

How Not to Theorize about the Language of Subjective Uncertainty

Eric Swanson
ericsw@umich.edu

Forthcoming in *Epistemic Modality*, Egan and Weatherson, eds.

When we theorize about linguistic communication it's routine to focus on cases in which someone who is certain or nearly certain that ϕ conveys that ϕ to someone else. But we often communicate from positions of significant subjective uncertainty, and it isn't obvious what features the 'communication of uncertainties' shares with the communication of certainties or near certainties. Indeed, the case of belief suggests that we should expect to see significant differences between these kinds of communication: to be uncertain whether ϕ is not to believe a proposition, and so it would be surprising if it turned out that to express uncertainty about ϕ was to express a proposition.

This paper argues that the following four desiderata must not be neglected when we theorize about the language of subjective uncertainty.

1. We must successfully explain the effects this language can have on addressees' subjective uncertainty.
2. We must successfully explain the effects this language can have on (what I will call) conversational uncertainty.
3. We must explain the 'third grade of modal involvement' exhibited by epistemic modals and by epistemic adjectives.
4. We must explain the norms governing the language of subjective uncertainty, and explain the differences between them and the norms governing the language of subjective certainty.

For helpful discussion, thanks to audiences at University of Michigan, Ann Arbor, Harvard University, and University of California, Berkeley, to Selim Berker, Aaron Bronfman, Andy Egan, Adam Elga, Thony Gillies, John MacFarlane, Dilip Ninan, Rich Thomason, Tim Williamson, Steve Yablo, Seth Yalcin, and to the MATTI reading group at MIT. Thanks especially to Kai von Fintel, Ned Hall, Sarah Moss, Mark Richard, and Bob Stalnaker.

Taken together, these constraints suggest that neither truth conditional nor traditional force modifier theories of the language of subjective uncertainty will be adequate.

Because some ways of theorizing about the language of subjective uncertainty are incompatible with ways of theorizing about language broadly construed, the force of these constraints is more general than it might appear to be. As constraints on theorizing about an important part of language, they are constraints on theorizing about language: a theory of language that rules out every plausible theory of the language of subjective uncertainty is ipso facto implausible. I will argue that important features of the language of subjective uncertainty do not sit well with standard ways of thinking about language broadly construed. And so I take these constraints to have purchase on anyone who theorizes about language.

1.

The literature on epistemic modality tends to focus on modals that can also be used to express non-epistemic modalities, like ‘can,’ ‘might,’ ‘have to,’ and ‘must’. But epistemic modals far outstrip the familiar operators ‘ \square ’ and ‘ \diamond ’:

‘It’s not unlikely that ϕ .’	‘It’s highly probable that ϕ .’
‘It’s a little more likely than not that ϕ .’	‘Probably ϕ .’
‘There’s at least a 10% chance that ϕ .’	‘Five to one that ϕ .’

To some extent Angelika Kratzer recognizes this, discussing some of the ways in which epistemic modality can be “graded” in her 1981. Unfortunately she seems not to recognize just how finely grained the language of subjective uncertainty can be. In many contexts it matters, for example, whether we use “Five to one that ϕ ” or “Six to one that ϕ ” to indicate our uncertainty with respect to the proposition that ϕ . Kratzer’s treatments of modals like ‘there is a good possibility’ and ‘there is a slight possibility’ cannot be extended to these explicitly quantitative expressions of uncertainty. So her approach leaves much of the language of subjective uncertainty unilluminated.

How broadly *should* we construe the language of subjective uncertainty? What phenomena should a theory of this language explain? Many modals have both epistemic and non-epistemic uses, but the quantitative aspects of the language of subjective uncertainty are quite distinctive. This suggests that in order to uncover ways of

expressing subjective uncertainty we would do well to look at the quantitative aspects of subjective uncertainty itself. Where there are aspects of one's doxastic state that are worth communicating, there are likely to be expressions that can help us communicate them.

There is a wide range of thought about what doxastic uncertainty is, how it can be measured, and how it should be represented. But by and large it is agreed that doxastic uncertainty cannot be characterized purely in terms of propositional content.¹ The prevailing theories use probability spaces to represent uncertainty, thereby offering straightforward correlates to quantitatively specified expressions of subjective uncertainty. Given a standard probabilistic theory of uncertainty, for example, it's natural to think that in order to express my 0.8 credence in the proposition that it rained in Seattle yesterday I say

- (1) There's an 80% chance that it rained in Seattle yesterday.

The most straightforward analyses of a sentence like (1) will give as its semantic value something that determines a function from the proposition that it rained in Seattle yesterday to 0.8.

It's compatible with such an analysis that (1) and other 'doxastically hedged' sentences have truth conditions, if those truth conditions are not supposed to give their meaning.² But the stronger view that truth conditions are sufficient to give the meaning of such sentences comes with heavy burdens to discharge. To begin with, we need the right kind of differences between the semantic value of (1), the semantic value of

- (2) There's an 90% chance that it rained in Seattle yesterday.

and so on. So to give truth conditions for doxastically hedged sentences like these we would need a to give a function $f(\cdot)$ from degrees of uncertainty and propositions into propositions, such that a believer is uncertain to degree n about a given proposition just in case she is certain or nearly certain of the proposition that is the image of that

1. Many also think it is a mistake to characterize "uncertain evidence" purely in terms of propositional content. (See especially JEFFREY 1968, 36.) For a contrary view on uncertain evidence, see WILLIAMSON 2000, 213–221.

2. LEWIS 1970, 220–226 takes this sort of view on imperatives.

degree/proposition pair under $f(\cdot)$. That is, to provide truth conditions for (1) and (2) is to provide propositions that one is (nearly) certain about just in case one is uncertain to degree 0.8 or 0.9 about the proposition that it rained in Seattle yesterday. Construing subjective uncertainty about whether ϕ in terms of near certainty about some other proposition seems wrongheaded. But unless the truth conditional theorist can find such propositions, there is no reason to suppose that an assertion of a doxastically hedged sentence will inculcate the appropriate partial belief in the addressee. And the project of finding such propositions looks quixotic if not impossible.

In fact I have considerably understated the challenge, in a few different ways. First, unless she provides a function from *intervals* in $[0, 1]$ and propositions into propositions, the advocate of a truth conditional theory cannot explain how to arrive at the semantic values of sentences like

- (3) There's an 80 to 90% chance that it rained in Seattle yesterday.

Second, we can say

- (4) It's likelier that it rained than that it snowed.
 (5) It's twice as likely that it rained than it is that it snowed.
 (6) It's between two and three times likelier that it rained than that it snowed.

To give such comparatives truth conditions we would need a function from intervals in the positive reals and pairs of propositions into propositions. Finally, consider

- (7) However likely it is that ϕ , it's every bit as likely that ϕ and ψ .

Epistemic comparatives of this form can be used to express the often important information that one's conditional probability of ψ on ϕ is high: if $P(\phi \wedge \psi) \geq P(\phi)$ then $\frac{P(\phi \wedge \psi)}{P(\phi)} \geq 1$, so $P(\psi|\phi) \geq 1$, so $P(\psi|\phi) = 1$. The familiar triviality results that follow from the claim that conditional probability is the probability of a proposition apply, *mutatis mutandis*, to such epistemic comparatives.³ So we have many relatively specialized bits of language that we use to communicate our subjective uncertainty that

3. For helpful presentations of various triviality results, see EDGINGTON 1995 and BENNETT 2003.

are hard to theorize about in purely truth conditional terms.

Ironically, the project of giving truth conditions for doxastically hedged sentences also threatens to undermine an important initial motivation for truth conditional semantics. Propositions—i.e., truth conditions—are supposed to represent ways the world could be.⁴ On the standard picture, to specify the content of a belief using a proposition is to say how the belief represents the world as being, by specifying the conditions under which the belief would successfully represent the world. The ways the world could be are, in principle, exactly the things eligible to be the content of a full belief. In light of subjective uncertainty we should not expect doxastic states to simply represent or misrepresent the world. In just the same way, we should not expect the language of subjective uncertainty to simply represent or misrepresent the world.

Here's one way to see why. As a heuristic, represent a given full belief using an ordered pair consisting of 1 and the proposition that is believed. Let \mathcal{F} be the set of all such ordered pairs. Represent partial beliefs by extending this set to the set \mathcal{P} of all the ordered pairs consisting of some $n \in [0, 1]$ and a proposition. Obviously \mathcal{F} is a proper subset of \mathcal{P} . By stipulation the elements of \mathcal{F} exactly suffice to represent all the ways the world could be. So the members of \mathcal{P} can represent *more* than this: for each way the world could be the members of \mathcal{P} can represent *all the degrees of belief that one could stand in* to that way the world could be. The advocate of a truth conditional theory of the language of subjective uncertainty holds, in effect, that there are not only propositions to play the representational role played by the elements of \mathcal{F} , but also propositions to play the role of the elements of $\mathcal{P} \setminus \mathcal{F}$. Such a view is quite distant from the thought that propositions represent ways the world could be. One fundamental ambition of truth conditional semantics—to help explain how language represents the world by providing the conditions under which language successfully represents the world—thus looks inconsistent with giving a truth conditional theory of the language of subjective uncertainty.

4. Obviously I am taking for granted a possible worlds conception of propositions. Many doubt that unstructured propositions are adequate to all the tasks propositions are supposed to discharge. But because the debate over whether propositions are structured has little to do with subjective uncertainty, I will freely talk in terms of possible worlds propositions. These can be thought of as equivalence classes of more finely-grained structured propositions.

The first desideratum, then, is that theories of the language of subjective uncertainty must not neglect its quantitative aspects. It will be difficult if not impossible for truth conditional theories to capture these aspects of the language of subjective uncertainty, and no extant truth conditional theory comes close.

2.

Conversation is fundamentally a group activity, and uncertainty within a group is quite different from the uncertainty of particular individuals. So although it's crucial that we connect the language of subjective uncertainty to subjective uncertainty itself, it's also important not to overlook *conversational* uncertainty and the ways in which the language of subjective uncertainty interacts with it.

Following Stalnaker, I hold that a conversational participant presupposes that ϕ just in case she takes it to be common belief among the conversational participants that for purposes of conversation they treat it as true that ϕ . On this analysis of presupposition, a conversational participant can presuppose that ϕ , can presuppose that $\neg\phi$, and can presuppose neither that ϕ nor that $\neg\phi$. This last state can be realized in a variety of ways; what's essential is simply that the participant *neither* take it to be common belief that all treat it as true that ϕ nor take it to be common belief that all treat it as true that $\neg\phi$.

In the simplest cases, where the conversational participants all make the same presuppositions, we can say that a conversation is uncertain as to whether ϕ just in case the participants presuppose neither that ϕ nor that $\neg\phi$. Conversational uncertainty of this sort directly affects what can be felicitously presupposed. For example, in a conversation that is uncertain about whether John is married, uses of 'John's wife' will generally be infelicitous unless presupposition accommodation occurs. Conversational uncertainty also makes a difference to which lines of inquiry it's natural to take up in a conversation. Whatever the individual credences of the participants in a conversation, a conversation that is uncertain about whether ϕ will be, *ceteris paribus*, more receptive to inquiry into the question whether ϕ than a conversation that is certain about whether ϕ .

Doxastically hedged assertions can influence not only the subjective uncertainty

of the individuals in a conversation, but also which propositions the conversation is uncertain about. In a normal conversational context in which no one demurs, for example, an utterance of ‘It might be that ϕ ’ ensures that the conversational participants do not presuppose that $\neg\phi$. To see this, consider the following dialogue:

BETTY: I saw Ron walking his dog last night with Sam.

CLARA: Are you sure it was Ron’s dog? It might have been a neighbor’s.

BETTY: # I think it was Ron’s dog, but I might be wrong. Anyhow, Ron’s dog was really misbehaving. . .

Betty’s response is infelicitous because the presuppositions carried by the definite expression ‘Ron’s dog’ are neither in place nor easily accommodated. Betty’s admission that it might not have been Ron’s dog ensures that the context set includes worlds in which Betty was wrong to think that the dog she saw was Ron’s dog. In this particular case, it ensures that the conversation is uncertain about whether the dog she saw was Ron’s dog. And this prevents Betty from appropriately presupposing that ‘Ron’s dog’ denotes the dog she saw.⁵

We can see the context-changing effects of ‘might’ in other places as well. We often use ‘might’ statements when we reject assertions:

SMITH: The weather report says it will definitely rain tomorrow, so it will rain tomorrow.

JONES: It might not rain tomorrow—weather reports are sometimes wrong.

Given a Stalnakerian picture of assertion, the conversational participants have ‘taken on board’ Smith’s assertive utterance that it will rain tomorrow only if the common ground comes to exclude worlds in which it doesn’t rain tomorrow, because to assertively utter a non-hedged sentence ‘ ϕ ’ is to propose that the common ground exclude worlds in which $\neg\phi$. Against this background, Jones exploits the context change potential of ‘It might not rain tomorrow’ to make her rejection of that conversational

5. Notice that these would-be failed presuppositions can be supplied by the antecedent of a conditional, as in:

BETTY: I think it was Ron’s dog, but I might be wrong. Anyhow, if it was Ron’s dog, his dog was really misbehaving. . .

proposal manifest. She in effect proposes that the common ground *include* some worlds in which it doesn't rain tomorrow. By making a proposal that is inconsistent with an intended effect of Smith's assertion, she rejects that assertion.

Stretching the original meaning of "context change potential" somewhat, I will call the aspiration of 'It might be that ϕ ' to ensure that it is not presupposed that $\neg\phi$ the context change potential of 'might' statements.⁶ This context change potential—combined with the crucial fact that speakers can often felicitously use 'might' statements without having much evidence that bears on the truth of the embedded claim—gives 'might' statements a surprising kind of power. Someone who gives very little credence to the proposition that ϕ may nevertheless still be obligated to admit that it *might* be that ϕ : "I might be a bodiless brain in a vat, but I really doubt it." So it's easy to make a *conversation* uncertain as to whether ϕ even if all the conversational participants think it's quite likely that $\neg\phi$. This is one reason why it's hard to argue with conspiracy theorists, skeptics, and the like. Give them an inch of credence, and they are entitled to take a mile of presupposition:

RICHARD: My hand hurts.

TOM: Are you sure you have a hand? You might be a bodiless brain in a vat.

RICHARD: # I think I have a hand, but I might be wrong. Anyhow, my hand has been hurting for several days now.⁷

Richard's response to Tom is not as marked as Betty's response to Clara ("Anyhow, Richard's dog was really misbehaving")—but only insofar as Richard is conveying that he'd prefer not to play the skeptic's game today. In a conversation that is persistently uncertain about a 'hinge' proposition like that expressed by 'I have a hand,' it will be unclear just what we *can* presuppose. And although this is merely a conversational

6. Heim's original way of thinking about context change potentials is significantly less inclusive than this one, in part because she assumes that applying a context change potential to a context will always yield a subset of the initial context. She writes, for example, that "There is an intimate connection between the CCP of a sentence and its truth conditional content: ... To be a true sentence is to keep the context true" (1983, 253).

7. Notice that the antecedent of a conditional can be used here, too, to supply the relevant presuppositions:

RICHARD: I think I have a hand, but I might be wrong. Anyhow, if I have a hand, my hand has been hurting for several days now.

effect—Richard needn't give any less credence to the proposition that he has a hand than he did before Tom raised his skeptical challenge—it makes it difficult if not impossible for Richard to converse normally without seeming to beg the question.

Although subjective uncertainty comes in degrees, conversational uncertainty does not. When I use a definite noun phrase like 'John's wife' I *simply presuppose* that John is married. Presupposition failure may have more or less serious consequences for the course of a conversation, and presupposition accommodation may be more or less surprising and more or less reasonable to expect, but what is actually presupposed is not a matter of degree. In fact this is fortunate: adding degrees to pragmatic presupposition would require a complete overhaul of the standard analysis, and it's unclear how such an overhaul would go. On the standard analysis, a conversational participant pragmatically presupposes that ϕ just in case she takes it to be common belief that all the conversational participants treat it as true, for purposes of conversation, that ϕ . Common p -belief—a notion that adds 'degrees of common-ness' to common belief⁸—is orthogonal to the language of subjective uncertainty: in ordinary conversations someone who says 'It might be that ϕ ' simply makes it common belief that for purposes of conversation we accept that it might be that ϕ . And analyzing degrees of presupposition in terms of common belief about degrees of treating as true for purposes of conversation would require, implausibly, that we coordinate not only on the content of presuppositions but also on the point-valued degrees (and, perhaps, intervals) to which a proposition is presupposed.

This difference between conversational and subjective uncertainty is reflected in the fact that expressions with (what for present purposes we can count as) the same context change potential can express different degrees of subjective uncertainty. Consider

(8) It's raining.

(9) It must be raining.

If I have inferred that it is raining merely on the basis of seeing some wet umbrellas, it is generally inappropriate for me to say (8), even though (9) is appropriate. And,

8. See MONDERER & SAMET 1989 and MORRIS & SHIN 1997.

unsurprisingly, (9) makes a relevantly uninformed addressee less sure that it is raining than (8) does. But despite this difference in strength, (8) and (9) *both* induce the presupposition that it's raining. Similarly for

(10) It isn't raining.

(11) It couldn't be raining.

Both (10) and (11) aspire to make it presupposed that it isn't raining, but they exhibit a difference in strength parallel to that between (8) and (9). My point is not that 'must' and 'couldn't' always signal uncertainty, but that they sometimes do. So a successful treatment of 'must' must at least accommodate such uses.⁹

Phenomena like these spell trouble for the idea that 'doxastic change potential'—be it truth conditions or something probabilistic—can be read off of context change potential.¹⁰ Because sentences like (8) and (9) and like (10) and (11) have the same effects on context (abstracting away from the different form of words), characterizations of their effects on context should make no distinction between them. But such characterizations will omit the differences between their effects on doxastic states. And doxastic change potentials (henceforth, 'DCPs') can't supplant context change poten-

9. F. R. Palmer classifies 'must' as a "Deductive" modal, noting that "it is the notion of deduction or inference from known facts that is the essential feature of *must*, not just the confidence of the speaker, which is expressed by the adverbs *certainly*, *definitely*, etc." (PALMER 2001, 34–35; see also COATES 1983, 41, 131, and 177 and WESTMORELAND 1998, 59–66). To my ear, epistemic 'have to' also has this feature, and under wide scope negation, epistemic 'can' and 'could' do as well. This rescues the hypothesis that 'can' and 'could' are duals of 'must' from the arguments to the contrary in WESTMORELAND 1998. For more on wide scope negation over epistemic modals, see VON FINTEL & IATRIDOU 2003, 184; cf. CINQUE 1999, 198. See also SWANSON 2008 for the evidentiality of epistemic 'should'.

10. The hypothesis that CCP determines doxastic change potential goes back to Irene Heim:

...I will suggest that, while the CCP [context change potential] of "if" cannot be derived from its other properties, one *can* derive the content property from the CCP. More generally, the truth conditional aspect of the meaning of any expression is predictable on the basis of its CCP. (1983, 253)

In later work Heim goes so far as to claim that "The meaning of a sentence is its context change potential" (1992, 185)—a fundamental principle of dynamic semantics. At this point it's not clear to me how separating doxastic change potential from CCP—as I am urging here—would affect dynamic semantics.

tial for all the reasons usually marshalled to think that truth conditions aren't up to the task: the DCP of ' ϕ and ψ ,' for example, does not determine its CCP. So although there are systematic relationships between DCPs and CCPs, the language of subjective uncertainty shows that neither determines the other.

Given that context and context change are non-degreed, there is no hope of recovering doxastic change potential from context change potential. And given that doxastic change potential should not encode—for example—the order of conjuncts, there is no hope of recovering context change potential from it. The second desideratum for theorizing about the language of subjective uncertainty, then, is that we must not neglect the phenomena of conversational uncertainty.

3.

'Force modifier' analyses of epistemic modals may look like good candidates for satisfying both of the desiderata I have discussed so far. According to such analyses, epistemic modals indicate "the speaker's assessment of the truth of the proposition expressed in the [sentence's] residue or the nature of the speaker's commitment to its truth" (HUDDLESTON & PULLUM 2002, 767); they are "modulators of assertive force" (YALCIN 2005, 251).¹¹ Perhaps subjective uncertainty could be expressed by an assertion with tempered force, and perhaps an assertion with tempered force could have the appropriate kind of context change potential. We then would have avoided the task of finding (inter alia) for each proposition P and each degree of credence n , a proxy proposition that garners high credence just in case one has credence n in P .

One familiar objection to force modifier approaches is that they give unsystematic, ad hoc stories about doxastically hedged clauses that are embedded in a larger linguistic context. For example, if 'believes' expresses a relation that holds between a believer and a proposition, then a force modifier account makes it obscure what (12) could mean.

(12) I believe it might have rained in Seattle yesterday.

11. For contemporary examples of such views, see WESTMORELAND 1998, DRUBIG 2001, VON FINTEL 2003, and YALCIN 2005. In the end von Fintel does not endorse a force modifier approach.

For present purposes I want to bracket this family of objections, because I think it is in fact not very hard to achieve compositionality in belief ascriptions (and perhaps even at the clausal level more generally). Obviously we have to drop the assumption that the semantic value of ‘believes’ is simply a relation between a believer and a proposition. But ‘believes’ could denote a relation between a believer, a proposition, and an interval in $[0, 1]$, and thus could be sensitive, in some way, to the ways in which epistemic modals are supposed to “modulate force.”

The really hard problems of compositionality for force modifier approaches arise at the “third grade of modal involvement”: the grade in which we see a quantifier scoping over a modal expression (QUINE 1953). The fact that epistemic modals can exhibit such scope relations shows that the essential doctrine of force modifier approaches—that in asserting a statement headed by an epistemic modal, a speaker puts forward a non-hedged proposition with less than the usual authority or certainty—is untenable.

Contra the “descriptive generalization” argued for in VON FINTEL & IATRIDOU 2003, quantifiers can scope over epistemic modals:

- (13) Al might be the best candidate, Betty might be the best candidate, and Clara might be the best candidate. So most people here might be the best candidate. (MOST PEOPLE HERE $>$ \diamond)
- (14) Lots of people we don’t know might be the murderer, so no one we know *has* to be the murderer. (LOTS OF PEOPLE WE DON’T KNOW $>$ \diamond , NO ONE WE KNOW $>$ \square)
- (15) Be careful where you step, because every inch of the floor might have paint on it. (EVERY INCH OF THE FLOOR $>$ \diamond)

To see why this kind of scope relation is important, note that a speaker who says (15) needn’t commit herself to the claim that it might be that every inch of the floor has paint on it. When I paint, even if I have been sloppy enough so that paint could be anywhere, I’m nevertheless absolutely sure that I haven’t splattered paint everywhere. But there is no proposition, put forward with *whatever* force, that gives the meaning of the relevant reading of (15). The proposition that at least one inch of the floor has paint on it is clearly too weak, and the proposition that every inch of the floor has paint on

it—again, put forward with whatever force—is too strong, since I may be certain that there are paint-free square inches of the floor. This shows that doxastically hedged statements cannot in general be analyzed as ways of putting forward a non-hedged proposition with less than the usual certainty, force, or authority. So the language of subjective uncertainty, appropriately broadly construed, cannot be handled by standard force modifier accounts, according to which a doxastic hedge modulates the force of the speaker’s commitment to a given proposition.

We see similar phenomena with epistemic adjectives:

- (16) This is an easy job; the person we hire for it doesn’t need any special qualifications. So even though only one person will be hired for the job, most of the applicants are possible hires.

The speaker here says that only one person will be hired for the job—thus *denying* that it’s possible that most of the applicants are hires—and yet consistently with that says that most of the applicants are *possible* hires. So

- (17) Most of the applicants are possible hires.

has a reading on which the quantifier scopes over the epistemic adjective. Similar problems arise in other domains where non-truth-conditional theories may be attractive: consider ‘Everyone should take a break, but not everyone should take a break at the same time.’¹²

A positive account of the meaning of sentences like these must pay close attention to the nature of quantification. Just as one can believe that most people are nice without knowing who is nice, one can believe that most of the applicants are possible hires without having any idea *which* of the applicants are possible hires. Some pointers toward a treatment can, I think, be gleaned from noticing the connections between

12. The as yet quite obscure view that quantifiers can ‘scope into speech acts’ might help here. (For work in this vein see KARTTUNEN 1977 and KRIFKA 2001 and 2004.) But one cannot consistently construe the logical form of (15) as (18) and hold that doxastic hedges simply serve to modify the ‘force’ associated with the assertion of a single proposition.

- (18) [For every inch of the floor]_{*i*}, it might be that [that inch]_{*i*} has paint on it.

quantification and disjunction. The proposition that most people in the set $\{a, b, c\}$ are hires is the proposition that exactly a and b are hires, or exactly a and c are hires, or exactly b and c are hires, or exactly $a, b,$ and c are hires. Given plausible assumptions, to believe that most of the applicants are hires is to believe this disjunction. Similarly, to believe that most of the people in the set $\{a, b, c\}$ are *possible* hires is to believe the disjunction ‘Exactly a and b are possible hires, or exactly a and c are possible hires, or exactly b and c are possible hires, or exactly $a, b,$ and c are possible hires.’

Note that one can believe that a and b are possible hires without believing that it’s possible that a and b are hires. Note also that one can believe an ordinary disjunction—one without any doxastic hedges, for example—while having significantly less than full belief in each of its disjuncts. But to *sustain* a belief in an ordinary disjunction a believer must be disposed to update her credences in a way that vindicates certain inferences. For example, a believer who sustains her belief that ϕ or ψ in the face of learning that $\neg\phi$ must come to believe that ψ . Similarly, a believer who sustains her belief that a is a possible hire or b is a possible hire in the face of learning that a is *not* a possible hire must come to believe that b is a possible hire. So a believer who sustains her belief in the disjunction that I am suggesting should be associated with ‘Most of the people in the set $\{a, b, c\}$ are possible hires’ must—to take just one example—come to believe that exactly a and b are possible hires in the face of learning that c is not a possible hire. This is so despite the fact that one can believe that exactly a and b or exactly a and c or exactly b and c or exactly $a, b,$ and c are possible hires without believing that a is a possible hire, or that b is a possible hire, or that c is a possible hire. Put a little less abstractly: you can consistently believe that most of the people in the set $\{a, b, c\}$ are possible hires without believing that any particular one of them is a possible hire. But if in such a state you learn that c is not a possible hire, and you sustain your belief that most of the people in the set $\{a, b, c\}$ are possible hires, you’re rationally constrained to believe that a and b are possible hires.¹³

13. Is believing that it might be that ϕ lending credence greater than zero to ϕ , as in YALCIN 2005? Or is it lending credence greater than or greater than or equal to some non-zero threshold to ϕ ? Suppose, for reductio, that to believe that it might be that ϕ it’s sufficient that one lend non-zero credence to the proposition that ϕ . Consider then a believer who believes that *either* it might be that ϕ or it might be that ψ , without believing that it might be that ϕ and without believing that it might be ψ . Either she assigns non-zero credence to ϕ or she does not, and either she assigns non-zero credence to ψ or she

The similarities between quantification and disjunction suggest one way of giving a compositional, non-truth conditional theory that can handle the third grade of epistemic modality. The thought is that quantified doxastically hedged sentences express something like disjunctions, each disjunct of which is a non-quantified doxastically hedged sentence. (For the details of a semantics that works in this way, see SWANSON 2006.) The important point for present purposes is that a compositional semantics—a semantics the effects of which could not be mimicked by a force modifier approach—needn't traffic solely in truth conditions.

Another point against force modifier approaches is the syntactic flexibility of epistemic adjectives. To begin with, consider

(19) Al is a possible hire.

The only plausible force modifier treatment of (19) that I can see gives it the logical form of (20).

(20) [It's possible that [Al is a hire]]

From a purely syntactic point of view it would be better (*ceteris paribus*) to say that 'possible' combines with 'hire' to form the complex predicate 'possible hire':

(21) [Al [is a possible [hire]]]

The costs of adopting syntactically revisionary theories of epistemic adjectives are even more evident in (22), which is doxastically hedged in two different ways.

(22) Al is a likely candidate and a possible hire.

In normal circumstances, an addressee's belief state after interpreting (22) will be no different (modulo beliefs about the mode of expression) than it would have been if the speaker had said

does not. But given our *reductio* assumption she cannot assign non-zero credence to either without contradicting the stipulation that she does not believe that it might be that ϕ and does not believe that it might be that ψ . So she must assign zero credence to both. But this is surely wrong: there is a difference between believing that either it might be that ϕ or it might be that ψ and believing that $\neg(\phi \vee \psi)$. As a result it's crucial that believers be able to assign sub-threshold credence to a proposition without assigning *zero* credence to that proposition.

(23) It's likely that Al is a candidate. It's possible that Al is a hire.

This suggests, plausibly enough, that in some sense (22) and (23) have the same or very similar *content*. But it is quite another thing to say that they have the same underlying syntactic structure. I see no way for force modifier views to avoid this implausible commitment.

Epistemic adjectives raise no special problems for truth conditional theories. But to my knowledge no advocates of force modifier approaches have tried to extend their theories to cover epistemic adjectives, and the prospects for such an extension look dim. So much the worse for force modifier approaches: epistemic modals and epistemic adjectives are equally a part of the language of subjective uncertainty, and they are interesting for many of the same reasons. For obvious reasons the presence of epistemic adjectives is closely related to the ability of quantifiers to scope over epistemic modals. These features of natural language make it crucial that a theory of the language of subjective uncertainty not ignore Quine's "third grade of modal involvement" in the epistemic realm.

4.

So far I tried to make it clear that it is not only counterintuitive but also deeply mistaken to theorize about the language of subjective uncertainty as though it were all of a piece with the language of subjective certainty (or near certainty; I will leave off this qualification in what follows). Taking the language of subjective certainty as paradigmatic will likely give us a distorted picture of the language of subjective uncertainty. But it is also deeply mistaken to theorize about communication as if it consisted solely of the communication of certainties. We can see this starkly when we consider the norms that govern assertion. The norms of assertion for doxastically hedged statements are not explained by—indeed, they are not even consistent with—many otherwise attractive accounts of the norms of assertion for non-hedged statements.

We can start by considering some distinctive features of the norms governing the use of 'might' statements. Suppose I have no idea where my car keys are, and neither does my housemate. He gets home from work—and so has no good sense of where

I've looked—and I ask him if he knows where my keys are. He says

(24) The keys might be on the kitchen table.

Now his utterance in this case may or may not be *helpful* to me, because I may have already scoured the kitchen table looking for my keys. But whether or not his 'might' statement is helpful to me, it is *appropriate*, and he *knows* that it is appropriate. It wouldn't be fair for me to say in response "No, I've already looked on the kitchen table. They're not there. So why did you say they might be there?" All I can say is something like "No, I've already looked on the kitchen table. They're not there."

Truth conditional semantics for 'might' have considerable trouble making the right predictions about this case. To see why, consider two simple semantics for 'might,' in the spirit of KRATZER 1977, 1981, and 1991.

- A *solipsistic* semantics: 'The keys might be on the table' is true iff it's consistent with what the speaker knows¹⁴ that the keys are on the table.
- A *non-solipsistic* semantics: 'The keys might be on the table' is true iff it's consistent with what the speaker and the addressee know, pooled together, that the keys are on the table.

Given standard assumptions about the norms governing assertion, both these semantics wrongly predict that my housemate's utterance was inappropriate. According to the solipsistic semantics, he asserted a proposition that (we can suppose) he rightly believed would be uninformative to me: I already knew that he didn't know whether the keys were on the table. But assertions that the speaker believes will be uninformative are generally not appropriate.¹⁵ According to the non-solipsistic semantics, my housemate asserted a proposition concerning not only what *he* knows about the location of the keys, but also what *I* know. Given standard assumptions about the norms

14. The arguments to come are also sound if we substitute, e.g., 'believes' for 'knows' throughout.

15. See, for example, GRICE 1987, 26. The right formulation of this constraint on appropriate assertion is a delicate matter, since it is appropriate for me to use

(25) Liem, you ate all the cookies.

simply to let my son know that *I* know that he ate all the cookies. But in such a circumstance I do believe that my utterance of (25) will change my son's beliefs and conversational presuppositions. Appropriate uses of 'might' statements do not require this.

governing assertion, on this semantics my housemate must be certain (or nearly certain) *that I am uncertain* as to whether the keys are on the table, if his utterance is to be appropriate. And in the situation as described he plainly is not certain about this. This brings out the crucial point that a semantics for the language of subjective uncertainty is not plausible unless it leaves room for uncertainty about epistemic states.

Moreover, it's nearly standard to think that a speaker cannot say truly that it might be that ϕ if relevant others know that $\neg\phi$ (HACKING 1967, 146, 148–149; see also TELLER 1972, 310–311, DEROSE 1991, 586–596, and VON FINTEL & GILLIES 2009, ##). But then my housemate would take a serious risk of saying something false with his 'might' statement, presumably would know of this risk, and presumably would be criticizable if I knew that the keys weren't on the table.

But even if we abstract away from particular semantic theories for 'The keys might be on the table,' this example raises acute problems for some putative norms of assertion. For suppose that my keys are in fact not on the table, that I know this, and that my housemate, while not being unreasonable to think that my keys might be on the table, has no special reason to think that they might be there. Then although he may well *believe* that they might be on the table, we wouldn't say that he *knows* that they might be on the table.¹⁶ Nevertheless he can appropriately say "The keys might be on the table." This suggests that the knowledge norm of assertion—"One must assert p only if one knows p " (WILLIAMSON 2000, 243)—is not right. Similarly, Brandom is not right to claim that in asserting speakers "undertake a specific task responsibility, namely the responsibility to show that they are *entitled* to the commitment expressed by their assertions, should that entitlement be brought into question" (1994, 173). Clearly my housemate can say that my keys might be on the table without any special epistemic entitlement—let alone the responsibility to show that he has such an entitlement—and without making any special commitments. Gary Watson voices a similar view—"To assert that p is, among other things, to endorse p , to authorize others to assume that p , to commit oneself to defending p , thereby (typically) giving others standing to criticize or challenge what one says. One exercises one's epistemic

16. Compare DeRose's example of knowing whether it's possible that ϕ : "I don't know whether it's possible that John has cancer; only the doctors know. I'll find that out tomorrow when the results of the test are revealed'" (1991, 584).

authority” (2004, 58). And McDowell claims that the contents of assertion are “epistemic surrogates for represented states of affairs” (1980, 45). These views are subject to similar objections. Obviously I could go on with examples like these, in part because these are *not* implausible things to say about the language of subjective certainty. But when we pay attention to the distinctive features of the language of subjective uncertainty, it becomes clear that these these cannot be the norms of assertion.

The defender of invariably stringent norms of assertion might well respond that ‘might’ statements are not, strictly speaking, assertions, and hence that they are not subject to the norms governing assertion. But this response would make it utterly mysterious what we should say about sentences like (26) and (27).

(26) Half these tires are worn, and might have punctures, too.

(27) Most of the candidates are possible hires, and well-credentialed too.

The norms governing the hedged parts of these sentences are different than the norms governing their non-hedged parts. But neither sentence naturally splits into a pair of syntactic units, so the claim that each involves two speech acts does not look promising. So we should count (26) and (27) both as assertions. They both aim to effect two kinds of doxastic change, which are governed by different norms.

Technically, ‘belief’ and ‘reasonable belief’ norms of assertion do better. As a general rule it seems inappropriate for my housemate to say that my keys might be on the table in circumstances in which he doesn’t even believe that they might be there. But I suspect that this is being found innocent on the basis of a technicality: advocates of belief norms typically take themselves to be proposing that full (or near full) belief is necessary for appropriate assertibility. Otherwise their disagreement with advocates of knowledge norms would be much more dramatic than they have taken it to be: it’s generally (though I think wrongly) presupposed that one knows that ϕ only if one gives full (or near full) credence to the proposition that ϕ .¹⁷

Why is it that the norms governing the appropriate use of ‘might’ statements are less stringent than the norms for statements that aren’t doxastically hedged? When

17. Though see DeRose’s examples of ‘knowing whether it’s possible that ϕ ’ (1991, 584), and see Moss 2009 for a clever argument that knowledge can be had at all levels of credence.

they make non-hedged assertions, at least, speakers seem to claim some authority (or to act as though they were claiming some authority). But a cooperative speaker who does not take herself to have much authority with respect to a given subject matter will signal her own ignorance by saying, for example, that it might be that ϕ , or that it's probably not true that ϕ (as opposed to simply saying that ϕ , or that $\neg\phi$). In so doing she claims less authority than she would have if she had used the language of subjective certainty. By claiming less authority she gives her addressees less stringent claims on her, and so the norms governing her assertion are looser as well.

The relationship between the language that a speaker uses and the authority she thereby claims is in fact a bit more complicated than what I have just said might suggest. To begin with, a speaker who says

(24) The keys might be on the kitchen table.

claims less authority (holding context, intonation, stakes, background conditions, and other relevant factors fixed) than does a speaker who says

(28) The keys are very likely on the kitchen table.

As I suggested earlier, even someone who is ignorant of relevant facts can appropriately (and in some cases helpfully) say (24). By contrast, a speaker has to be in a relatively good epistemic position to appropriately say (28). Another way to see this contrast is to compare (29) and (30) as responses to "Have you seen my keys?"

(29) I don't know, they might be on the kitchen table.

(30) I don't know, they're very likely on the kitchen table.

I find (29) fine, and (30) quite strange. It's plausible that the explicit disavowal of epistemic authority in these sentences—"I don't know..."—is compatible with 'might' and clashes with 'very likely.'

However, many uses of the language of subjective uncertainty do not indicate that the speaker claims less than the usual epistemic authority. A climatologist who, after years of painstaking study, gives precise odds for the average temperature in southeastern Michigan 5,000 years ago justly claims a significant amount of authority when

she says

(31) It's 90% likely that the average was between 46 degrees and 50 degrees.

She is using the language of subjective uncertainty *although she is an expert* on the subject matter at hand. (Similarly, we can use such language to make authoritative claims about objective chance.) On the other hand, if I, as a cooperative non-expert, were to opine on the climate, I might say something like

(32) The average temperature was probably between 46 degrees and 50 degrees.

My “probably” is more vague than the climatologist’s “90% likely,” and also specifies a wider range of acceptable credences. Of course it’s difficult to say just where the lower bound on the range is, and its precise location is no doubt partly a function of context, the interests and values of the conversational participants, and so on. But a believer with *any* credence between that lower bound (say, 0.6, in a particular case) and 1 can be truthfully counted as thinking that probably the average temperature was between 46 and 50 degrees. And it would be misleading, at least, for me to say (32) if my credence weren’t in that range.

When I do say (32), however, I am *not* aiming to ensure that my addressees give a credence within that range to the proposition that the average temperature was between 46 and 50 degrees. By saying only (32) I present myself not as an expert, but as a person with a not wholly unreasonable hunch. If I were to say (32)—unwittingly—to a group of climatologists, I would happily defer to them if they made their expertise manifest and demurred. But in such a case I would not have violated norms of conversation as badly as I would have if I’d presented myself as an expert, giving *precise* odds on the climate.

Or recall the car keys case. I do not know where my car keys are, and neither does my housemate; he does not know where I’ve looked; he says “Your keys might be on the kitchen table.” In many cases he will have spoken appropriately even if I have already searched the kitchen table and know that my keys are not there. I can’t criticize him for giving bad advice about the kind of doxastic state to have. So my housemate intends his advice to have *no force* if I already know that the keys are not on the table. He is attempting to ensure only that I not *inadvertently* rule out or overlook the possibility

that my keys are on the kitchen table.

One important lesson to draw from these examples is that often speakers have quite modest intentions when they use doxastically hedged statements. Often, their hedging indicates that they are communicating from a position of ignorance. But this indication not only conveys a credence other than certainty or near certainty. It also indicates a kind of epistemic and communicative modesty—a disposition to take one’s own credence as less than authoritative, and an intention that one’s addressees take the expression of that credence as less than authoritative—that attenuates the authority that she claims with her assertion, and hence the norms that govern it.

Moreover, different expressions of subjective uncertainty attenuate the speaker’s claim to authority to different degrees. With respect to authority claimed, ‘might’ is like ‘doesn’t have to be,’ ‘must’ is like ‘couldn’t be,’ and so on: holding other relevant factors fixed, it’s more committal to say that the keys couldn’t be in the living room than it is to say that they don’t have to be in the living room, and it is still more committal to come right out and say that they *aren’t* in the living room. This leads to the following generalization.

AUTHORITY REFLECTS RANGE:

The *authority* that a speaker claims in asserting that ϕ decreases with increases in the size of the *range* of credences such that ‘*S* believes that ϕ ’ is true (holding fixed context, content of the prejacent,¹⁸ vagueness of expression, intonation, stakes, background conditions, and other factors that help determine the authority that a speaker claims).¹⁹

A simple assertion with no doxastic hedges is maximally specific. Doxastically hedged statements exhibit degrees of specificity: ‘*Might* ϕ ’ is less specific than ‘*Very likely* ϕ ’, which is less specific than ‘*There’s a 50% chance that* ϕ .’ Accordingly, a speaker who says ‘*Might* ϕ ’ claims less authority than one who says ‘*Very likely* ϕ ’, and so on. ‘*It might be that* ϕ ’ admits a wide range of credence assignments, so that by saying it a cooperative speaker signals that she does not have the epistemic authority to express

18. *Ceteris paribus*, “*Judy might be in Tangiers*” generally claims more authority than “*Judy might be in her office*.”

19. For simplicity, here I’m abstracting away from mixed cases like (22), (26), and (27).

anything that is particularly committal about the right credence to have in ϕ . This generalization also lets us explain why my housemate's suggestion that the keys might be on the table is not criticizable in the ways that non-hedged assertions are.

By appealing to the ways in which the language of subjective uncertainty can modulate the authority of an assertion, we can explain aspects of cases that are commonly used to motivate relativist theories of the content of doxastically hedged statements. Consider

EAVESDROPPING:

The White spies are spying on the Red spies, who are spying on the gun for hire. The gun for hire has left evidence suggesting that he is in Zurich, but one clever White spy knows that he is in London. After finding the planted evidence, one Red spy says to the others, "The gun for hire might be in Zurich," and the others respond "That's true." The clever White spy says "That's false—he's in London" to the other White spies, and explains how he knows this.²⁰

One argument a relativist could give here is that in order to explain the ways in which judgments about the Red spy's utterance are affected by the assessor's epistemic position, the truth value of 'what is said' must be sensitive to the assessor's epistemic position. But note that the Red spy significantly tempers the authority he claims when he says that the gun for hire might be in Zurich, by using an epistemic 'might.' As a result, I respond to the hypothetical relativist, we're inclined to judge the Red spy in a lenient way. It's nevertheless appropriate for the White spy—an observer with more relevant information, and hence more justifiable claim to authority—to say that because the gun for hire is in London, he *couldn't* be in Zurich, and to say that (there is a sense in which) the Red spy spoke falsely. Indeed, I find that my relativist-friendly truth value judgments, such as they are, weaken as the authority claimed is strengthened: I'm more inclined to judge "The gun for hire can't be in London" as straightforwardly

20. A more complicated but in some respects better way to set up this example is with the Red spy using wide-scope negation over a necessity modal. Suppose he had thought that the gun for hire was in London, but on finding the planted evidence says "The gun for hire doesn't have to be in London." This modification to the example undercuts the thought that the demonstratives in the "That's true" / "That's false" responses target the modal's prejacent clause.

false—even if the case is set up so that Red spy’s belief is fully justified—than I am inclined to judge false “The gun for hire might be in Zurich.” AUTHORITY REFLECTS RANGE can help explain why we have relativist-friendly judgments about expressions that claim less than the usual authority: the less authority we claim when making an assertion, the more lenient the norms that govern the assertion.

5.

In theorizing about natural language there is an understandable tendency to take one problem at a time. In many cases this approach is helpful—even essential—if we want to make any progress. But we can get a false impression of understanding when we toil away on one problem without considering the constraints on solutions that other problems impose. So in general it’s important not to work on problems in *too* piecemeal a way. I worry that much theorizing about the language of subjective uncertainty has been too piecemeal. Focusing on one issue at a time has made the space of plausible solutions look bigger than it in fact is.

In particular, the first two constraints—that we explain both the doxastic and conversational changes associated with the language of subjective uncertainty—together suggest that doxastic change potential is not determined by context change potential, and thus that semantics should not deliver one or the other, but both. The third constraint shows that even as we move on from truth conditional semantics, we must do so without abandoning compositionality. Whatever one’s thoughts about the plausibility of compositionality ‘across the board,’ it’s not plausible that compositionality breaks down in simple structures of the form ‘Most *F*’s are *G*’s.’ And although philosophers and linguists often talk as though compositionality demands a truth-conditional semantics, clearly it doesn’t: truth-conditional semantics yields objects that are functions from possible worlds into truth values, but a compositional semantics (that proceeds, as is familiar, via functional application) might well yield functions from sets of possible worlds into real values, or intervals, or the characteristic functions of sets of such functions.²¹ But this does not mean—recalling the second constraint—that

21. In SWANSON 2006 I treat sentences as denoting functions of type $\langle\langle st, v \rangle, t\rangle$, where $D_v = [0, 1]$ (67–68), but other denotations would work equally well.

changes to the context need reflect such fine-grained semantic values. And it does not mean that we must abandon truth-conditional semantics: indeed, we should generalize to the worst case while effectively *embedding* truth-conditional semantics within a probabilistic theory. Finally, because the norms governing a speech act are a function of (inter alia) its meaning, semantic theories of the language of subjective uncertainty need to deliver objects that can be well used by a theory of the norms that govern it. This provides a further constraint on those theories. And again, it would be unproductive, I think, to try to satisfy this desideratum while ignoring the others.

In the 1970s, semantic proposals were generally given for a fragment of a natural language. Since then, the tacit codification of the methods of formal semantics has made it easy to forget that fruitful semantic techniques and frameworks are *fruitful relative to* such a fragment. Indeed, the worth of a framework for a particular fragment may be downright misleading when we begin to consider other and larger fragments of a language. I suspect that this is the case with the language of subjective uncertainty. Purely truth-conditional semantics has given us considerable insight into the language of subjective certainty. But that is no reason to think that we must strictly adhere to that framework in theorizing about the language of subjective uncertainty. The next step is to develop frameworks and theories that aspire to illuminate the language of subjective uncertainty and the language of subjective certainty together.

References

- BENNETT, JONATHAN. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford.
- BRANDON, ROBERT B. 1994. *Making It Explicit*. Harvard University Press, Cambridge.
- CINQUE, GUGLIELMO. 1999. *Adverbs and Functional Heads: A Cross-Linguistic Perspective*. Oxford University Press, Oxford.
- COATES, JENNIFER. 1983. *The Semantics of the Modal Auxiliaries*. Croom Helm, London.
- DEROSE, KEITH. 1991. "Epistemic Possibilities." *Philosophical Review*, vol. 100 (4): 581–605.

- DRUBIG, HANS BERNHARD. 2001. "On the Syntactic Form of Epistemic Modality." Ms., University of Tübingen, URL <http://www.sfb441.uni-tuebingen.de/b2/papers/DrubigModality.pdf>.
- EDGINGTON, DOROTHY. 1995. "On Conditionals." *Mind*, vol. 104: 235–329.
- VON FINTEL, KAI. 2003. "Epistemic Modals and Conditionals Revisited." Ms., Massachusetts Institute of Technology, URL <http://web.mit.edu/fintel/www/umass-handout.pdf>.
- VON FINTEL, KAI & ANTHONY S. GILLIES. 2009. "Might Made Right." In *Epistemic Modality*, ANDY EGAN & BRIAN WEATHERSON, editors. Oxford University Press, Oxford.
- VON FINTEL, KAI & SABINE IATRIDOU. 2003. "Epistemic Containment." *Linguistic Inquiry*, vol. 34 (2): 173–198.
- GRICE, PAUL. 1987. "Logic and Conversation." In *Studies in the Way of Words*. Harvard University Press, Cambridge.
- HACKING, IAN. 1967. "Possibility." *Philosophical Review*, vol. 76 (2): 143–168.
- HEIM, IRENE. 1983. "On the Projection Problem for Presuppositions." In *Formal Semantics: The Essential Readings*, PAUL PORTNER & BARBARA H. PARTEE, editors, 249–260. Blackwell Publishers Ltd., Oxford.
- . 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics*, vol. 9: 183–221.
- HUDDLESTON, RODNEY & GEOFFREY K. PULLUM, editors. 2002. *The Cambridge Grammar of the English Language*. Cambridge University Press, Cambridge.
- JEFFREY, RICHARD C. 1968. "Probable Knowledge." In *Probability and the Art of Judgment*, 30–43. Cambridge University Press, Cambridge.
- KARTTUNEN, LAURI. 1977. "Syntax and Semantics of Questions." *Linguistics and Philosophy*, vol. 1: 3–44.
- KRATZER, ANGELIKA. 1977. "What *Must* and *Can Must* and *Can Mean*." *Linguistics and Philosophy*, vol. 1: 337–355.

- . 1981. “The Notional Category of Modality.” In *Words, Worlds, and Contexts: New Approaches in Word Semantics*, HANS JURGEN EIKMEYER & HANNES RIESER, editors, 38–74. W. de Gruyter, Berlin.
- . 1991. “Modality.” In *Semantics: An International Handbook of Contemporary Research*, ARNIM VON STECHOW & DIETER WUNDERLICH, editors, 639–650. W. de Gruyter, Berlin.
- KRIFKA, MANFRED. 2001. “Quantifying into Question Acts.” *Natural Language Semantics*, vol. 9: 1–40.
- . 2004. “Semantics Below and Above Speech Acts.” Ms., Humboldt Universität, URL <http://amor.rz.hu-berlin.de/~h2816i3x/StanfordLecture2004.pdf>.
- LEWIS, DAVID K. 1970. “General Semantics.” In *Philosophical Papers*, vol. 1. Oxford University Press, Oxford.
- MCDOWELL, JOHN. 1980. “Meaning, Communication, and Knowledge.” In *Meaning, Knowledge, and Reality*, 29–50. Harvard University Press, Cambridge.
- MONDERER, DOV & DOV SAMET. 1989. “Approximating Common Knowledge with Common Beliefs.” *Games and Economic Behavior*, vol. 1: 170–190.
- MORRIS, STEPHEN & HYUN SONG SHIN. 1997. “Approximate Common Knowledge and Co-ordination: Recent Lessons from Game Theory.” *Journal of Logic, Language, and Information*, vol. 6: 171–190.
- MOSS, SARAH. 2009. “Non-Propositional Factives.” Ms., University of Michigan, Ann Arbor.
- PALMER, F. R. 2001. *Mood and Modality*. Cambridge University Press, Cambridge, second edn.
- QUINE, W. V. O. 1953. “Three Grades of Modal Involvement.” In *The Ways of Paradox*, 158–176. Random House, New York.
- SWANSON, ERIC. 2006. *Interactions with Context*. Ph.D. thesis, Massachusetts Institute of Technology.
- . 2008. “Modality in Language.” *Philosophy Compass*. DOI: 10.1111/j.1747-9991.2008.00177.x.

- TELLER, PAUL. 1972. "Epistemic Possibility." *Philosophia*, vol. 2: 302–320.
- WATSON, GARY. 2004. "Asserting and Promising." *Philosophical Studies*, vol. 117: 57–77.
- WESTMORELAND, ROBERT R. 1998. *Information and Intonation in Natural Language Modality*. Ph.D. thesis, Indiana University.
- WILLIAMSON, TIMOTHY. 2000. *Knowledge and its Limits*. Oxford University Press, Oxford.
- YALCIN, SETH. 2005. "Epistemic Modals." In *New Work on Modality*, JON GAJEWSKI, VALENTINE HACQUARD, BERNARD NICKEL & SETH YALCIN, editors, 231–272. MIT Working Papers in Linguistics, vol. 51.