

A Dynamic Probabilistic Approach to Epistemic Modality*

Manuel Kríž

University of Vienna

Abstract. This paper develops a very simple probabilistic dynamic formalism for a semantics of epistemic modal expressions that is essentially expressivist in nature, with a notion of information states based on probability spaces. It differs from existing approaches to this task (in particular [22] and [12]) in certain features that make it more appropriate for modelling actual conversational behavior, especially as regards ignorance and unawareness. It will be argued that the latter is important for an adequate treatment of the alleged non-monotonicity of conditionals. Directions for further research that follow from the approach will be pointed out.

Introduction

Recent work ([19], [20], [11]) has emphasized the importance of probability in epistemic modality in natural language. At the same time, there is a tradition in dynamic semantics (e. g. [16], [6], more recently [13]) which is successful in dealing with the subjective, non-truth conditional aspects of their meaning and its non-monotonicity in conversation. In this paper, we make a first attempt at doing the obvious thing and bringing them together. We will explore a possibility that has also been hinted at in [22], and developed quite fully, though with a different focus, in [12]: we will give a semantics for epistemic modals that is both probabilistic and essentially dynamic, using sets of probability functions to model information. In section 1, we will define our simple technical apparatus. In the next section, we will explore some applications to the behavior of epistemic modals in discourse and some linguistically relevant aspects of internal states of speakers and hearers. Section 3 concerns the non-monotonic properties of conditionals, and in the final section, we point to some further limitations of our model and directions for future research that follow from our approach.

1 A Toy Model

Definition 1. A *probability space over possible worlds* is a triple $\langle S, \mathcal{S}, P \rangle$, where S is a set of worlds, \mathcal{S} is some Boolean algebra over S , and P is a probability function on \mathcal{S} .

*My thanks are due to Viola Schmitt, Daniel Büring, and especially Johannes Schmitt for stimulating and encouraging conversations.

Let our language \mathcal{L} be a simple propositional language with modal operators \diamond , \square and \triangle (the latter for *probably*) and a conditional connective \succ . We call a formula *non-modal* if it contains none of these.

Definition 2. A *model* for \mathcal{L} is a pair $\langle W, \llbracket \cdot \rrbracket \rangle$, where W is a set of possible worlds and $\llbracket \cdot \rrbracket$ an interpretation function that assigns to every non-modal formula of the language a subset of W , subject to the usual conditions for \wedge and \vee .

Definition 3. We define an *information state* (or *context*) c to be a pair $\langle \mathcal{S}, \mathcal{P} \rangle$, where \mathcal{S} is a boolean algebra over W and \mathcal{P} is a set of functions P such that $\langle W, \mathcal{S}, P \rangle$ is a probability space.

A preliminary definition of update, which serves to illustrate what properties we want it to have, is given in the following:

Definition 4. (to be revised) If ϕ is a non-modal formula, then the **update** of a context c is defined as follows:¹

$$\begin{aligned} c[\phi] &:= \{P(\cdot|\phi) \mid P \in c \wedge P(\phi) \neq 0\} \\ c[\diamond\phi] &:= \{P \in c \mid P(\phi) > 0\} \\ c[\square\phi] &:= \{P \in c \mid P(\phi) = 1\} \\ c[\triangle\phi] &:= \{P \in c \mid P(\phi) > .5\} \\ c[\phi \succ \psi] &:= \{P \in c \mid \{P\}[\phi][\psi] = \{P\}[\phi]\} \end{aligned}$$

Definition 5. A formula ϕ is **accepted** in an information state c ($c \models \phi$) if $c[\phi] = c \neq \emptyset$. It is **consistent** with c if $c[\phi] \neq \emptyset$ and **inconsistent** if it is not consistent. It is **compatible** with c if there is a sequence Δ (which may include ϕ) of formulae such that successive update with the members of Δ yields a c' in which ϕ is accepted. ϕ is **incompatible** with c if it is not compatible.

The recursive definition of negation in such a system is a tricky matter due to the qualitative difference between conditionalization and eliminative update. It is not even clear that the definition of negation really needs to be given recursively instead of as a set of rules for negative update, as in [13] (cf. also [10]). However, in our case a recursive definition *can* be given by adapting the procedure developed in [12].

We use a slightly simplified version of Schmitt's notion of a Bayesian closure of a set of probability functions. It is the set that contains all the functions in the original set, and all those that are the result of conditionalizing a function in the original set on some non-modal sentence in the language.

Definition 6. The *Bayesian closure* c^{cl} of a context c is defined as

$$\{P \mid \exists P' \in c: \exists \phi \in \mathcal{L}: P(\cdot) = P'(\cdot|\phi)\}.$$

Schmitt proves that the function $(\cdot)^*$ is the reverse of the Bayesian closure function if c fulfills the condition of *weak regularity*: if a function in c assigns 1 to any formula ϕ , then all functions in c do. Our contexts are, in general, not weakly regular, but as will be clear in a moment, that is not a problem for our final definition of update.

¹We write $P(\phi)$ for $P(\llbracket \phi \rrbracket)$. Furthermore, since \mathcal{S} is not affected by updates at this point, we will for ease of exposition pretend that a context is just a set of probability functions.

Definition 7. The $*$ -function returns for every set C of probability functions that is closed under conditionalization in the above sense the unique weakly regular set of probability functions that C is the Bayesian closure of.

$$C^* = \{P \in C \mid \neg \exists P' \in C: \exists \phi \in \mathcal{L}: P(\cdot) = P'(\cdot \mid \phi)\}$$

On the level of Bayesian closures, a recursive definition for negation can be given. These are Schmitt's update rules:

Definition 8. The update of a Bayesian closure C with a formula is defined as follows, where ϕ is a non-modal formula and ψ is an arbitrary formula.

$$\begin{aligned} C \uparrow \phi &:= \{P \in C \mid \exists P' \in C: P(\cdot) = P'(\cdot \mid \phi)\} \\ C \uparrow \neg \psi &:= \{P \in C \mid \exists y \subseteq C: P \in y \wedge y^{cl} \uparrow \psi = \emptyset\} \\ C \uparrow \diamond \phi &:= \{P \in C \mid C \uparrow \neg \phi \neq C\} \\ C \uparrow \square \phi &:= \{P \in C \mid C \uparrow \phi = C\} \\ C \uparrow \Delta \phi &:= \{P \in C \mid \forall P \in C^*: P(\phi) > .5\} \\ C \uparrow (\phi \succ \psi) &:= \{P \in c \mid (C \uparrow \phi) \uparrow \psi = C \uparrow \phi\} \end{aligned}$$

But Bayesian closures are not what correctly represents our information, and furthermore, the update with modals here is only a test. What we will do is to make the update of a context distributive: we take the singleton of every function in P , form the Bayesian closure of it, apply Schmitt's rules to it, feed it to the $*$ -function, and collect all the results.

Definition 9. The update of a context c with an arbitrary formula ϕ (written $c[\phi]$) is given by the following:

$$c[\phi] = \bigcup \{ \mathcal{P} \mid \exists P \in c: \mathcal{P} = (\{P\}^{cl} \uparrow \phi)^* \}$$

The effects of this update are generally as in Definition 4. Now the negation of a formula ϕ is accepted in c iff ϕ is inconsistent with c .²

Norm of Assertion. A speaker S may assert ϕ iff ϕ is accepted in their information state c_S .

Norm of Contradiction. A hearer H may contradict a speaker's assertion of ϕ iff ϕ is incompatible with their information state c_H .

2 Some Applications

2.1 Question Sensitivity

These definitions allow us to capture elegantly a number of linguistic phenomena. The first (perhaps more of a cognitive phenomenon) is what is called *question sensitivity*

²We cannot, however, derive the fact, pointed out in [14], that the negation of a conditional is actually the negation of the consequent. Neither are the neg-raising properties of *probable* explained. This is a defect that the present approach shares with all comparable ones that I am aware of.

in [21]. There are people who don't know about the existence of a town by the name of Topeka. It seems very plausible to say that they are entirely *insensitive* to the question of whether it's raining in Topeka. The possibility of encoding this falls out directly from our use of probability spaces. Nothing forces \mathcal{S} to be the whole power set of W , so there may well be propositions that the probability functions in \mathcal{P} are not defined on.³ Note that this means that sensitivity is not closed under logical consequence, although due to our use of proper probability spaces over possible worlds, we do, of course, still incur the problem of logical omniscience and, as it were, logical omnisensitivity.

2.2 Ignorance about Possibilities

Using sets of probability functions is the standard way of capturing that probability assignments are typically vague (cf. e. g. [7], among many others): someone who says that it's probably raining doesn't have to have a specific probability, say, .7, in mind.⁴ But there is an interesting special case: writing $\mathcal{P}(\phi)$ for the set of values assigned to ϕ by some $P \in \mathcal{P}$, we can consider the possibility that $\mathcal{P}(\phi) = [0, 1]$. An agent with such an information state seems strange at first: she is sensitive to ϕ , and she doesn't have any evidence to exclude it. Still, she doesn't believe that ϕ is possible. But we can use this to make sense of certain acts of communication that Yalcin ([21]) subsumes under question sensitivity, but which we would like to distinguish from for reasons that will become clear later (cf. footnote 10). For instance, in well-known examples from [1], a person who is anxiously awaiting the results of John's cancer test may be disposed to say "I don't know whether John might have cancer, we're still waiting for the test results." Of course, John's having cancer is consistent and compatible with her information state, and she has considered the question, which means she is sensitive to it. In fact, the person could just as well say "Yes, John might have cancer, that's why they're running a test." But by saying she doesn't know whether John might have cancer, she portrays herself as accepting neither that he might nor that he might not have cancer; in our terms, this means that she presents herself as being in an information state where the range of probabilities assigned to John's having cancer is $[0, 1]$ (or $[0, 1)$). Similarly, a speaker who assigns probabilities $(0, 1]$ to ϕ is disposed to say that she doesn't know how probable ϕ is (at least if the interval is dense).

It could be suggested that possibility operators embedded under *know* have no semantic effect (at least at the level of at-issue content), so that an agent who says she doesn't know if something is possible is just expressing that she is ignorant as to whether it does in fact obtain. However, evidence for an analysis along the lines presented above comes from the fact that while the agent may, in a given situation, either say that something is possible or that she doesn't know if it's possible, she can't say both at the same time; she has to decide how to present herself:

³A reviewer suggested that this might just be a case of a familiarity presupposition of a proper name. But that is not an *alternative* explanation: the fact that his probability function is undefined for sentences involving the name is just an effect of the agent's being unfamiliar with it! In addition, it is a well-known fact from the literature of belief that it does often not seem appropriate to ascribe beliefs about a matter that the agent is simply not thinking of, even if she would have no trouble forming them were the proposition to come to his mind. See section 3 for some further discussion of this.

⁴Naturally, there is also higher-order vagueness, which will be ignored.

- (1) ??John might have cancer, so they ran a test to rule that out. I haven't seen the results yet, so I don't know if he might have cancer.

Our analysis is similar in spirit to the treatment that the phenomenon receives in [17]. However, Willer, using a non-probabilistic system, has to introduce special-purpose machinery: for him, information states are sets of sets of worlds, i. e. sets of Stalnakerian contexts, which is not something that is readily interpretable. In contrast, in our probabilistic framework, we have the requisite technical apparatus in place because it also models the vagueness of probability. The only question that remains is what it means for an agent to assign to a proposition a vague probability represented by an interval that include both extreme (0 or 1) and intermediate values. Isn't there simply a fact of the matter about whether or not their evidence excludes the possibility? Yes, there is; and if pushed, the agent will probably agree that the proposition under discussion is a possibility after all. But by presenting herself, for the purpose of the conversation, as assigning to ϕ (e. g. that John has cancer) such a weird set of probability values, she expresses a disposition: should an issue come up the resolution of which depends on whether ϕ is possible or not (e. g. the question of whether certain precautions should be taken), the agent would not just assume the possibility, but rather try to get hold of the lab report before deciding anything. Only if she could not obtain the lab report would she resort to assuming that John's having cancer is, after all, possible, and proceed to action directly.⁵ The crucial thing here is that an agent's information state only encodes her dispositions (or the dispositions she wants to portray herself as having) at a given time, and it may change not only as a result of utterances of others, but also be manipulated by processes *internal* to the agent. It is, however, very unclear to what extent such internal processes are amenable to logical modeling; they certainly aren't at *this* point.

2.3 Contradiction and the Search for Evidence

The reader may have been puzzled by the norm of contradiction given above; the usual way to state it is, of course, to say that a hearer may (or even must) contradict an assertion of ϕ if ϕ is inconsistent with their information state. There are, however, some cases where a hearer can neither accept an assertion nor contradict it. Assume that A assign's to John's having been there probabilities in the interval [.6, .9].

- (2) A: Smith is probably the murderer.
 B: He **MUST** be.
 A¹: Why?
 A²: ??I see.
 A³: ??No.

We can make sense of this if we assume that A is in an information state where the range of probabilities assigned to John's having been there (which we abbreviate as χ) is, say, [.6, .8]. If he updated with $\Box\chi$, he would end up with the absurd context, so he cannot just accept B's assertion. On the other hand, he cannot contradict it either, because his information state is, presumably, not incompatible with $\Box\chi$: from

⁵Note that in order for this to work technically, the person's probability estimate about the contents of the lab report also have to be totally vague. But that seems about right: she is not disposed to speculate about the lab report. Rather, if questions about its content come up, she would try to settle them by having a look at it.

an information state that accepts the negation of a modal formula, one can, in the usual case, reach one that accepts its unnegated version through update with certain non-modal formulae, i. e. conditionalization on the right kind of evidence. That is why, in the most typical case, the only route open to A is to *ask* for such evidence.

As predicted, unmodalized assertion behaves differently:

- (3) A: John was probably there.
 B: Yes, he was there.
 A¹: I see.

To be sure, there are cases that can be updated with $\Box\phi$ directly, viz. if the hearer's information state contains a P such that $P(\phi) = 1$. In that case, we would expect her to just accept the assertion. Indeed, if a hearer is very weakly opinionated about ϕ , or thinks it possible that ϕ follows from what she knows, then such a reaction seems appropriate intuitively. Only when the hearer does have an opinion about the probability of ϕ , the question for evidence is necessary.

Of course, the classic kind of examples of “subjective” use of epistemic modals are covered by our theory as well, as well as the usual asymmetries.

- (4) A: Your keys might be in the car.
 B: No, they can't be, I still had them when we came into the house.
 B': #Okay, but I know that they're not there.

(after [3])

Here, speaker A asserts a possibility on the basis of her own information state, and speaker B, following the norm of rejection, denies it. Her information state assigns probability 0 to the keys' being in the car, so she cannot update in any way to make A's assertion accepted in his information state, thus fulfilling the norm of rejection. What she cannot do is accept A's assertion and point out that she, on the other hand, knows better — such a reaction is just not sanctioned by the norms of conversation.

- (5) a. A: John might be/is probably in the garden. B: He's not.
 b. #A: John isn't in the garden. B: He might be.
 etc.

Our system also includes the explanation from [6] for why the sequence in (5a) is infelicitous when uttered by a single speaker: it is not possible to be in an information state that accepts both of these statements, so the norm of assertion would be violated.⁶ But with two speakers involved, there is no problem, because A can update with the information from B and end up in a state where she no longer accepts his initial assertion.

In (5b), this is not an option. Of course, (5b) is not an impossible sequence in actual discourse, but the point is that B is prompting A to *revise* her information state — by removing the information that John isn't in the garden — and not merely to update.

⁶Contra Yalcin ([22]), who considers sequences like (5a) dubious without taking into account the question of whether they are uttered by one or two speakers. That is why Yalcin is hesitant about recommending a system like the present one, where the update with non-modal formulae is non-eliminative, although his formalism could accommodate that in a way similar to our own.

2.4 A Word on Worlds

The technical setup we have seen might look very similar to that of [22]. However, we believe it is only as similar as any system that deals with the same phenomena will be. The crucial difference is that for Yalcin, information states are not sets of probability functions, but sets of pairs $\langle s, P \rangle$, where s is a set of worlds, the “live possibilities”, and P is a probability function so that $P(s) = 1$. For him, the only expression that properly uses a probability function is *probably*; possibility and necessity modals merely quantify over the elements of s , and update with a non-modal formula restricts s (with concomitant conditionalization of P). By having suitable variation within a context, it is possible to reproduce the total vagueness we used to model a person who claimed not to know if something is possible. But insensitivity is now lost on us; or rather, it is restricted to *probable*. There is, however, no way that $\diamond\phi$ or $\Box\phi$ could be undefined in Yalcin’s system.

Yalcin’s reason for sticking with a quantificational semantics for modals, is given in [18]: there are events that have, as a mathematical fact, probability 0 without being impossible: if we pick a number from the interval $[0, 1]$ completely at random, then the probability of that number being any particular number is 0; still, it is, by hypothesis, possible for every number to be picked. We have gone with [11] in identifying possibility with non-zero probability. Lassiter’s reply to Yalcin’s objection is that in natural language, there is always granularity. If we want to assess the probability of the randomly picked number being .5, they do not think of it as absolutely precise; there will always be some numbers that are indistinguishable from it. We can, as it were, be arbitrarily, but not infinitely, precise. Therefore, natural language statements are never about a continuous sample space in the intended way — only about arbitrarily precise sample spaces. We are sympathetic to this point of view; but we may note that, as a last resort, we would even be prepared to postulate infinitesimal probabilities in order to harvest the benefits of doing away with quantification over worlds. Primarily, this allows us the notion of insensitivity we have developed, which will be put to further use in the next section. A secondary reason to disfavor possible worlds is that they may eventually get in the way of avoiding logical omniscience (cf. [4], where probabilities are assigned directly to sentences).

3 Conditionals and Sensitivity

There is a well-known kind of conjunctions of conditionals which are not reversible, sometimes called *Sobel sequences*. The classic example is this:

- (6) If the USA were to throw its nukes into the sea tomorrow, there would be war.
Of course, if the USA and all the other superpowers were to throw their nukes into the sea tomorrow there would be peace.
- (7) #If the USA and all the other superpowers were to throw their nukes into the sea tomorrow, there would be peace. Of course, if the USA were to throw its nukes into the sea tomorrow, there would be war.

Sobel sequences are generally discussed in the context of counterfactual conditionals, but they are just as possible with epistemic conditionals as with counterfactuals.⁷ In

⁷The existence of indicative Sobel sequences is in fact acknowledged in [5].

a few thousand years, a historian researching the history of the twentieth century, and so oblivious of human nature that she actually considers the possibility of any country throwing its weapons into the sea, could utter:

- (8) If the USA threw its nukes into the sea, there was war. Of course, if the USA and all the other superpowers threw their nukes into the sea, there was peace.
- (9) #If the USA and all the other superpowers threw their nukes into the sea, there was peace. Of course, if the USA threw its nukes into the sea, there was war.

Intuitively, it is clear what happens here: at first, the speaker didn't even consider the possibility that all superpowers could throw their weapons into the sea. Then it comes to her attention and she corrects herself.

The traditional analysis (e. g. [2], [5]) is formulated in terms of domain of quantification: the domain of quantification for a conditional is not the whole set of possible worlds, but a certain subset not including any possibilities that one considers too remote or just forgets to think about, and while the domain can be widened in the course of a conversation, it cannot (at least not without significantly more effort) be narrowed, hence the irreversibility.⁸

Theories based on domain widening have a problem: they cannot distinguish between disregarding the possibility of ϕ (i. e. failure to consider it) and accepting that ϕ is impossible. Assume that in a context⁹ c_1 , a speaker has accepted *if ϕ , ψ* . At the same time, she is disposed to accept *if ϕ and χ , then $\neg\psi$* (which is the situation in Sobel sequences), update with which would put her in the information state c_2 . On the domain restriction theory, it follows that in c_1 , no χ -worlds are in the relevant domain. But then in c_1 , $\neg\Diamond\chi$ is accepted, and so is *if ϕ , $\neg\chi$* . But this is not what we want: a speaker who has failed to consider a possibility is not necessarily disposed to accept its negation.

Conversely, domain restriction theories predict that propositions deemed to be impossible can still feature as the antecedent of a subsequent conditional. But sequences like the following have a strong contradictory flavor:

- (10) #The USA can't have thrown their weapons into the sea, because human nature wouldn't have allowed them to; of course, if they did, there was peace.

It suggests itself to apply the ideas in [8] to this, where pragmatic halos are assumed for domains of modal quantification. We do not doubt that the empirical connection between conditionals and imprecision is correct (cf. also [9]), but such an analysis is only as good as the pragmatic halos framework, which has trouble incorporating negation and providing an explanation for the directionality effects.

In our approach, there is a notion of insensitivity that captures the fact that disregarding a possibility is to disregard its negation and various combinations of it with

⁸Schulz's ([13]) analysis derives the non-monotonicity of counterfactual conditionals in a different way and fails to predict this dynamic, directional effect.

⁹Of course, in these theories, a context is something different than in ours.

others as well.¹⁰ While we cannot analyze complete Sobel sequence, we believe that the manner of failure here is very interesting.

Our update rules do not change \mathcal{S} , which means that we cannot model *becoming* sensitive to a question in the course of a conversation. This is obviously inadequate. In fact, when a question is mentioned in a conversation that we weren't sensitive to, we adjust our information state often (though not always) silently. In our formalization, this would mean that if an agent is insensitive to ϕ and ϕ (or some modalized version of it) is mentioned, they extend \mathcal{S} to the smallest superset of itself that contains $\llbracket\phi\rrbracket$ and still fulfills all the closure conditions that come with being part of a probability space. Of course, the probability functions in \mathcal{P} would have to be altered so as to assign some value to ϕ . This could be done by just *extending* them (which will result in some constraints on the probability of the new ϕ), but this is quite obviously inadequate: upon becoming aware of a possibility, we sometimes *lower* the probability we assign to other possibilities. This is also what is needed in Sobel sequences. For assume that in $c, \models \phi \succ \psi$ is accepted and ϕ and ψ are non-modal formulae, so that $\mathcal{P}(\psi|\phi) = 1$. As long we are not allowed to alter existing probability assignments, that will not change no matter what probability we assign to the new possibility χ , since $\mathcal{P}(\phi|\phi \wedge \chi) = 1$, and so $\mathcal{P}(\psi|\phi \wedge \chi)$. But the conditional that prompted us to become sensitive to χ was $(\phi \wedge \chi) \succ \psi$, with which our context is then inconsistent.

What is missing is a representation of the background knowledge on the basis of which we alter probability assignments when we become aware of an additional possibility; presumably the kind of information that is encoded in *generic* conditionals.¹¹ The modeling of this is a challenging task: one may suspect with [10] that this is actually an issue of activation in the associative memory of the speaker/hearer.

4 Further Limitations and Outlook

There are other issues we cannot treat with our simplistic approach; in particular, the evidential effect of modals. While $\Box\phi$ and ϕ have different update effects and deniability conditions, they still have identical acceptance conditions, and so should be assertable by an agent in the same circumstances (and be interchangeable under attitude operators and in the consequent of conditionals). However, as noted in [15], this is not true:

- (11) Context: *Rain is falling outside and the agent is looking out of the window.*
- a. It's raining outside.
 - b. #It must be raining outside.

Evidential effects are equally present under attitude operators:

- (12) a. Peter believes that aliens were here.

¹⁰This is the reason why we want to distinguish insensitivity to ϕ from $\mathcal{P}(\phi) = [0, 1]$. The first conjunct of a Sobel sequence behaves, in a way, as if the probability of the disregarded possibility were 0; this couldn't be captured if insensitivity were just total vagueness of the probability. On our conception, however, insensitivity does not survive even mention of the possibility (and neither do conditionals survive the mention of an exception!), which makes it unfit for the explanation of the DeRose kind of examples.

¹¹Of course, one could just use an ordered set of contexts with successively larger event spaces \mathcal{S} , but that would be entirely uninteresting.

- b. Peter believes that aliens must have been here.

(12b) clearly implies that Peter's belief about the presence of aliens is inferred, and that he neither actually saw, nor, indeed, hallucinated them. (12a) is neutral about either possibility. That $\Box\phi$ and ϕ are therefore not in complementary distribution either is what makes the incorporation of the evidential impact of modals very non-trivial. One can of course extend information states to include a special body of directly evidenced propositions, but the intricate interplay of implicatures and presuppositions with respect to these will then have to be explored. This we leave to future research as well.

In summary, we have demonstrated that the probabilistic dynamic approach has a number of useful features. Modeling an information state as a vague probability space over possible worlds, we have provided a conspicuous picture of the workings of modals in conversation regarding acceptance, rejection and to some extent the information conferred by them, and makes use of the formal possibility of a totally vague probability distribution to model an agent who presents herself having no commitment at all as to the probability of a proposition. Furthermore, a superior account of question sensitivity falls out directly from the formal tools that are employed.

We have reached the limits of our simple model at two points, giving rise to questions for further research. First, the evidential effect of modals is not captured, which points to the necessity of extending information states to designate some propositions of special status so that evidential presuppositions may be formulated. Second, while we have seen that our notion of insensitivity is in fact superior to that employed in standard accounts of Sobel sequences of conditionals, we have not ourselves provided an analysis of these. It turned out that in order to treat the dynamics of sensitivity that are involved, a representation of generic probabilistic knowledge is needed, and it stands to reason that this ultimately leads us deep into cognitive science.

References

1. DeRose, K.: Epistemic possibilities. *The Philosophical Review* 100, 581–605 (1991)
2. von Stechow, P.: Counterfactuals in a dynamic context. In: Kenstowicz, M. (ed.) *Ken Hale: A Life in Language*. MIT Press, Cambridge (2001)
3. von Stechow, P., Gillies, A.: 'might' made right. In: Weatherson, B., Egan, A. (eds.) *Epistemic Modality*. Oxford University Press (2011)
4. Gaifman, H.: Reasoning with limited resources and assigning probabilities to arithmetical statements. *Synthese* 140, 97–119 (2004)
5. Gillies, A.: Counterfactual scorekeeping. *Linguistics and Philosophy* 30, 329–360 (2007)
6. Groenendijk, J., Stokhof, M., Veltman, F.: Coreference and modality in multi-speaker discourse. In: Kamp, H., Partee, B. (eds.) *Context Dependence in the Analysis of Linguistic Meaning*. I. M. S., Stuttgart (1997)
7. Joyce, J.M.: How probabilities reflect evidence. *Philosophical Perspectives* 19, 153–178 (2005)
8. Klecha, P.: Shifting modal domains: An imprecision-based account. Paper presented at LSA 2012 (2012)
9. Kriz, M., Schmitt, V.: Semantic slack in plural predication. Talk at the Semantics Workshop, Vienna (May 2012)
10. van Lambalgen, M., Stenning, K.: *Human Reasoning and Cognitive Science*. MIT Press (2008)

11. Lassiter, D.: Measurement and Modality: The Scalar Basis of Modal Semantics. Ph.D. thesis, NYU (2011)
12. Schmitt, J.: Iffy Confidence. Ph.D. thesis, USC (2012)
13. Schulz, K.: Minimal Models in Semantics and Pragmatics: Free Choice, Exhaustivity, and Conditionals. Ph.D. thesis, University of Amsterdam (2007)
14. Stalnaker, R.: A defense of conditional excluded middle. In: Stalnaker, R., Harper, W., Pearce, G. (eds.) *Ifs: Conditionals, Belief, Decision, Chance and Time*, pp. 87–104. Reidel, Dordrecht (1981)
15. Swanson, E.: How not to theorize about the language of subjective uncertainty. In: Weatherson, B., Egan, A. (eds.) *Epistemic Modality*. Oxford University Press (2011)
16. Veltman, F.: Defaults in update semantics. *Journal of Philosophical Logic* 25(3), 221–261 (1996)
17. Willer, M.: Dynamics of epistemic modality. *Philosophical Review* (forthcoming)
18. Yalcin, S.: Epistemic modals. *Mind* 116, 983–937 (2007)
19. Yalcin, S.: Modality and Enquiry. Ph.D. thesis, MIT (2008)
20. Yalcin, S.: Probability operators. *Philosophy Compass* 5(11), 916–37 (2010)
21. Yalcin, S.: Nonfactualism about epistemic modality. In: Weatherson, B., Egan, A. (eds.) *Epistemic Modals*. Oxford University Press (2011)
22. Yalcin, S.: Context probabilism. *Proceedings of the 18th Amsterdam Colloquium* (2012)