

In press at *Mind and Language*.

# Clarity and the grammar of skepticism<sup>1</sup>

Chris Barker, *NYU*

**Abstract.** Why ever assert clarity? If *It is clear that p* is true, then saying so should be at best superfluous. Barker and Taranto (2003) and Taranto (2006) suggest that asserting clarity reveals information about the beliefs of the discourse participants, specifically, that they both believe that *p*. However, mutual belief is not sufficient to guarantee clarity (*It is clear that God exists*). I propose instead that *It is clear that p* means instead (roughly) ‘the publicly available evidence justifies concluding that *p*’. Then what asserting clarity reveals is information concerning the prevailing epistemic standard that determines whether a body of evidence is sufficient to justify a claim. If so, the semantics of clarity constitutes a grammatical window into the discourse dynamics of inference and skepticism.

## 1. The paradox of asserting clarity

We have two photographs before us. One, labeled “Nawal”, shows a grave young woman in ordinary clothes. The other, labeled “Abby”, shows a different young woman wearing a white lab coat with a stethoscope around her neck.

- (1) a. Nawal is a doctor.
- b. Abby is a doctor.

---

<sup>1</sup>Draft of June 10, 2008. Thanks for valuable discussion to Ariel Cohen, Andy Egan, Peter Pagin, James Pryor, Stephen Schiffer, Christopher Tancredi, Lavi Wolf, participants at the Conference on Vagueness and Language Use, and my editors and referees.

Address for correspondence: 726 Broadway, 7th floor / New York, NY 10003 USA

Email: chris.barker at nyu.edu

These sentences may or may not be true. Certainly nothing in either of the photographs settles the question.

- (2) a. It is clear that Nawal is a doctor.  
 b. It is clear that Abby is a doctor.

The status of the sentences in (2) is quite different, both from their counterparts in (1), and from each other. (2a) is simply false: even if we have unimpeachable private evidence that Nawal is a doctor, it is not *clear* that she is a doctor, at least, not based on the photograph alone.

(2b), on the other hand, might be true. What does its truth depend on? We have a certain amount of evidence that Abby is a doctor—the white coat, the stethoscope—and we have to decide whether that evidence is sufficiently compelling to justify concluding that she is a doctor. This decision will depend in part on the stakes involved (v., e.g., Stanley 2005): if Abby's being a doctor has serious consequences, we will be more reluctant to assert or assent to (1b). Conversely, if we are in a credulous mood, (1b) becomes easier to accept. In other words, given the evidence in the photograph, the truth of (1b) depends on the prevailing standard for skepticism.

The falsehood of (2a) demonstrates a fact that will be crucial in the discussion through this paper: that private evidence is not enough to justify asserting clarity. Not only must there be sufficient evidence supporting the truth of the target proposition, that evidence must be available to all the parties concerned. So asserting clarity makes explicit that the conclusion in question follows from publicly available evidence.

Now, it is not always obvious who the parties concerned will be. Sometimes the relevant group can be spelled out by an overt *to* phrase (e.g., *It is clear to me that Nawal is a doctor*; see the discussion of personal clarity below in section 6). However, in the absence of a *to* phrase, the assessors can be some salient group not necessarily including the listener. For instance, the department chair can return from a meeting and tell her secretary *It is clear that the next chancellor will be Jones*, even if the secretary was not present at the meeting and does not have access to the relevant evidence. Indeed, the group concerned need not even contain the speaker. In a faculty meeting, a colleague can lean close and ask *Is it clear that I'm drunk?*, meaning roughly *It is clear*

*to our colleagues that I'm drunk?*. (Thanks to Stephen Schiffer for this scenario.) Under normal circumstances, however, the assessors of the evidence will include at least the speaker and the listener, and in the absence of a *to* phrase, I will assume that this default is always one possible interpretation.

We can now state the central puzzle: if the evidence supporting the conclusions in (2b) is public, and if that evidence is sufficient to support the conclusion that Abby is a doctor, then the proposition in question should already be obvious to everyone, and hardly needs stating. The paradox of asserting clarity, then, is that if (2b) is true, it is necessarily uninformative!

One clue to the solution of this puzzle is that assertions of clarity are typically used when the evidence is not overwhelmingly compelling. It would be odd to say, for instance, based on the same photograph, *It is clear that Abby is wearing a stethoscope*, since that truly is too clear for clarity to be asserted. Note that the infelicity of asserting clarity in this instance does not follow from a general ban on stating the obvious, since we can perfectly felicitously state *Abby is wearing a stethoscope* if we want to draw the attention of our addressee to some important fact (perhaps, for instance, in preparation for an utterance of (2b)). Somewhat paradoxically, then, clarity is best asserted only when the proposition in question is in genuine doubt. In the analysis below, this observation—that assertions of clarity are most useful in situations of borderline clarity—will be closely related to the fact that clarity, unlike other forms of meta-assertion (*must p*, *so p*, *therefore p*, etc.) is fully gradable (e.g., *very clear*, *reasonably clear*, etc.).

Barker and Taranto (2003) and Taranto (2006:91) develop a theory on which “clarity depends directly on belief, and only indirectly on truth” of the embedded proposition. On their view, (1) can be paraphrased roughly as follows: ‘based on publicly available evidence, the discourse participants believe that Abby is a doctor’. If so, then an assertion of clarity reveals nothing new about the world under discussion (i.e., whether Abby is or is not a doctor), but does reveal something new about the beliefs of the discourse participants.

One correct prediction of the belief theory of clarity is, perhaps surprisingly, that *p* can be clear without *p* being true. This is for the simple reason that believing that *p* does not entail *p*. On the implementation of the belief account in Barker and Taranto (2003) and Taranto (2006), the truth

conditions of *It is clear that p* would guarantee that all of the discourse participants believed that *p*, at the same time that updated discourse context would include worlds in which *p* was false.

- (3) a. We know that *p*, though *p* might be false.
- b. We believe that *p*, though *p* might be false.
- c. It is clear that *p*, though *p* might be false.

Thus for a factive such as *know*, (3a) is internally inconsistent. In contrast, in (3b) we can allow that our belief might be mistaken without any inconsistency. Clarity behaves more like belief in this respect than like knowledge, as expected on the belief theory of clarity. On the analysis below, this behavior will follow from the fallible nature of evidence and inference (see especially section 6.2 for further discussion).

However, the belief theory is too strong. Assume we are both in the presence of some limited evidence that God exists. Perhaps we are considering the intricate complexity of the octopus' eye, and it strikes us that the most plausible explanation for that complexity is that it was designed by an intelligent being. We consider this to be a strong argument, but not in itself conclusive.

- (4) It is clear that God exists.

Imagine further that we both believe with all our hearts that God exists. The shared evidence, even in the presence of our full mutual belief, does not make (4) true. The reason is that the evidence alone is not sufficient for justifying the conclusion that God exists.

In the other direction, we can imagine asserting (5) while harboring a secret irrational but definite belief that Mars supports life.

- (5) It is reasonably clear that Mars is barren of life.

If the evidence justifies concluding (however marginally) that Mars is barren, then it is only rational to believe (however weakly) that Mars is barren. But nothing in the grammar guarantees rational behavior, and it is all too humanly possible to assent to (5) while nonetheless harboring a clairvoyant belief that life exists on Mars.

Yet it would be a mistake to sincerely assert (5) without being prepared to believe that Mars is barren of life. On the analysis here, this follows from the fact that it would be a mistake for the standard for evidence to be set so low as to lead to conclusions you believe may be incorrect. In effect, then, competent assertions of clarity require belief. But this is not because clarity is (shared) belief—it is clarity that compels belief, not the other way around.

In any case, belief is certainly not sufficient for clarity, and so we need to go beyond the belief theory. I will offer here a different analysis on which clarity depends directly on justification, without the mediation of belief.

## 2. Context update: main effects and side effects

I will argue that understanding the meaning of *It is clear that p* requires making a semantic distinction between main effect and side effects. This terminology comes from the theory of programming languages, but has fruitful application in natural languages. The distinction between a main effect and a side effect can be stated quite precisely and explicitly (see, e.g., Barker 2002:222, Potts 2005, Shan 2005, etc.), but an informal development will suffice here. Roughly, the main effect of an assertion is its informational payload, the facts at issue about the world that the assertion is intended to convey. The main effect of an assertion of *John's cat is sick*, then, is the proposition that a certain cat is sick. Side effects involve conventional implications other than the main truth-conditional effect, including presuppositions, the creation of discourse referents, and certain other information routinely maintained in a context (though certainly not including conversational implicatures). The side effects of an assertion of *John's cat is sick* will include the fact that John has a cat, as well as the proposition that the context provides a discourse referent for that cat.

Beaver (2001) observes that sentences can be uttered purely for the sake of their side effects.

- (6) a. Exactly two of the three frogs are in the pond.
- b. One of the frogs is not in the pond.
- c. It is in a tree.

(6a) entails (6b). Therefore (6b) adds no new information—at least, no new information about the world. More precisely, every world in which (6a) is true will also be a world in which (6b) is true, so if a context were modeled simply as a set of worlds, then update with (6b) would have no effect whatsoever on any context that has already been updated with (6a).

However, (6b) clearly does useful work, since it establishes a discourse referent to serve as the antecedent of the pronoun in (6c). To see this, note that if (6b) is omitted from the discourse, (6c) becomes infelicitous (on the intended interpretation). Therefore, reasons Kamp (1988), there must be more to updating context than keeping track of possible worlds—we must also track side effects, including at least a list of discourse referents.

In order to track side effects, I will assume that expressions are evaluated in contexts whose evaluation points have two coordinates. For each point  $\langle w, d \rangle$  in the context set, where  $w$  is a possible world and  $d$  is a possible discourse,  $w$  will be a world that is a candidate for the actual world, and  $d$  will be a candidate for the conversation under way.

Note that the discourse coordinate in our evaluation points will need to keep track of more aspects of the discourse than just variable assignments. In particular, it will be important here to keep track of vague standards. We can think of  $d$ , then, as a little mini-world containing only the discourse underway. Perhaps  $d$  could be a situation (a partial world) in the sense of Kratzer (1989). Just as we can peer into a world  $w$  and determine whether it is raining or whether Abby is a doctor in that world, we will peer into a discourse situation  $d$  in order to see whether a certain index refers to a man or a woman, or what the height cutoff for counting as tall happens to be in  $d$ . Given an evaluation point  $\langle w, d \rangle$ , we will not only be able to answer questions about the world (is Abby a doctor?), but to answer questions about the discourse (is it clear that Abby is a doctor?).

A technical detail: as, e.g., Stalnaker (1998:6) points out, discourses are themselves parts of worlds. If you and I are talking, we will naturally presuppose that we are talking. Then any appropriate candidate for the actual world will necessarily contain within it some version of our conversation. Therefore, he suggests, discourse conditions are simply a particular kind of condition on worlds, namely, constraints on the part of the world in which the conversation currently under way is taking place. If so, then for any appropriate evaluation point  $\langle w, d \rangle$ ,  $d$  must be a subpart of

w. We can think of such a point as a world centered on a conversation, in the spirit of Egan's (e.g, in press) centered worlds. However, we do not need to take a stand on this issue, as long as our evaluation points provide separate access to the world at large and to the discourse.

Calling these 'effects' emphasizes the dynamic aspects of meaning. On traditional dynamic treatments (e.g., Heim 1983 or Groenendijk and Stokhof 1991), the meaning of a sentence is conceived as a function that maps a context of utterance to an updated context incorporating the information conveyed by the sentence. However, it is not necessary to conceive of sentences as context update functions in order to make the required distinction between main effects and side effects. For our purposes, it is perfectly adequate to conceive of both main effects and side effects as static conditions, *world conditions* and *discourse conditions*. This has the advantage of remaining neutral with respect to type of speech act: for assertions, the main effect corresponds to truth conditions; for questions, the main effect corresponds to answerhood conditions; for imperatives, fulfilment conditions, and so on. But for the discussion here, we can assume that main effects amount to truth conditions.

### 3. Epistemic *must* and the missing entailment hypothesis

Perhaps an assertion of clarity merely calls to the attention of the addressee some important fact already entailed by the common ground, but which has somehow not yet been added to the common ground. The evidence is available, the conclusion follows, but for some reason the addressee is hesitating to make that last step to the final conclusion. Perhaps they have a logic deficit, and can't compute the consequences of their own beliefs. "Hmmm," says your interlocutor, "Abby is wearing a lab coat, she is holding a stethoscope, what's going on here?" *Clearly, you idiot, Abby is a doctor.*

I will call this the missing entailment hypothesis. Recently, von Stechow and Gillies have advocated just such an approach for epistemic *must*. Could their approach account for clarity as well? After all, there is a striking similarity between *must* and clarity:

- (7) a. Abby must be a doctor.  
 b. It is clear that Abby is a doctor.

To the extent that (7a) and (7b) are true and usable in similar situations, it makes sense to consider the extent to which a missing entailment analysis could be used for both construction types. However, we will see that despite their semantic similarity, epistemic *must* and clarity are different in important and relevant ways.

On the traditional modal account (e.g., Hughes and Cresswell 1996), *must* p holds in a world  $w$  just in case every epistemically possible world accessible from  $w$  is a p world. Since epistemic accessibility is reflexive, it follows that if *must* p is true, then p is true, hence *must* p is at least as strong a claim as p.

However, in the linguistic literature going back at least to Karttunen 1972, *must* p is treated as weaker than p. On Kratzer's (e.g., 1991) analysis, p is required to hold only in those epistemically possible worlds that are sufficiently normal, where a world is more 'normal' if it conforms to what Kratzer calls a stereotypical ordering source. I will make use of Kratzer's stereotypical ordering source below to provide a suitably gradable notion of clarity.

But von Fintel and Gillies reject the weak theory of epistemic *must*, asserting that *must* strictly entails its prejacent:

(8) It must be raining.

The 'prejacent' is the clause in which *must* occurs, minus *must*: here, the proposition that it be raining. If (8) entails that it is raining (as von Fintel and Gillies argue at length), we have a puzzle similar to our original puzzle concerning the assertion of clarity: if (8) entails that it must be raining, why not simply assert that it is raining? The answer given by von Fintel and Gillies is as follows:

Epistemic modals are also evidential markers: they signal that the prejacent was reached through an inference rather than on the basis of direct observation or trustworthy reports.

If we extend the missing entailment approach to clarity, then what asserting clarity adds beyond the embedded proposition is explicit notification of the inferential nature of the conclusion.

So is *clear* a kind of evidential marker? Many languages have an obligatorily marked distinction among phrase types signalling the source of the information conveyed. A typical evidential system might use verbal morphology to distinguish reports based on direct perception (seeing John) from indirect evidence (muddy footprints) from hearsay. Certainly adjectives such as *clear*, *obvious*, *apparent*—what Taranto (2006) calls discourse adjectives—do not form any sort of grammatical evidential system in the official typological sense. Nevertheless, it is possible that the meaning of *clear* could be similar to the denotation of a genuine evidential marker. After all, evidential systems are so common in the world’s languages because marking the source of evidence is such a useful thing to do. If it turns out that the meaning of clarity shares properties with some evidential markers, so much the better for the general and cross-linguistic implications of our conclusions here.

More technically, von Stechow and Gillies follow Kratzer (e.g., 1991) in elaborating on the usual epistemic accessibility relation by replacing it with a modal base  $f$ , where a modal base is a function mapping each evaluation world  $w$  to a set of propositions  $f(w)$ : those propositions that are known relative to  $w$ .

Borrowing from theories of belief (e.g., Hansson 1999), in addition to the modal base, von Stechow and Gillies associate each evaluation world  $w$  with a KERNEL,  $K_w$ : the set of trusted propositions for that world. The kernel generates the modal base, in the sense that the modal base is the closure of  $K_w$  under logical entailment. If we model propositions as sets of worlds, then we can recover the traditional set of epistemically accessible worlds by intersecting the sets of worlds in the kernel propositions ( $f(w) = \cap K_w$ ). (I assume that the kernel is consistent.)

What von Stechow and Gillies propose, then, is that epistemic *must* presupposes, or perhaps implicates, that the prejacent is in the modal base, but not in the kernel.

### 3.1 Unlike *must*, clarity requires public evidence

We shall see that there are at least two respects in which the missing-entailment analysis of *must* does not immediately lead to an adequate analysis of clarity. The first respect is that it does not require that the relevant evidence be public.

Because *must* can be asserted on the basis of private knowledge, it follows that a use of *must* can reveal new information about the world:

- (9) A. John must be home by now.  
 B. How do you know?  
 A. He left half an hour ago.

If B does not have enough evidence to conclude that John is at home (for instance, B does not know that John left), A's assertion reveals to B that there is some information that A has that B lacks. In update terms, a use of *must* can rule out worlds whose modal base fails to entail the prejacent. Because we've ruled out worlds, a use of *must* can be informative. In terms of main effect versus side effects, *must* can have a non-trivial main effect.

Not so for clarity. As we've seen, for *clearly*, the relevant body of evidence must be public. Since public information is by definition in the common ground, evaluation points must be consistent with public information. If it is public information that Abby is wearing a stethoscope, then every world in the context will be a world in which she is wearing a stethoscope. If it follows from the fact that Abby is wearing a stethoscope (in combination with other public evidence) that she is a doctor, then she must be a doctor in each of the evaluation points. Therefore in any situation in which it is true that it is clear that Abby is a doctor, the common ground already entails that she is a doctor. In other words, the main effect of asserting clarity is necessarily trivial (no worlds are ruled out).

Can there really be an expression type whose main effect is necessarily trivial? Perhaps the information is just more subtle. For instance, it has been suggested to me that what an assertion of clarity reveals is the existence of background information supporting the inference that connects the public evidence with the embedded proposition. If we jointly consider the photo of Abby, wearing her lab coat and her stethoscope, and you assert that it is clear that Abby is a doctor, I may conclude that wearing stethoscopes is diagnostic of doctorhood. If I did not know this before, then your assertion of clarity has informed me of something roughly like *Stethoscopes are worn by doctors*. As long as there are some circumstances under which an assertion of clarity can

provide new information about the world, it is not, as claimed here, an expression type incapable of providing new information about the world after all.

However, if I was unaware of the correlation between stethoscopes and doctors, then it was not in fact clear that Abby was a doctor. You have relied on private information, just as in the *must* case given above. Since clarity requires evidence to be shared, I am perfectly within my rights to reject your assertion of clarity. If you respond by informing me that usually doctors wear stethoscopes, it is only at that point that Abby's doctorhood has become clear.

If I deduce that you must be relying on some private knowledge, I need not object to your assertion of clarity, of course; I can guess what background generalizations you must have relied on, and silently add them (at least provisionally) to my store of knowledge. I might be embarrassed by my ignorance, or unwilling to derail your flow of conversation, much in the same way that I might decline to object to your use of a sentence that presupposes information that I don't possess (*Isn't it awful that John got fired yesterday?*). But the fact that presuppositions can be informative does not change the fact that they are not at-issue entailments. So too with background information that facilitates clarity inferences. Thus I maintain that an assertion of clarity cannot be informative (at least, can't be informative about the world), and we are back to square one.

Therefore even if a missing entailment analysis were correct for clarity, it would support my claim that clarity sentences assert facts about the state of the discourse, not about the state of the world. If the conclusion follows from public evidence, then any world consistent with the public evidence will already be a world consistent with the conclusion. On a missing-entailment analysis, then, the only job that asserting clarity could accomplish would be to adjust which facts already present in the context worlds happen to be salient in the discourse.

### **3.2 Unlike *must*, clarity is essentially a matter of degree**

According to von Stechow and Gillies, *must* logically entails its prejacent. But if we're looking at our photographs, and I wonder why Abby is wearing a lab coat and a stethoscope, you can reply by asserting *She must be a doctor*, despite the fact that wearing a lab coat and a stethoscope does not logically guarantee doctorhood. That Abby is a doctor may be the most plausible explanation, but

it is hardly inevitable. She might be an actor in a television series, or else she might be a fraud, or any of a thousand even more outlandish—but certainly not impossible—explanations. Does this mean we must abandon the strong entailment theory of *must*?

Not necessarily. Strong meanings always give rise to loose talk. Carston (2002: chapter 5) develops an account of how Sperber and Wilson’s notion of an ad-hoc concept explains the way that absolute meanings systematically give rise to approximate meanings. We can decide that *raw* means ‘completely uncooked’, yet still describe a partially cooked steak as raw if it is sufficiently undercooked. This is an instance of forming an ad-hoc concept **raw\***, a “particular relaxing” (p. 336) of the absolute concept, that is, a form of conceptual broadening.

Perhaps, then, we can suppose that clarity as well requires in its basic meaning that the conclusion follows from the evidence by strict entailment, but that we naturally find it useful to assert clarity in situations that do not fully meet the strict entailment standard. However, the second main respect in which clarity differs from epistemic modality suggests that we can usefully say more about the fine-grained semantic behavior of clarity:

- (10) It is becoming clear that p.      very clear, reasonably clear, etc.  
       \*It is becoming must that p.      \*very must, \*reasonably must, etc.

Furthermore, *clear* but not *must* has a comparative and a superlative form (e.g., *It is more clear that Abby is a doctor than that Nawal is a doctor*). Even if this contrast is a mere morphological or syntactic fact, *clear* (along with its close cohorts *obvious*, *apparent*, etc.) is fully grammatically gradable, and *must* is not.

This last claim deserves some clarification. Certainly *must* can be hedged or weakened (examples involving *pretty much must*, *virtually must*, *almost must* are easily found). But even if we decide that *must* were (however unsystematically) gradable, it would be a different sort of gradability than that of *clear*, not only morphosyntactically, but semantically: a gradable *must* would be what Kennedy (2007) calls an absolute predicate, that is, a predicate that picks out the end point of a scale. Kennedy argues that such predicates do not participate in vagueness, and in particular do not give rise to sorites sequences (see section 5.2 establishing a sorites sequence for clarity).

In any case, we shall see that an analysis that recognizes the full gradability of clarity can capture the inferential nature of clarity without needing to distinguish between a kernel and a modal base. Furthermore, we shall see that a suitably gradable account of clarity also explains the fact that clarity, unlike epistemic modals (at least on von Fintel and Gillies's account), does not strictly entail the truth of its target proposition. First, however, we must consider the main effects and side effects of vague expressions.

#### 4. Update with vague predicates

The analysis of clarity below builds on a particular view developed in Barker (2002) of how vague predicates interact with context.

Whether a sentence such as *John is tall* is true depends on a number of factors, including the maximal degree of John's height, and the prevailing standard for tallness. (It also depends on identifying a comparison set, i.e., tall in comparison with basketball players or with jockeys. See, e.g., Kennedy (2007) for discussion. But comparison sets are independent of our concerns here.)

What is of particular interest here is that the standard for tallness depends in part on the discourse situation, and not just on the world at large. Especially in borderline cases, the decision between tall and not tall depends on prevailing conventions, on the judgment of the discourse participants, and on the goals of the conversation (Fara 2000, Kyburg and Morreau 2002).

This dependence on properties of the discourse gives rise to two different ways for an assertion containing a vague predicate to contribute information: by main effect or by side effect. Assume for the sake of argument that people count as tall if they are taller than six feet.

- (11) a. I've never met Bill. What's he like?  
 b. Bill is tall.

In this use, the assertion of (11b) adds information about the world under discussion. In particular, we learn that Bill's height is at least 6 feet. The update effect will be to eliminate all those worlds from the initial evaluation context in which Bill's height fails to exceed 6 feet. This, of course, is

a main effect, since updating with the truth conditions of (11b) tells us something new about the part of the world that is under discussion.

But there is another way that such a sentence can be used.

- (12) a. I'm new in town. What counts as tall around here?  
 b. See Bill over there? Bill is tall.

In this situation, there is no uncertainty concerning Bill's height. Perhaps he is standing in front of us at a party. Imagine even that we just measured Bill's height with a tape measure. Then asserting *Bill is tall* reveals no new information about Bill's height. Instead, in this use, it communicates something about the prevailing standard for tallness in the community in which the discourse is taking place. More technically, the update effect is to eliminate all those worlds in which the standard for tallness in this discourse is less than Bill's height.

Peter Pagin observes that in addition to assertions, questions can also be asked purely for the sake of their side effects. He reports that an acquaintance stood with him observing a light rain through a window. "Is it drizzling?" she asked. Obviously, she wasn't asking to be informed about the state of the weather. Rather, she was asking for information concerning the standard against which degrees of rain intensity are compared in order to qualify as drizzling.

In these uses, update with these sentence adds information purely by means of a side effect. That is, the standard for tallness is a part of the discourse, not of the world under discussion. Because the special context in (12) eliminates the possibility of learning anything new about the world under discussion, the only thing left for the update of (12) to achieve is by means of side effects, in this case, updating our knowledge of the prevailing standard for tallness.

#### **4.1 Relative stupidity**

Barker (2002) argues that some expressions conventionally have no update effect apart from update of a vague standard, i.e., no main effect.

- (13) a. Rumsfeld is stupid.  
 b. Rumsfeld is stupid to defend torture.

In addition to dispositional (individual-level) stupidity, as in (13a), there is stupidity relative to a specific event, as in (13b). That is, (13a) may very well be false at the same time that (13b) is true. The class of adjectives that participate in this per-event evaluation (including *stupid*, *smart*, *lucky*, *rude*, etc.) are discussed by Wilkinson (1970, 1976), Rivière (1983), Barker (2002) and Kertz (2006).

Clearly part of the update effect of (13b) is the usual side effect due to a use of a vague predicate: to eliminate all worlds in which the standard for per-event stupidity exceeds the degree of stupidity of Rumsfeld's defending torture. I argued in Barker (2002) at some length that there are no other update effects.

To see this, first note that the entailment that Rumsfeld defended torture is a presupposition:

- (14) a. Rumsfeld was stupid to defend torture.  
 b. Rumsfeld wasn't stupid to defend torture.

Whether affirmed or denied, both forms in (14) entail that Rumsfeld defended torture, which is the hallmark of presupposition.

There is in addition a presupposition that the subject of the sentence is sentient:

- (15) ?The carpet was(n't) stupid to be cleaned before the party.

Whether affirmed or denied, (15) is odd unless we assume that the carpet is volitionally involved in the cleaning event.

Furthermore, there are no other entailments beyond these two presuppositions. In particular, there are no entailments concerning the mental state of the subject (in contrast with, say, *eager*, as in *Rumsfeld was eager to answer questions*).

According to Barker (2002), then, the only update effect of asserting relative stupidity is to ensure that the degree of stupidity of Rumsfeld's actions exceeds the prevailing standard for relative stupidity. In other words, the only purpose in uttering a sentence like (13b) is in order to negotiate constraints on the standard for a vague predicate. If so, then an utterance of (13b) provides no new information (i.e., no non-presupposed information) about the events under discussion, and instead provides new information only concerning the nature of the discourse underway.

Other expression types that arguably share this property is *You did well to hide the evidence*. Fleisher (2008) offers *a long book to read in one sitting*. Another candidate, of course, as suggested by Barker and Taranto (2003) and Taranto (2006), is asserting clarity.

## 5. The vagueness of clarity

Given that clarity is a question of degree, we must decide what *clear* measures degrees of. Certainly clarity does not involve degrees of probability, as shown by the standard lottery scenario (Kyburg 1961). If your chances of winning a lottery are inversely proportional to the total number of tickets, I cannot appropriately assert *You will lose* (at least, I can't assert it on the basis of knowledge), no matter how many other lottery tickets have been sold. Nor can I appropriately say *It is clear that you will lose*, even if you and I both believe that you will lose (contrary to the predictions of the belief theory of clarity, incidentally).

I propose instead that clarity concerns degrees of justification: the degree to which the relevant evidence supports concluding that *p*.

One argument that clarity deals in degrees of justification is that it is possible to explicitly state the source of the evidence that justifies the assertion of clarity, usually with a *from* phrase:

(16) It is clear from the look on your face that you don't agree.

The *from* phrase describes some source of evidence that is sufficient to justify the claim that the *that*-proposition is true.

A second argument that clarity is about degrees of justification is what it means for a claim to become clear. A couple of naturally-occurring examples will illustrate the claim:

(17) From observations of its behavior in captivity and in the wild, it is becoming clear that the White Shark, too, needs its large, complex brain to function effectively in its day-to-day life.

(18) With the publication of the DART study and more recently the GISSI-Prevenzione Study, it is becoming clear that small intakes of omega-3FAs can significantly impact CHD risk.

In each case, the degree to which the conclusion is probable or true does not change. That is, the degree to which the shark needs its brain, and the degree to which omega acids impact CHD risk remains constant. Rather, what changes is the degree of justification we have for each of these claims. Clarity is achieved when our epistemic state contains a sufficiently high degree of justification for the claim in question.

Now, we can never have enough evidence to be absolutely sure of any contingent fact. Nevertheless, we must boldly decide that some propositions are true anyway, or else allow ourselves to be paralyzed by skepticism. Evidence can take us most of the way towards belief in some proposition, but the last little step will always be a leap of faith. Asserting clarity is a way of announcing that it's time to take that leap, and inviting the addressee to make the leap with you.

### **5.1 Epistemic ignorance versus epistemic uncertainty**

The problem, of course, is to figure out how to think about degrees of justification in terms of truth conditions. As a starting point, consider Williamson's (1994) epistemic theory of vagueness. Williamson claims that if John is borderline tall, there is a fact of the matter: either John is tall, or he is not. However, our knowledge of the exact location of the cutoff point for tallness is incomplete, and we may not (and indeed, perhaps will never) be able to discover which side of the line John falls on. Thus vagueness is a kind of incomplete knowledge. More specifically, it is incomplete knowledge about the state of our discourse, namely, the standard for tallness in that discourse.

Uncertainty about whether a proposition is clear differs from uncertainty about whether an individual (say, Bill) has some property (tallness): uncertainty about clarity is uncertainty about whether some proposition is in the common ground. Perhaps, then, echoing Williamson on tallness, we should think of the common ground as a set of worlds with a crisp, clear boundary: a world is either in the common ground, or is not. However, our knowledge about where precisely that boundary falls is incomplete. After all, what is in the common ground depends on what information is shared with your interlocutor, and it is difficult to know what assumptions are held in common with complete precision.

Uncertainty about the membership of the common ground naturally leads to uncertainty about which propositions follow from the common ground. Thus we can make a distinction here between ignorance and uncertainty. If we don't know whether Abby is a doctor, there will be worlds in our common ground that make that proposition true, and other worlds that do not. We are sure that it might be so, and we are equally sure that it might not be so, given our current state of knowledge. We are ignorant about the truth of the proposition, and we are sure that we are ignorant.

On the other hand, if we have a certain body of evidence already present in the common ground that supports the claim that Abby is a doctor, we may wonder whether there are any worlds in the common ground in which she is not a doctor. They will be unlikely worlds, worlds in which something unexpected or unsuspected is taking place. If we are unsure about exactly which assumptions we share with our interlocutor, we may be uncertain which worlds ought to be included in the common ground.

Note that the more unlikely the worlds we are entertaining in our discourse, the more skeptical we are being. So uncertainty about the boundaries of the common ground amounts to uncertainty about the precise degree of skepticism required for present conversational purposes. This, in turn, depends on our level of interest in the outcome, and in particular the consequences for guessing wrong. If we are deliberating over allegations in a serious criminal investigation, for instance, we may be more inclined to wonder whether the lab coat and the stethoscope are conclusive indications, or merely circumstantial evidence.

In the same way described above that *Bill is tall* can be used (solely) to reduce uncertainty concerning the boundary between tall and not-tall, assertions of clarity can be used to reduce uncertainty concerning the boundary between what is known (mutually assumed) and what is not known.

## **5.2 Sorites for clarity**

If clarity is gradable and vague, presumably it is susceptible to soritical reasoning. Considering in detail a sorites sequence for clarity will support the claim that clarity is sensitive to degrees of justification.

In a standard sorites scenario, we have some vague predicate  $F$ , and we consider the claim that for all  $i$ ,  $F(i) \rightarrow F(i - 1)$ . For instance, it is natural to assent to the claim that if a man with  $i$  hairs is not bald, then a man with  $i - 1$  hairs is not bald. But this gives rise to the sorites paradox: given that a man with 1000003 hairs is not bald, repeated application of the major premise leads us to conclude that a man with 1 hair is not bald, contrary to intuitions.

The major sorites premise says that there are no sharp boundaries in the applicability of a vague predicate: plucking out one hair does not make a man suddenly bald. As expected given the justification analysis here, we shall see that the sorites premise for clarity denies sharp boundaries in the amount of evidence required to justify a claim: removing just a tiny bit of evidence does not make a claim suddenly unclear. In contrast with more usual sorities scenarios, which usually turn on facts about the world (how much hairs, how tall, etc.), sorities for clarity is about our epistemic state, i.e., facts about the state of the discourse.

If we consider an assertion of clarity in which the target proposition itself contains vagueness (*It is clear that Bill is tall*), we have to worry about separating the soritical properties of clarity from those of the target proposition. Therefore I will consider a case in which the target proposition is (arguably) not vague.

(19) It is clear that 1000003 is prime.

In order to construct a sorites sequence, we must introduce incremental evidence in favor of the proposition that 1000003 is prime. It turns out that there are several probabilistic primality tests that can provide exactly the sort of incremental evidence we need.

The Lehmann primality test (e.g., Schneier 1996:259) is one such. This test is based on the fact that whenever  $n$  is prime,  $a^{(n-1)/2} \bmod n$  is either 1 or  $n - 1$  for any positive integer  $a$ . The test works like this: given some possibly prime integer  $n$ , choose an integer  $a$  at random between 0 and  $n$ . If  $a^{(n-1)/2} \bmod n$  is neither 1 nor  $n - 1$ , we know that  $n$  cannot be prime. If the result is consistent with  $n$  being prime, however, we cannot conclude that  $n$  must be prime, since there are choices for  $n$  and  $a$  for which satisfy the test even though  $n$  is not prime. But for fixed non-prime  $n$ , it is known that the probability of choosing an  $a$  at random that satisfies the test is bounded

above by .5. If we wonder whether some  $n$  is prime, each time we choose an  $a$  at random that passes the test, our doubt that  $p$  is not prime should be reduced by at least a factor of 2 (assuming the only thing we know about  $n$  is the result of performing this primality test).

Crucially, the repeated application of the test does not bear on the probability that  $n$  is prime. After all, either  $n$  is prime or not, so the probability of the proposition is either 0 or 1. Rather, repeated tests reveal the probability that our epistemic state unluckily supports an incorrect conclusion. That is, we are gaining information about our epistemic state, not about  $n$ .

For instance, if we choose  $n = 7$  and  $a = 2$ ,  $2^{(7-1)/2} \bmod 7 = 2^3 \bmod 7 = 8 \bmod 7 = 1$ . If 7 were not prime, the probability of choosing an  $a$  with this property is at most .5. Since 7 is in fact prime, the evidence provided by this application of the test supports the correct conclusion.

Similarly, if we choose  $n = 21$  and  $a = 8$ , we find that  $8^{(21-1)/2} \bmod 21 = 8^{10} \bmod 21 = 1$ . If we weren't aware that  $21 = 3 * 7$ , we would be justified in concluding that the chances that 21 is not prime are at most .5. (More pedantically, the probability that if 21 were not prime that our test could fail to reveal that fact.) We might choose again at random and unluckily test with  $a = 13$ , and gain more evidence in favor of the proposition that 21 is prime, so that given our updated epistemic state, we now believe that the chances that 21 is not prime are at most .25. But if we then go on to choose  $a = 2$ , we discover that  $2^{10} \bmod 21 = 16$ , and 16 is neither 1 nor 20, in which case we learn that 21 is definitely not prime.

In the case of 1000003, for each  $a$ , we find that the computation comes out as either 1 or 1000002 each time, consistent with 1000003 being prime. Each time we test a new choice for  $a$ , we learn that the probability that we might have unluckily failed to discover that 1000003 is not prime decreases by a factor of (at least) 2. We become more and more confident that 1000003 is prime. At some point along the sequence of tests, it becomes clear that 1000003 is prime.

At what point will it become clear that 1000003 is prime? In the Wikipedia article for Solovay-Strassen primality test (a different test of the same sort, also characterized by a probability factor of .5), we have

For purposes of cryptography if we pick a sufficiently large value of  $k$ , such as 100, the chance of the algorithm failing is so small that we can use the prime in cryptographic applications without worrying.

Here  $k$  is the number of times the test is iterated with freshly chosen random  $a$ 's. What it means for the algorithm to “fail” is for us to be so unlucky that we choose 100 random numbers that all produced test values of 1 or  $n - 1$  despite  $n$ 's not being prime.

Our degree of certainty that a number is prime changes by a small amount after each iteration of the test. (Indeed, we can easily engineer a modified test so that the degree of certainty increases by an arbitrary small amount.) Given some  $n$ , after only a few applications of the test, it is not clear that  $n$  is prime. For instance, given our experience testing the primality of 21, we would be reluctant to conclude that a number is prime after only one or two iterations of the test. But after 100 tests, depending on what is at stake in our application (e.g., cryptographically securing a low-value credit card transaction), we might conclude that we have enough supporting evidence to conclude that  $n$  is prime. If the evidence is publicly available among some relevant group of assessors, it would be clear to them that  $n$  is prime. Furthermore, we might very well assent to some version of the major sorites premise, e.g., *If it is clear after  $k$  tests that  $n$  is prime, it is clear after  $k-1$  tests that  $n$  is prime.* If so, we have arrived at the sorites paradox for clarity.

Unlike standard sorites sequences, we are not considering a sequence of objects (paint chips that range from red to orange, larger and larger heaps of sand, etc.) that differ continuously along some gradable dimension. Nor are we considering a sequence of distinct propositions that differ, say, in their likelihood. (A given number is either prime or not.) Rather, we are considering a sequence of epistemic states. The facts concerning the state of the world do not change: each  $n$  either is or is not prime, and each choice of  $a$  either satisfies the test or not. Rather, each position in the sequence differs in that we have slightly more information than we had before. Clarity is reached when we have enough evidence. If so, then asserting clarity depends only on our epistemic state. Thus assertions of clarity provide information only about the state of the discourse, and not about the state of the world.

### 5.3 Degrees of justification

It remains only to figure out how to model the degree to which a proposition follows from a body of knowledge.

Let  $f$  be an epistemic modal base that maps each world  $w$  onto a set of propositions representing the common ground in  $w$ . Then  $\cap f(w)$  will be the set of all worlds that are consistent with the common ground in  $w$ . Let  $g$  be a stereotypical ordering source that maps each world  $w$  onto a set of assumptions about the normal course of events in  $w$ . Then  $g(w)$  imposes a partial order on  $\cap f(w)$  in the way described by Kratzer (e.g., 1991):  $w \leq_{g(w)} w'$ —a world  $w$  is at least as normal as a world  $w'$ —just in case  $w$  satisfies all of the assumptions in  $g(w)$  that  $w'$  does. Finally, let  $\mu$  be a function mapping worlds to degrees that is consistent with  $g$  in the following sense:  $\mu(w) \leq \mu(w')$  only if  $w \leq_{g(w)} w'$ . Then *It is clear that*  $p$  will be true at an evaluation point  $\langle d, w \rangle$  just in case every world  $x \in \cap f(w)$  such that  $\mu(x) \leq d(\text{clear})$  is a world that satisfies  $p$ , where  $d(\text{clear})$  is the operative standard for skepticism at that evaluation point.

For instance, *It is clear that Abby is a doctor* will be true of those evaluation points  $\langle d, w \rangle$  when all those worlds consistent with the common ground in  $f(w)$  that are sufficiently normal are also worlds in which Abby is a doctor. Here, a world is “sufficiently normal” if the degree to which it agrees with stereotypical assumptions exceeds the standard assigned to *clear* at that evaluation point (i.e., exceeds  $d(\text{clear})$ ).

In contrast, in the lottery example, *It is clear that you will lose* is correctly predicted to be false. The reason is that that world in which you improbably win the lottery is just as normal as the many worlds in which you do not win. It follows that there is no way for there to be only losing worlds above the clarity threshold.

## 6. Personal clarity versus predicates of personal taste

One of the prominent properties of adjectives like *clear* is that they can take modifiers restricting the scope of the clarity to a specific person (or group of people):

- (20) a. It is clear that Abby is a doctor.  
 b. It is clear **to me** that Abby is a doctor.  
 c. Is it clear **to you** that Abby is a doctor?

I'll call (20b) and (20c) expressions of personal clarity, as opposed to simple clarity in (20a). (Compare to *\*It is probable to me that Abby is a doctor.*)

As discussed above in section 1, I have been assuming that the people who must have access to the relevant evidence defaults to a group containing at least the speaker and the listener. Personal clarity differs from simple clarity by specifying overtly who is entailed to have access to the relevant evidence.

Yet in asserting personal clarity, I am not merely claiming that if you knew what I knew that you would come to the same conclusion I have. For that would assume that you and I had the same standard of clarity. Asserting personal clarity entails only that my personal degree of justification exceeds the relevant standard, and remains neutral about whether you and I have the same standards. In fact, one of the situations in which one is most disposed to resort to personal clarity is precisely when all evidence is shared, but there is some doubt about whether the standard is shared. After all, I can be in a credulous mood without assuming that you will be so easy to convince.

This is by no means the way vagueness normally works. If the standard of tallness is uncertain enough, I can point to Bill and ask you whether he is tall. But if I first assert *Bill is tall*, immediately asking *Is Bill tall?* is inconsistent in a way that differs from the situation with personal clarity. In contrast, the statement in (20b) can easily be followed immediately by the question in (20c).

### **6.1 Predicates of personal taste**

There are some vague predicates that appear to allow discourse participants to maintain different standards, the ones that Lasersohn (2005) calls predicates of personal taste. According to Lasersohn, it is perfectly acceptable for one person to declare *This chili is tasty*, and for another person to declare *This chili isn't tasty*. Note that these predicates (unlike other cases that have been argued to motivate assessor-relative notions of truth, notably epistemic modality, MacFarlane 2006) also

allow explicit specification of the assessor group (*This chili seems tasty to me, This chili tastes good to me*).

Lasersohn views such disputes as faultless disagreement: there is genuine contradiction, yet neither participant has any privileged access to the truth. Indeed, there may not be a fact of the matter. On Lasersohn's analysis, the truth of the statement *this is tasty* is relativized to the tastes of the assessor. Then a proposition can be true relative to one participant but false relative to another—disagreement, but without either participant speaking falsely in any absolute sense.

The analysis here suggests a different approach on which the disagreement is characterized without recourse to relativized truth. Instead, we must recognize that judgments involving gradable vague predicates operate simultaneously along two dimensions, as described above in section 4: they make claims about the world (their main effect), and also about the discourse (their side effects). Even assuming that two discourse participants agree on the facts in the world, including the degree of tallness, clarity, or tastiness of some object, they can still differ on what they consider to be the prevailing vague standard.

If two speakers differ in their beliefs concerning Bill's maximal degree of height, the facts of the world are conclusive, and at most one of them will be right. However, if two speakers differ in their beliefs concerning the prevailing standard for tallness or tastiness, when they are considering borderline cases, the only relevant facts are facts about the discourse underway. Neither participant has any greater natural authority than the other with respect to imposing their preferences for vague standards upon the discourse.

On this view, then, disagreements involving predicates of personal taste or of clarity are not faultless, merely faultless with respect to the facts of the world. They can arise only in a defective context, in a context in which the vague standards are not sufficiently in agreement to give similar results on issues relevant for conversational purposes.

## **6.2 Epistemic incompetence**

Personal clarity explicitly relativizes the judgment expressed by the assertion to the standards of the overt assessors. This makes assertions of personal clarity immune to certain kinds of error:

- (21) A. Abby is a doctor.  
 B. No, she's just an actor.  
 A. Whoops, then I said something false.

If I assert Abby is a doctor, and you contradict me, then I must admit that I have said something false.

- (22) A. It is clear to me that Abby is a doctor.  
 B. No, actually, I asked her, and she said she's just an actor.  
 A. #Whoops, then I said something false.

But if I assert personal clarity, and you provide new information, then I have made a certain kind of mistake, to be sure. But it doesn't follow that I said something false. Based on the evidence I had available to me at the time, I made an inference in line with the appropriate epistemic standard in force at that point in the conversation. So I must admit that I arrived at a conclusion that turned out to be incorrect (Abby is not a doctor). But not only was I justified in asserting that it was clear to me that Abby was a doctor, I spoke truly and accurately, and I do not need to retract my assertion of clarity.

I do have to admit, however, to making a mistake: I set my epistemic standard too low. By relaxing my epistemic standards too far, I reached a conclusion that is not sufficiently reliable for present conversational purposes. As a consequence, after your contribution in (22B), the standard for clarity will increase, and I would no longer be willing to assert (22A) even based on the old evidence (i.e., setting the information conveyed in (22B) aside).

Similarly, in section 2 I suggested that the clarity version of Moore's paradox (*It is clear that Abby is a doctor, though she might not be a doctor*) stems from the fact that the only way for the first clause to be true is if the standard for clarity is set low enough for the evidence to justify the proposition that Abby is a doctor. If it is known or believed that she is not in fact a doctor, that means that the standard for clarity was set too low, since it ruled in conclusions that were evidently not sufficiently reliable for present conversational purposes. Thus the paradoxical sentence can in fact be true, but only in a discourse in which the standard for clarity is set incompetently low.

## 7. Conclusions

Why ever assert clarity? Because doing so reveals information about the epistemic standard of evidence that is operative in a discourse.

Like assertions of relative stupidity, asserting clarity does not have any at-issue entailments that concern the world under discussion. Clarity may depend indirectly on background assumptions about what is the normal course of events in a world (usually, only doctors wear stethoscopes), but precisely because these are background assumptions, they are not what is at issue. As a result, assertions of clarity are conventionally and necessarily used entirely for the sake of their side effects, namely, providing constraints on the location of a vague standard.

More technically, an evaluation point will satisfy a claim of clarity for a proposition  $p$  just in case all of the sufficiently normal worlds epistemically accessible from that evaluation point are worlds in which  $p$  obtains. ‘Sufficiently normal’ means satisfying the expectations about the normal course of events in the evaluation point to a degree that exceeds the contextually-determined standard associated with the adjective *clear* relative to that evaluation point. As a result, asserting clarity does not reveal anything new about the fact of the matter under discussion; rather, it distinguishes among possible states for the discourse. More specifically, it excludes from further consideration evaluation points at which the discourse has an inappropriately stringent epistemic standard.

Thus asserting clarity explicitly rejects excessive skepticism.

## References

- Barker, Chris. 2002: The dynamics of vagueness, *Linguistics and Philosophy* **25.1**, 1–36.
- Barker, Chris and Gina Taranto. 2003: The paradox of asserting clarity. In Paivi Koskinen (ed), *Proceedings of the Western Conference on Linguistics (WECOL) 2002*, Volume 14. Fresno, CA: Department of Linguistics, California State University. 10–21.
- Beaver, David. 2001: *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI Publications.
- Egan, Andy. In press: Billboards, Bombs and Shotgun Weddings. Forthcoming in *Synthese*.

- Fara, Delia Graff. 2000: Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics* **20**, 45–81.
- Groenendijk, Jeroen and Martin Stokhof. 1991: Dynamic predicate logic. *Linguistics and Philosophy* **14**, 39–100.
- von Fintel, Kai and Anthony Gillies. 2007: Must...Stay...Strong.  
Manuscript. <http://kaivonfintel.org/must-stay-strong/>
- Fleisher, Nicholas. 2008: *Adjectives and Infinitives in Composition*. PhD dissertation, UC Berkeley.
- Hansson, Sven Ove. 1999: *A Textbook of Belief Dynamics: Theory Change and Database Updating, Applied Logic Series*. Volume 11. Dordrecht: Kluwer Academic Publishers.
- Heim, Irene. 1983: On the projection problem for presuppositions. WCCFL 2. 114–126. Reprinted in Davis (ed) 1991: *Pragmatics: a reader*. Oxford. 397–405.
- Hughes, G.E. and Max Cresswell. 1996: *A new introduction to modal logic*. Routledge.
- Karttunen, Lauri. 1972: Possible and Must. In J. Kimball (ed), *Syntax and Semantics*, vol. 1. New York: Academic Press. 1–20.
- Kamp, Hans. 1988: Comments. In R. Grim and D. Merrill (eds), *Contents of Thoughts*. Tucson, Arizona: University of Arizona Press. 156–181.
- Kennedy, Christopher. 2007: Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy* **30.1**, 1–45.
- Laura Kertz. 2006: Evaluative adjectives: an adjunct control analysis. In Donald Baumer, David Montero, and Michael Scanlon (eds), *Proceedings of the 25th West Coast Conference on Formal Linguistics*. Somerville, MA: Cascadilla Proceedings Project. 229-235.
- Kratzer, Angelika. 1989: An investigation into the lumps of thought. *Linguistics and Philosophy* **12**, 607–653.
- Kratzer, Angelika. 1991: Modality/Conditionals. In A. von Stechow and D. Wunderlich (eds), *Semantik. Ein internationales Handbuch der zeitgenössischen Forschung*. 639–659.
- Kyburg, Alice and Michael Morreau. 2000: Fitting Words: Vague Language in Context. *Linguistics and Philosophy* **23.6**, 577–597.

- Kyburg, Henry. 1961: *Probability and the logic of rational belief*. Middletown: Wesleyan University Press.
- Laserson, Peter. 2005: Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* **28.6**, 643–686.
- MacFarlane, John. 2006: Epistemic modals are assessment sensitive. To appear in Brian Weather-son and Andy Egan (eds), *Epistemic Modality*, Oxford University Press.
- Potts, Christopher. 2005: *The logic of conventional implicature*. Oxford.
- Rivière, Claude. 1983: Modal adjectives: transformations, synonymy, and complementation. *Lingua* **59:1**, 1–45.
- Schneier, Bruce. 1996: *Applied Cryptography*. Second edition. Wiley.
- Shan, Chung-chieh. 2005: *Linguistic Side Effects*. PhD, Harvard.
- Stalnaker, Robert. 1998: On the representation of context. *Journal of Logic, Language, and Information* **7**, 3–19.
- Stanley, Jason. 2005: *Knowledge and Practical Interests*. Oxford.
- Taranto, Gina. 2006: *Discourse Adjectives*. Routledge.
- Wilkinson, Robert. 1970: Factive complements and action complements. *CLS* **6**, 425–44.
- Wilkinson, Robert. 1976: Modes of predication and implied adverbial complements. *Foundations of Language* **14.2**, 153–194.
- Williamson, Timothy. 1994: *Vagueness*. Routledge.