical principle of Modus Tollens, according to which *If A then B* and *not-B* should together imply *not-A* . Let *A* be *The die came up even* and *B* be *There is a 1/3 chance that the die came up 6*. Since (9a) entails that *B* is false—there is not a 1/3 chance that the die came up 6—we should be able to infer by Modus Tollens that *A* is false—the die did not come up even. But this is clearly not an appropriate inference when we have no idea at all how the die came up!

This kind of example can be used to motivate a different connection between conditionals and Bayesian inference: when a probabilistic expression like *1/3 chance* occurs in the consequent of a conditional, its interpretation is tied to a certain conditional probability—the probability generated by conditioning on the truth of the antecedent. On this account, the premises of (9) are interpreted as

(11) a. $\mathcal{P}(\textbf{die comes up 6}) = 1/6$

b. $\mathcal{P}(\textbf{die comes up 6} \mid \textbf{die comes up even}) = 1/3$

Clearly, nothing follows about whether or not the die came up even from these premises. In effect, the consequent of (9b) is interpreted not as *There's a 1/3 chance ...*, but as *There's a 1/3 chance*$_{\textbf{\textit{even}}}$, where the interpretation of *chance* has been modified to exclude the possibility that the die did not come up even. This explains the sense that the premises in (9) are talking about different things: one instance of *chance* takes into account both even and odd possibilities, while the other excludes the odd.

This kind of interaction between conditionals and probabilistic language is expected on either a restrictor (Kratzer 1991a,b; Yalcin 2012a; Lassiter 2017b) or a suppositional (Edgington 1995; Schulz 2017) theory of conditionals, both of which indicate in various ways that a conditional is evaluated by temporarily adding the information that the antecedent is true. This result can also be encoded within the traditional style of syntactic and semantic analysis of conditionals where *if* is a binary function on propositions (a connective), but doing so may require these theories to take on aspects of the restrictor and/or suppositional theories.

4.2 Experimental and quantitative work

There is much more experimental work on the probabilistic analysis of conditionals than on epistemic items like *certain* and *might*. The question of how people reason with indicative conditionals has been treated in an extensive body of work in experimental psychology, some of which engages deeply with probabilistic analyses drawn from the philosophical literature (e.g., Evans & Over 2004).

Much work on indicative conditionals has assumed the material conditional analysis of indicative conditionals, where *If A then B* is true whenever *A* is false or *A* and *B* are both true. On this analysis, $\mathcal{P}(\textit{If A then B})$ should be equal to $\mathcal{P}(\bar{A}) + P(A \cap B)$. In many cases, this value will be greater than the value $\mathcal{P}(B \mid A)$ that we expect by Stalnaker's thesis. Recycling an earlier example, the probability that I throw a 6 if I throw an even number is obviously $1/3$. But the material conditional analysis suggests the value $2/3$—the sum of $\mathcal{P}(\overline{\textbf{even}}) = 1/2$ and $\mathcal{P}(\textbf{even} \cap \textbf{six}) = 1/6$. While it may be possible in some cases to mount a pragmatic defense of this odd prediction (cf. Lewis 1976), many philosophers and psychologists have taken this and similar discrepancies to tell decisively against the material conditional analysis.

Moving beyond intuition, Evans, Handley & Over (2003) and Oberauer & Wilhelm (2003) (working independently) pitted Stalnaker's thesis against the predictions of the material conditional analysis in a series of experiments, using problems structurally similar to the die-throwing example just discussed. Both sets of researchers found results that tell decisively against the predictions of the material conditional analysis, and both found an interesting pattern of individual differences. For instance, in one of Evans et al.'s (2003) experiments, half of participants conformed with Stalnaker's thesis, but nearly as many (43%) appeared to use the probability of the *conjunction* of antecedent and consequent! In our die-throwing example, this would correspond to judging the probability of *If the die is even it is 6* to be equal to *The die is even and 6*, which is equivalent to *The die is 6*—probability $1/6$. In a follow-up experiment, Evans, Handley, Neilens & Over (2007) investigated the source of these differences. They found that participants who scored poorly on a test of general intelligence were more likely to use the conjunctive interpretation. This suggests that the striking pattern of conjunctive responses may not reflect genuine linguistic variation, but may rather be an artifact of the experimental method due to some participants' poor understanding of the questions or background scenarios (though see Girotto & Johnson-Laird 2004; Oberauer, Geiger, Fischer & Weidenfeld 2007 for interpretations in terms of reasoning strategies cashed out in Johnson-Laird's (1983) Mental Models framework).

A variety of further experiments support Stalnaker's Thesis: in many cases, people will judge the probability of an indicative conditional to be equal to the corresponding conditional probability (e.g., Over & Evans 2003; Over, Hadjichristidis, Evans, Handley & Sloman 2007). Evans & Over (2004) survey the relevant experimental findings up to 2004 and use them to argue for the suppositional theory of conditionals advocated by Edgington (1995). On this account, we evaluate the probability of a conditional by temporarily assuming that the antecedent is true and considering the probability of the consequent in this light. Note that this sort of account focuses on the reasoning potential of conditionals rather than issues of truth and falsity; indeed Edgington (1995) argues extensively in favor of a non-propositional theory of conditionals. While the experimental evidence in favor of Stalnaker's thesis is quite strong, there are several problem cases in the semantic literature that have not yet been subject to experimental investigation: see Kaufmann 2004; Khoo 2016; Moss 2018.

Now consider an intuitively plausible generalization of Stalnaker's thesis involving right-nested conditionals. On face, it would seem that a sentence of the form *If A then (C if B)* should be equivalent to one of the form *If A and B then C*, and that they should have the same probabilities. For instance, the probability of *If the die is even, then it's a 2 if it's not a 4* seems to be $1/2$, since an even die that isn't a 2 is either a 4 or a 6, with equal probabilities. Similar reasoning yields the value $1/2$ for *If the die is even and it's not a 4, it's a 2*. This suggests a generalization of Stalnaker's thesis:

$$\mathcal{P}(\text{If } A\text{, then } C \text{ if } B) = \mathcal{P}(\text{If } A \text{ and } B\text{, then } C)$$

The generalized principle makes sense, given that—in the probability calculus—conditioning sequentially on A and B is equivalent to conditioning on their conjunction. However, Douven & Verbrugge (2013) provide experimental evidence that may problematize the generalized thesis.

Participants were asked to repond on a 1-to-7 scale, from "highly improbable" to "highly probable", to prompts with this form (but slightly more complex):

(12) a. Suppose it rains tomorrow. How probable is the following outcome? "My boots will get muddy if I go outside." ["probability of conditional"]

b. Suppose it rains tomorrow and I go outside. How probable is the following outcome? "My boots will get muddy." ["conditional probability"]

They found that the "conditional probability" condition was associated with a small, but highly statistically significant, increase in probability ratings (4.94 vs. 4.56). While replication with different and more diverse materials is needed, it is puzzling for advocates of Stalnaker's thesis. One possibility, suggested by Douven & Verbrugge (2013), is that the meaning of a conditional is sensitive to a belief state, as in van Fraassen's (1976) theory mentioned above. If so, we might expect the denotation of the conditional to be affected by whether it is embedded or not: for instance, the probability measure relevant to its interpretation would be conditional on the truth of the antecedent. If so, the usual assumption that *If A, then C if B* is always equivalent to *If A and B, then C* is not correct. (See also Khoo & Mandelkern to appear for further arguments against this assumption.)

A different, deflationary interpretation of Douven & Verbrugge's (2013) results might invoke their use of *Suppose A ...* in place of *If A, ...* in the experimental prompts, assuming that their effects are equivalent. Perhaps there are subtle differences between these items that could explain the differences. In any case, their findings are in need of explanation, and point again both to the importance of experimental work and to the potential for data involving nested conditionals to help us choose among theories.

There is a large body of experimental work on counterfactual reasoning, much of it from social psychology. This field of study largely explores the relationship between counterfactual reasoning and socially important concepts such as explanation, causation, responsibility, blame, and regret, including much pragmatically interesting work about when and why people generate counterfactual explanations. In addition, there is important work in this area on the psychological processes involved in counterfactual reasoning and what biases these processes might induce: see Roese 1997; Byrne 2007; Epstude & Roese 2008 for surveys. This body of research seems to offer many insights for semantics and pragmatics that remain mostly unexplored to date, but is largely beyond the scope of this article since it does not focus on probabilistic topics.

5 Conclusion

The quantitative turn in linguistics and many neighboring fields, and the accompanying attention to phenomena involving uncertainty and gradation, have caused many to wonder whether the logical methods of formal linguistics are really the right kind of tools for the analysis of language. But probability theory is built on classical logic, and it allows us to accommodate gradation and uncertainty without abandoning the highly structured, logically well-behaved models that have proved so productive in the past. The enriched representation of information associated with probabilistic models may also shed new light on the many areas in which uncertainty and informational dynamics are crucially implicated in natural language semantics and pragmatics.

This article began by outlining the semantics of probability and conditional probability, emphasizing the continuity of these ideas with familiar models of information and belief update as well as empirical motivations for the enriched representation that probability theory provides. We then surveyed some applications of probabilistic ideas to two areas of lexical semantics and pragmatics—epistemic vocabulary and *if*—with attention to both theoretical and experimental work that has been inspired by this perspective. The probabilistic perspective on the language of uncertainty has already generated many new theoretical and empirical insights, and it promises many more as we explore new interdisciplinary connections with psychology, philosophy, and beyond.

References

Adams, Ernest W. 1966. Probability and the logic of conditionals. *Studies in Logic and the Foundations of Mathematics* 43. 265–316.

Adams, Ernest W. 1975. *The logic of conditionals: An application of probability to deductive logic*. Springer.

Anderson, John R. 2009. *How can the human mind occur in the physical universe?* Oxford University Press.

Barker, Chris. 2009. Clarity and the grammar of skepticism. *Mind and Language* 24(3). 253–273.

Barker, Chris & Gina Taranto. 2003. The paradox of asserting clarity. In *Wecol proceedings*, .

Barker, Stephen. 1998. Predetermination and tense probabilism. *Analysis* 58(4). 290–296.

Bennett, Jonathan F. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press.

Brasoveanu, Adrian & Jakub Dotlacil. To appear. *Formal linguistics and cognitive architecture: Integrating generative grammars, cognitive architectures and bayesian methods*. Springer.

Byrne, Ruth MJ. 2007. *The rational imagination: How people create alternatives to reality*. MIT press.

Charlow, Nate. 2016. Triviality for restrictor conditionals. *Noûs* 50(3). 533–564.

Chater, Nick & Christopher D Manning. 2006. Probabilistic models of language processing and acquisition. *Trends in cognitive sciences* 10(7). 335–344.

Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2013. Inquisitive semantics: a new notion of meaning. *Language and Linguistics Compass* 7(9). 459–476.

Copley, Bridget. 2004. So-called epistemic should. *Snippets* 9. 7–8.

Crocker, Matthew W. 2010. Computational psycholinguistics. In Alexander Clark, Chris Fox & Shalom Lappin (eds.), *Handbook of computational linguistics and natural language processing*, Wiley Blackwell.

Crone, Philip. 2016. Asserting clarity as managing awareness. To appear in *Journal of Semantics*.

Douven, Igor & Sara Verbrugge. 2013. The probabilities of conditionals revisited. *Cognitive Science* 37(4). 711–730.

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

Edgington, Dorothy. 1997. Vagueness by degrees. In Rosanna Keefe & Peter Smith (eds.), *Vagueness: A reader*, 294–316. MIT Press.

Edgington, Dorothy. 2003. Counterfactuals and the benefit of hindsight. In Phil Dowe & Paul

Noordhof (eds.), *Cause and chance: Causation in an indeterministic world*, Routledge.

Edgington, Dorothy. 2008. Counterfactuals. In *Proceedings of the aristotelian society*, vol. 108 1, 1–21.

Edgington, Dorothy. 2014. Indicative conditionals. In Edward N. Zalta (ed.), *The stanford encyclopedia of philosophy*, Metaphysics Research Lab, Stanford University winter 2014 edn.

Egré, Paul & Mikaël Cozic. 2011. If-clauses and probability operators. *Topoi* 30. 17–29.

Epstude, Kai & Neal J Roese. 2008. The functional theory of counterfactual thinking. *Personality and Social Psychology Review* 12(2). 168–192.

Evans, Johathan St. B. T. & David E. Over. 2004. *If.* Oxford University Press.

Evans, Jonathan St. B.T. 2008. Dual-processing accounts of reasoning, judgment, and social cognition. *Psychology* 59(1). 255.

Evans, Jonathan St BT, Simon J Handley, Helen Neilens & David E Over. 2007. Thinking about conditionals: A study of individual differences. *Memory & Cognition* 35(7). 1772–1784.

Evans, Jonathan St B.T., Simon J Handley & David E Over. 2003. Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29(2). 321.

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

Fox, Craig R. & Gülden Ülkümen. 2011. Distinguishing two dimensions of uncertainty. In G. Kirkeboen W. Brun, G. Keren & H. Montgomery (eds.), *Perspectives on thinking, judging, and decision making*, Universitetsforlaget.

van Fraassen, Bas C. 1976. Probabilities of conditionals. In W.L. Harper & C.A. Hooker (eds.), *Foundations of probability theory, statistical inference, and statistical theories of science*, vol. 1, 261–308. Reidel.

Frank, M.C. & N.D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998.

Franke, Michael & Gerhard Jäger. 2016. Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft* 35(4). 3–44.

Frazee, Joey & David Beaver. 2010. Vagueness is rational under uncertainty. In Maria Aloni, Harald Bastiaanse, Tikitu de Jager & Katrin Schulz (eds.), *Logic, language and meaning: 17th amsterdam colloquium, amsterdam, the netherlands, december 16-18, 2009, revised selected papers*, 153–162. Springer.

Girotto, Vittorio & Philip N Johnson-Laird. 2004. The probability of conditionals. *Psychologia* 47(4). 207–225.

Goodman, Noah D. & Daniel Lassiter. 2015. Probabilistic semantics and pragmatics: Uncertainty in language and thought. In Shalom Lappin & Chris Fox (eds.), *Handbook of Contemporary Semantic Theory*, Wiley-Blackwell 2nd edn.

Hawthorne, John. 2004. *Knowledge and lotteries*. Oxford University Press.

Hawthorne, John, Daniel Rothschild & Levi Spectre. 2016. Belief is weak. *Philosophical Studies* 173(5). 1393–1404.

Heit, E. 1998. A Bayesian analysis of some forms of inductive reasoning. *Rational models of cognition* 248–274.

Heit, E. & C.M. Rotello. 2010. Relations between inductive reasoning and deductive reasoning.

Journal of Experimental Psychology: Learning, Memory, and Cognition 36(3). 805.
Johnson-Laird, Philip N. 1983. *Mental models*. Cambridge University Press.
Jurafsky, Dan. 2003. Probabilistic modeling in psycholinguistics: Linguistic comprehension and production. In Rens Bod, Jennifer Hay & Stefanie Jannedy (eds.), *Probabilistic linguistics*, MIT Press.
Kahneman, Daniel. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux.
Kahneman, Daniel & Amos Tversky. 1982. Variants of uncertainty. *Cognition* 11(2). 143–157.
Kaufmann, Stefan. 2001. Tense probabilism properly conceived. In *Thirteenth amsterdam colloquium*, 132–137.
Kaufmann, Stefan. 2004. Conditioning against the grain. *Journal of Philosophical Logic* 33(6). 583–606.
Kaufmann, Stefan. 2005. Conditional predictions: A probabilistic account. *Linguistics and Philosophy* 28(2). 181–231. doi:10.1007/s10988-005-3731-9.
Kaufmann, Stefan. 2009. Conditionals right and left: Probabilities for the whole family. *Journal of Philosophical Logic* 38(1). 1–53. doi:10.1007/s10992-008-9088-0.
Kaufmann, Stefan. 2013. Causal premise semantics. *Cognitive science* 37(6). 1136–1170.
Kennedy, Chris. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1). 1–45.
Khoo, Justin. 2016. Probabilities of conditionals in context. *Linguistics and philosophy* 39(1). 1–43.
Khoo, Justin & Matthew Mandelkern. to appear. Triviality results and the relationship between logical and natural languages. *Mind* .
Klecha, Peter. 2014. *Bridging the Divide: Scalarity and Modality*: University of Chicago dissertation.
Kratzer, Angelika. 1991a. Conditionals. In A. von Stechow & D. Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, 651–656. Walter de Gruyter.
Kratzer, Angelika. 1991b. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, 639–650. Walter de Gruyter.
Kyburg, Henry. 1961. *Probability and the logic of rational belief*. Wesleyan University Press.
Lagnado, David A & Steven A Sloman. 2004. Inside and outside probability judgment. *Blackwell handbook of judgment and decision making* 157–176.
Lassiter, Daniel. 2010. Gradable epistemic modals, probability, and scale structure. In Nan Li & David Lutz (eds.), *Semantics & Linguistic Theory (SALT) 20*, 197–215. CLC Publications.
Lassiter, Daniel. 2011. Measurement and Modality: The Scalar Basis of Modal Semantics. Ph.D. thesis, New York University.
Lassiter, Daniel. 2016. *Must*, knowledge, and (in)directness. *Natural Language Semantics* 24(2). 117–163.
Lassiter, Daniel. 2017a. *Graded Modality*. Oxford University Press.
Lassiter, Daniel. 2017b. Probabilistic language in indicative and counterfactual conditionals. In Dan Burgdorf, Jacob Collard, Sireemas Maspong & Brynhildur Stefánsdóttir (eds.), *Semantics & Linguistic Theory (SALT) 27*, 525–546. CLC Publications.
Lassiter, Daniel. 2018. Talking about (quasi-)higher-order uncertainty. In Cleo Condoravdi & Tracy Holloway King (eds.), *Tokens of meaning: Papers in honor of Lauri Karttunen*, CSLI

Publications.

Lassiter, Daniel & Noah D. Goodman. 2017. Adjectival vagueness in a Bayesian model of interpretation. *Synthese* 194(10). 3801–3836.

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

Leitgeb, Hannes. 2014. The stability theory of belief. To appear in *The Philosophical Review*.

Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 85(3). 297–315. doi:10.2307/2184045.

Lewis, David. 1979. Attitudes *de dicto* and *de se*. *The philosophical review* 88(4). 513–543.

Lin, Hanti & Kevin T Kelly. 2012. Propositional reasoning that tracks probabilistic reasoning. *Journal of philosophical logic* 41(6). 957–981.

Løhre, Erik & Karl Halvor Teigen. 2015. There is a 60% probability, but I am 70% certain: Communicative consequences of external and internal expressions of uncertainty. *Thinking & Reasoning* 1–28.

Lucas, Christopher G. & Charles Kemp. 2015. An improved probabilistic account of counterfactual reasoning. *Psychological Review* 122(4). 700–734.

Macmillan, N.A. & C.D. Creelman. 2005. *Detection theory: A user's guide*. Lawrence Erlbaum.

Meek, Christopher & Clark Glymour. 1994. Conditioning and intervening. *The British journal for the philosophy of science* 45. 1001–1021.

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

Oaksford, Mike & Nick Chater. 2007. *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.

Oberauer, Klaus, Sonja M Geiger, Katrin Fischer & Andrea Weidenfeld. 2007. Two meanings of "if"? individual differences in the interpretation of conditionals. *The Quarterly Journal of Experimental Psychology* 60(6). 790–819.

Oberauer, Klaus & Oliver Wilhelm. 2003. The meaning (s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29(4). 680.

Over, David E & Jonathan St BT Evans. 2003. The probability of conditionals: The psychological evidence. *Mind & Language* 18(4). 340–358.

Over, David E, Constantinos Hadjichristidis, Jonathan St BT Evans, Simon J Handley & Steven A Sloman. 2007. The probability of causal conditionals. *Cognitive psychology* 54(1). 62–97.

Pearl, Judea. 2000. *Causality: Models, reasoning and inference*. Cambridge University Press.

Perfors, A. 2012. Bayesian models of cognition: What's built in after all? *Philosophy Compass* 7(2). 127–138.

Perfors, Amy, Joshua B Tenenbaum, Thomas L Griffiths & Fei Xu. 2011. A tutorial introduction to Bayesian models of cognitive development. *Cognition* 120(3). 302–321.

Pólya, George. 1954. *Mathematics and plausible reasoning, volume 1: Induction and analogy in mathematics*. Princeton University Press.

Portner, Paul. 2009. *Modality*. Oxford University Press.

Roberts, Craige. 2004. Context in dynamic interpretation. In L. Horn & G. Ward (eds.), *Handbook of Pragmatics*, 197–220.

Roese, Neal J. 1997. Counterfactual thinking. *Psychological bulletin* 121(1). 133.
Romeijn, Jan-Willem. 2017. Philosophy of statistics. In Edward N. Zalta (ed.), *The stanford encyclopedia of philosophy*, Metaphysics Research Lab, Stanford University spring 2017 edn.
Rotello, C.M. & E. Heit. 2009. Modeling the effects of argument length and validity on inductive and deductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35(5). 1317–1330.
Schulz, Katrin. 2011. "If you'd wiggled A, then B would've changed": Causality and counterfactual conditionals. *Synthese* 179(2). 239–251.
Schulz, Moritz. 2017. *Counterfactuals and probability*. Oxford University Press.
Sloman, Steven A. 1996. The empirical case for two systems of reasoning. *Psychological bulletin* 119(1). 3.
Sloman, Steven A. 2005. *Causal models: How we think about the world and its alternatives*. OUP.
Slote, Michael A. 1978. Time in counterfactuals. *The Philosophical Review* 87(1). 3–27.
Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Syntax and semantics 9: Pragmatics*, Academic Press.
Stalnaker, Robert & Richard Jeffrey. 1994. Conditionals as random variables. In *Probability and conditionals: Belief revision and rational decision*, 31–46. Cambridge University Press.
Stalnaker, Robert C. 1970. Probability and conditionals. *Philosophy of science* 37(1). 64–80.
von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3(1). 1–77.
Sutton, Peter R. 2017. Probabilistic approaches to vagueness and semantic competency. *Erkenntnis* 1–30.
Swanson, Eric. 2015. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* 1–26.
Tenenbaum, Joshua B., Charles Kemp, Tom L. Griffiths & Noah D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331(6022). 1279–1285.
Ülkümen, Gülden, Craig R Fox & Bertram F Malle. 2015. Two dimensions of subjective uncertainty: Clues from natural language. To appear in *Journal of Experimental Psychology: General*.
Wallsten, Thomas S, David V Budescu, Amnon Rapoport, Rami Zwick & Barbara Forsyth. 1986. Measuring the vague meanings of probability terms. *Journal of Experimental Psychology: General* 115(4). 348.
Williams, Peter M. 1980. Bayesian conditionalisation and the principle of minimum information. *The British Journal for the Philosophy of Science* 31(2). 131–144.
Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
Yalcin, Seth. 2007. Epistemic modals. *Mind* 116(464). 983–1026.
Yalcin, Seth. 2012a. Context probabilism. In M. Aloni, V. Kimmelman, F. Roelofsen, G. W. Sassoon, K. Schulz & M. Westera (eds.), *Logic, language and meaning*, vol. Lecture Notes in Computer Science 7218, 12–21. Springer.
Yalcin, Seth. 2012b. A counterexample to modus tollens. *Journal of Philosophical Logic* 41(6). 1001–1024.
Yalcin, Seth. 2016. Modalities of normality. In Matthew Chrisman & Nate Charlow (eds.), *Deontic modals*, 230–255. Oxford University Press.