



HAL
open science

What should "should" mean?

Bridget Copley

► **To cite this version:**

Bridget Copley. What should "should" mean?. Language Under Uncertainty Workshop, Kyoto University, January 2005, 2006. halshs-00093569

HAL Id: halshs-00093569

<https://shs.hal.science/halshs-00093569>

Submitted on 13 Sep 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

What should *should* mean?*

Bridget Copley
CNRS/Université Paris 8

January 3, 2006

Abstract

One analysis of non-deontic *should* treats it as having less-than-universal quantification over the epistemically accessible worlds – the worlds that, for all the speaker knows, could be the actual world. This analysis is based on the intuition that *should* assertions are weaker than are assertions of epistemic *must* sentences. Problems with the traditional analysis, however, indicate that there must be a different reason why these *should* sentences express weaker propositions. This paper argues that non-deontic *should* can involve either epistemic or metaphysical modality; both are weaker than epistemic *must* because *should* does not trigger a presupposition that things work out normally, while *must* does. An initially problematic attempt to extend this analysis to deontic *should* prompts a revision to Kratzer’s theory of modals, in which the division of labor between the modal base and the ordering source is rethought.

1 The traditional view

English *should*, like many other modals, has more than one flavor of modal meaning, as demonstrated by the examples in (1). (1a) conveys the speaker’s assessment that the hearer would do well to go to school tomorrow. (1b) conveys the speaker’s assessment that it will likely rain tomorrow.

*I am indebted to Nathan Klinedinst for helping me to realize a number of consequences of my initial idea. Thanks as well to attendees of the *Language Under Uncertainty* workshop for much helpful discussion, and to everyone who (advertently or inadvertently) provided me with data. All errors and omissions are mine.

- (1) a. You should go to school tomorrow.
 b. It should rain tomorrow.

The same sentence can have both kinds of readings, as in (2), with two paraphrases given in (2a) and (2b):

- (2) Jenny should be there tomorrow.
 a. If she wants to get the job, she really needs to be there tomorrow.
 b. That’s what I gather from what I know about her plans.

In each paraphrase, there are different kinds of facts that are presented as relevant. In (2a), as in (1a), the relevant information seems to be about ideals of some kind: either there is a rule saying what the ideals are, or there is in effect a rule about what one would ideally do if one had certain ideals. This kind of reading is often referred to as a “deontic” or “practical necessity” reading. We will leave this reading aside until the very end of the paper, and instead concentrate on the kind of reading exemplified in (1b) and (2b).

In these examples, the relevant information that the speaker takes into consideration seems to consist of facts about the world; the speaker’s evidence or grounds for believing the proposition to be true. This kind of reading is commonly known as the “epistemic” or “logical necessity” reading of *should* sentences.

Behind the choice of the term “epistemic” is the suspicion that the kind of modality involved has to do with speaker knowledge. That is, it is the same kind of modality that is involved in epistemic readings of *must* sentences as in (3) (the parenthetical in (3) is intended to rule out the deontic reading of *must*, on which Xander is required to be there).

- (3) Xander must be there (his car is outside, his lights are on, etc.).

The starting idea is that the epistemic reading of a sentence *must p* says that all epistemically accessible worlds are *p* worlds (Kratzer, 1991). Let C be a context of utterance, an n -tuple including (at least) the speaker, the time of utterance, and the world of utterance.

- (4) $C := \langle x_C, t_C, w_C \rangle$

Let \mathcal{E}_C be the epistemically accessible set of worlds, the worlds that for all the speaker knows could be the actual world:

- (5) $\mathcal{E}_C :=$ the set of worlds compatible with what x_C knows at t_C in world w_C .

And let **All** be a function with the usual (Barwise and Cooper-inspired) semantics as follows:

- (6) $\mathbf{All} := \lambda p \lambda q . p \cap q = p$

We can assume then that the meaning of *must* is as given in (7); it says of a proposition that it is true in all the worlds that are epistemically accessible to the speaker at the world and time of speech.¹

- (7) $\llbracket \textit{must} \rrbracket(C)(p) := 1$ iff $\mathbf{All}(\mathcal{E}_C)(p)$

Actually, as Kratzer points out, this denotation is too strong; it predicts that the proposition expressed by *must p* entails the one expressed by *p*, since the actual world has to be among the set of epistemically accessible worlds. The solution (following Lewis (1968, 1975)) is to say that the set quantified over is not the set of epistemically accessible worlds, but a subset of those worlds; the set of best epistemically accessible worlds according to some plausibility metric (or “ordering source,” in Kratzer’s terminology). Let us suppose that we can construct a function, relativized to the context *C*, that takes a set of worlds as its argument and returns the subset that contains exactly the worlds in which the least out-of-the-ordinary things happen. We will be coy for the moment about precisely how to construct this function.² We change the denotation of *must* to reflect this new set:

- (8) $\llbracket \textit{must} \rrbracket(C)(p) := 1$ iff $\mathbf{All}(\mathbf{highest-plausibility}_C(\mathcal{E}_C))(p)$

Then, since it need not be true that the actual world is one of the most plausible epistemically accessible worlds, *must p* is weaker than *p*, as desired.

Can this epistemic analysis be extended to non-deontic *should*? If so, how? Horn (1989) proposes that *should* is to *must* as *most* is to *all*, based on the contrast shown in (9).

- (9) a. #Xander must be there, in fact, he should be.

¹For our purposes we are abstracting away from the temporal dimension; propositions are sets of worlds, and “*p* is true in *w*” is the same as “*w* is an element of *p*”.

²For instance, at this point our context will have to contain more than what we have put in it so far.

- b. Xander should be there, in fact, he must be.

Horn argues that the contrast in (9) shows that a *should* sentence expresses a weaker proposition than its *must* counterpart. The argument is based on a Gricean implicature: the proposition expressed by the second clause must not be entailed by that expressed by the first, or else the “Be informative” maxim would be violated. Horn argues that what goes wrong in (9a) is exactly that: The proposition expressed by *he should be* is entailed by the proposition expressed by *Xander must be there*.

This is exactly parallel to the behavior of *all* and *most*.

- (10) a. #I ate all of the raisins, in fact, I ate most of them.
- b. I ate most of the raisins, in fact, I ate all of them.

Horn argues that this weakness is due to a weaker quantifier in *should*. And it is parallel to *most*, not, for instance, *some*, assuming that something like *may* corresponds to *some*, as shown by (11) and (12).

- (11) a. #I ate most of the raisins, in fact, I ate some of them.
- b. I ate some of the raisins, in fact, I ate most of them.
- (12) a. #Xander should be there, in fact, he may be there.
- b. Xander may be there, in fact, he should be there.

To express the desired analysis of *should*, reflecting Horn’s idea in combination with our Kratzer-style modal semantics, we define a function **Most** with the usual semantics:

$$(13) \quad \mathbf{Most} := \lambda p \lambda q . | p \cap q | > | p - q |$$

And we give a denotation for *should* that is true if on most of the epistemically accessible worlds, p.

$$(14) \quad \llbracket \textit{should} \rrbracket (C)(p) := 1 \text{ iff } \mathbf{Most}(\text{highest-plausibility}_C(\mathcal{E}_C))(p)$$

This analysis of *should* I will call the “traditional view.”

2 Evaluating the traditional view

As it turns out, the traditional view does not adequately account for the meaning of *should*.

Unlike *must p* sentences, *should p* sentences are possible with a continuation expressing the speaker's absolute ignorance as to whether *p* is true or not. This contrast is shown in (15).

- (15) a. #The beer must be cold by now, but I have absolutely no idea whether it is.
b. The beer should be cold by now, but I have absolutely no idea whether it is.

The judgment for *must* in (15a) makes sense on the traditional view; if you use *must*, and thereby convey that on all of the most plausible epistemically possible worlds the beer is cold, it would be strange to then comment that you have no idea whether it is or not, giving rise to an instance of Moore's Paradox.³

The question is why the sentence in (15b) is not also an instance of Moore's Paradox. By the traditional view of *should*, the speaker is conveying that on most of the most plausible epistemically accessible worlds, the beer is cold. So if you utter (15b), there must be some reason why most of your most plausible epistemically accessible worlds are *p*-worlds. Perhaps you saw someone put the beer in the fridge. But the fact that there is some reason that the beer is cold on most of the worlds you are considering, is reason enough why you should not be able to assert that you have absolutely no idea if it is cold or not. You do have some idea.

The problem is even worse, however: it's not just that the speaker can continue that they have no idea whether *p*, but they can continue by asserting that not-*p* is true (thus entailing, for our purposes, that *p* is false).

- (16) a. #The beer must be cold by now, but it isn't.
b. The beer should be cold by now, but it isn't.

Again, *must* behaves as expected. If on all the most plausible worlds that for all the speaker knows, could be the actual world, the beer is cold, it is an instance of Moore's Paradox for the speaker to then assert that the beer is actually not cold. *Should*, however, again behaves unexpectedly. It would be strange for the speaker to assert that on most of the most plausible worlds that for all they know could be the actual world, the beer is cold, but that in the actual world it isn't. If the speaker knows that the beer isn't cold, there

³Moore's Paradox is, in its simplest form, the fact that the sentence *p but I don't believe that p* is contradictory.

is no way they can assert that on most of the worlds that for all they know could be the actual world, it is cold now.

A possible objection might be that really we should be using dynamic semantics; the context might be getting changed. After all, (17) (i.e., (16b) with the conjuncts switched), is decidedly odd, so dynamic semantics may well be relevant.

(17) #?The beer isn't cold, but it should be.

The hope is that perhaps dynamic semantics could save the traditional view from (16b). A dynamic semantics story for (17) would proceed roughly as follows. The meaning of *but* requires that we update a context first with the proposition that the beer is not cold, and then update that context with the proposition that it should be. However, if *should p* has the traditional meaning, that on most of the most plausible epistemically accessible worlds, *p*, then it is not informative to update with *should p*, so (17) sounds strange.⁴

Now for the contrast in (16). For the *must* sentence in (16a), the context is first updated with the proposition that on all of the most plausible epistemically accessible worlds, the beer is cold. Then it is not possible to update the context with the proposition that the beer is not cold. If, on the other hand, we are merely uttering the *should* sentence in (16b), there is no problem; first the context is updated with the proposition that on most of the most plausible epistemically accessible worlds, the beer is cold; and then the context is updated with the proposition that on the actual world the beer is not in fact cold.

Note that this explanation hinges on the idea that (16a) is bad exactly because *must* involves universal quantification. Less than universal quantification should pose no problem. However, consider (18).

(18) #The beer may be cold, but it isn't.

The sentence in (18) ought to be felicitous, as the treatment of (18) ought to be similar to that of (16b). We ought to be able to update the context first with the proposition that on some of the most plausible epistemically accessible worlds, the beer is cold, and then with the proposition that the beer is not actually cold. But (18) is not felicitous. If we believe that *may* is

⁴There also seems to be a problem with a presupposition of *but* of contrast, but the same issue arises with *and*: *The beer should be cold, and it isn't* is much better than #?The beer isn't cold, but it should be.

epistemic, this is a problem for the traditional view of epistemic *should*, because the less-than-universal-quantification explanation for the acceptability of (16b) cannot be correct.⁵

It seems an inescapable conclusion that while both *must* and *may* preclude the speaker’s having knowledge about whether *p* holds of the actual world, *should* does not. Why is this so? And what, if anything, does the answer to this question have to do with the reason why *should* is somehow less informative than *must*?

Let us retrace our steps. Recall the initial evidence in (9), repeated below as (19), that a *should* sentence is less informative than the corresponding *must* sentence.

- (19) a. The beer should be cold; in fact, it MUST be.
 b. #The beer must be cold; in fact, it SHOULD be.

The traditional view, we saw, attributes this property to *should* having a weaker force of quantification than *must*, though with both of them quantifying over (the most plausible) epistemically accessible worlds. However, this analysis was shown to wrongly predict that *should* triggers a version of Moore’s paradox, which is not the case.

What could be done to make *should* weaker, aside from giving it a less than universal force of quantification? Let us return to the denotation we gave for *must* in (7) above, repeated below as (20).

$$(20) \quad \llbracket \textit{must} \rrbracket(C)(p) = 1 \text{ iff } \mathbf{All}(\mathbf{highest-plausibility}_C(\mathcal{E}_C))(p)$$

The denotation of *must* suggests two ideas for how to proceed. The first idea is that perhaps *should* constrains the context less than *must* does, by presupposing less. In that case, a *must* sentence would be more informative than a *should* sentence because it would narrow down the context more. Another idea is that *should* does not quantify over the same set of worlds as *must* does; instead it quantifies over a differently-constructed set of worlds that happens to be smaller than $\mathbf{highest-plausibility}_C(\mathcal{E}_C)$. In that case, a *must* sentence would be more informative for reasons that would have to do with the relation between $\mathbf{highest-plausibility}_C(\mathcal{E}_C)$ and whatever this other set of worlds was.

⁵The story is the same for a story in which epistemic modality is defined in terms of the outcome of updating the context in various ways (as in Veltman (1996) and Beaver (2001)). The idea is that there is no problem updating the information state for (17).

Only one of these would have to be true to ensure that *must* sentences are more informative than *should* sentences. In what follows, I will argue that the first idea is correct: *must* sentences presuppose that the actual world will turn out to be among the highest-plausibility worlds, while *should* sentences only presuppose that it is possible for the actual world to turn out to be such a world. This difference, along an axis we shall call *efficacy*, is responsible for the fact that *should* sentences are less informative than *must* sentences.

But the second idea, that *should* quantifies over a different set of worlds than does *must*, also has a role to play. I will argue that some examples of what has been called “epistemic” *should* actually quantify over metaphysically accessible worlds. While a reading of *should* that quantifies over epistemically accessible worlds does exist, it appears to be marked typically by an different pattern of focus. Since it is not clear that the set of metaphysically accessible worlds has any relation to the set of epistemically accessible worlds, this idea does not do the job of making *should* weaker than *must*, but it is worth sorting out the facts.

As a postscript, of course *should* has also the deontic or practical necessity reading as well; extending the theory to account for this reading of *should* will motivate a change to Kratzer’s theory of modals.

First, however, let us take a closer look at the intuitions regarding “epistemic” *should*, to argue that the first idea is correct: *should* presupposes less than *must*. The conclusion we will come to is that these modals differ in the presupposition that they make about whether the expected course of events will actually come to pass. This axis we will call “efficacy”, since it reflects whether the things that are supposed to happen actually manage to happen.

3 Efficacy

So, what has been said about non-deontic readings of *should*? Leech (1971), for one, takes “logical necessity” *should* to indicate that “the speaker has doubts about the soundness of his/her conclusion.” To illustrate the difference between *should* and *must*, he provides the following contrasting glosses for the minimal pair of sentences in (21).

- (21) Leech p. 101
- a. Our guests must be home by now.
(‘I conclude that they are, in that they left half-an-hour ago,

- have a fast car, and live only a few miles away.’)
- b. Our guests should be home by now.
 (‘I conclude that they are, in that . . ., but whether my conclusion is right or not I don’t know – it’s possible they had a breakdown, for instance.’)

What is enlightening about this example is the particular flavor of doubt that surfaces in the gloss of (21b). It’s not that the speaker doubts their own inference from the facts to the conclusion, all else being equal; it’s that *they doubt that all else really will be equal*. The relevant intuition seems to be that *should* sentences say what ought to happen if things proceed normally, whatever “normally” means; this is our “highest plausibility” condition, and also Kratzer’s “ordering source”. So it seems that *must p* presupposes that the actual world is going to be one of the most plausible worlds, while *should p*, if it presupposes anything, presupposes merely that it is (still) possible that the actual world is (going to turn out to be) one of the favored worlds.

How can we model this intuition about the difference between *must* and *should*? What we’d like is for *must* to require that the actual world be in the set **highest-plausibility**_C(\mathcal{E}_C), while *should* requires only that the actual world could be in that set. While the requirement for *must* is sensical, the requirement for *should* is non-sensical, in the current framework. What does it mean for a world to not be in that set, but possibly be in it? Set membership is an in or out affair. What we are looking for seems to be not a relationship between a world and a set of worlds, but between two sets of worlds, which would provide the needed flexibility. That is, we are looking for a mystery set ϕ such that *must* requires it to be a subset of **highest-plausibility**_C(\mathcal{E}_C), while *should* requires only that ϕ and **highest-plausibility**_C(\mathcal{E}_C) have a non-empty intersection. This gets us closer, because it gives us a sense in which a world “could” be in **highest-plausibility**_C(\mathcal{E}). If ϕ is a subset of **highest-plausibility**_C(\mathcal{E}) any world in ϕ must be in **highest-plausibility**_C(\mathcal{E}). However, if ϕ is merely required to have a non-empty intersection with **highest-plausibility**_C(\mathcal{E}), then an arbitrary world in ϕ might be in **highest-plausibility**_C(\mathcal{E}), or it might not.

- (22) a. presupposition of *must*: **highest-plausibility**_C(\mathcal{E}_C) $\subseteq \phi$
 b. presupposition of *should*: **highest-plausibility**_C(\mathcal{E}_C) $\cap \phi \neq \emptyset$

What could play the role of the mystery set ϕ ? Remember the intuition: We want *must* and *should* to convey, via (22a) and (22b), whether the ac-

tual world turns out to be one of the most plausible epistemically accessible worlds. So perhaps an appropriate ϕ would be a set of worlds that are epistemically accessible in a later context: a set just like \mathcal{E}_C , but with a different (later) context. We can call these sets “epistemic states” for short.

In that case, the presupposition of *must* should say that a later epistemic state is a subset of the most plausible epistemically accessible worlds. That is not yet quite right; notice that the speaker of a *must p* sentence need not believe that they will get enough information to ever learn the truth, since (23) is non-contradictory.

- (23) The murderer must have thrown the murder weapon into the Seine, but we’ll never find it.

So it seems we need to say something a bit weaker: that any more informative epistemic state (we assume that any later epistemic state is more informative) is a subset of the best (currently) epistemically accessible worlds. What it means for a later epistemic state to be more informative is that it rules out more possible worlds; it is necessarily a subset of the current epistemic state. However, in general, a later epistemic state need not be a subset of the set of most plausible worlds in the current epistemic state. The proposed presupposition for *must* adds exactly that requirement. Thus:

- (24) for all C, p:
- a. $\llbracket \text{must} \rrbracket (C)(p)$ asserts that **highest-plausibility** $_C(\mathcal{E}_C) \subseteq p$, and presupposes that $\forall \mathcal{E}$ more informative than \mathcal{E}_C :
 $\mathcal{E} \subseteq \text{highest-plausibility}_C(\mathcal{E}_C)$
 - b. $\llbracket \text{should} \rrbracket (C)(p)$ asserts that **highest-plausibility** $_C(\mathcal{E}_C) \subseteq p$, and presupposes that $\forall \mathcal{E}$ more informative than \mathcal{E}_C :
 $\mathcal{E} \cap \text{highest-plausibility}_C(\mathcal{E}_C) \neq \emptyset$

This difference⁶ ensures that Moore’s Paradox does obtain for *must p*, but does not obtain for *should p*. If one asserts *must p*, one is committed to the idea that any more informative epistemic state is a subset of the set of most plausible worlds, as well as the proposition that the most plausible currently accessible worlds are a subset of the set p worlds. Therefore it is contradictory to continue by saying that you do not believe that a more

⁶See Copley (2002) for arguments that *will*, and future modals in general, have a similar presupposition. Werner (2005) also refers to it.

informative epistemic state will not be a subset of the set of p worlds. Since *should p* does not presuppose that any more informative state is a subset of the set of most plausible worlds, this contradiction does not arise for *should p* utterances.

Does this solve the problem of the weakness of *should*? It does, because *should*'s presupposition is weaker than *must*'s.

If this is all correct, we are done.

4 Metaphysical and epistemic modality

So, is it all correct? Seemingly. We have epistemic analyses of both *must* and *should*, explaining why *should p* is weaker than *must p*, also keeping *must p* weaker than p . But as it turns out, we have not yet accounted for the meaning of *should*, as we will see now.

Suppose you are indoors, and a friend who has been outside has just told you that, although it is not raining, the ground is wet outside. It would be entirely appropriate for you to utter the sentence in (25a) as a response. It would not, however, be appropriate to utter the sentence in (25b) (assuming that the only grounds for asserting it was what your friend had just told you).

- (25) a. It must have rained.
b. #It should have rained.

This difference between *must* and *should* cannot be explained by a difference in efficacy, i.e. whether or not the actual world turns out to be one of the best worlds. The reason is that the difference between (25a) and (25b) involves the time at which evidence can be admitted. While (25a) admits evidence of the results of the rain event, (25b) does not. The only way you can utter (25b) is if you have evidence about what the world was like *before* the (putative) rain event. If, for instance, you know that the clouds had been building up, and that a thunderstorm had been approaching, then you would be able to say (25b). But the knowledge (or belief) that the ground is now wet is not relevant.

Note as well that it is not a matter of past knowledge, either. You could felicitously say (25b) even if in the past, before the putative rain event, you didn't know that the thunderstorm was approaching. It seems that that there are two differences between the inference evoked by *must have* and the

inference invoked by *should have*. While *must have* reasons backwards from current evidence, or our knowledge about it, *should have* reasons forward from earlier events. And while *must have* reasons from what is known, *should have* reasons not from what was known in the past, but what was actually the case in the past.

So there are two differences: the direction of reasoning from the evidence, and the kind of evidence (what is known vs. what is actually the case).⁷ The latter difference indicates that we are dealing with two accessibility relations, not one. *Must* is indeed epistemic, since it takes into consideration what is known, and *should* in these examples is not epistemic, since it does not take into consideration what is known. We might call it instead a “metaphysical” modality:⁸ it takes all the facts that are actually the case at a particular time into consideration.

Assuming that we adopt this distinction, why then there is a difference in the direction of inference? Is there any reason why epistemic inferences proceed from present evidence to past or present events, while metaphysical inferences proceed from earlier facts to later facts?

Metaphysical causation, the garden-variety causation of one event causing another, requires that the causing event begin earlier than the caused event. So metaphysical inference is as follows: We infer, given the causing event, that the caused event happened.⁹ Epistemic inference seems at first to be just the opposite: surely there we reason from effect to cause (the wet ground to the raining) instead of the other way around. But in another sense, we are still inferring from cause to effect, using a different flavor of causation. Learning that the ground is wet causes us to conclude that it has rained, assuming that there is no other relevant information. This kind of causation we might term “epistemic causation:” information causes a certain change in a belief state, as long as no other information intervenes. This way of looking at epistemic inference makes it parallel to metaphysical inference, in which an event causes a certain change in the state of the universe, as

⁷A third difference is whether *have* scopes over or under the modal. This difference has been dealt with at length by, for instance, Condoravdi (2001), Stowell (2004), and Demirdache and Uribe-Etxebarria (2005). I will not address it here.

⁸Kratzer would derive such an accessibility relation from a “totally realistic” modal base; the idea is the same. I prefer the term “metaphysical” because of its use in the philosophical literature (see Thomason (1970), e.g.) to talk about the future, and because of the causal distinction discussed just below.

⁹Or, in the case of *should*, that it would have happened if nothing else intervened.

long as no other events intervene. The directionality of epistemic causation, from later events to earlier events, then would have to follow from what can be “information,” a topic I will not pursue here. In any case, modulo any type-shifting that a real definition of “information” might necessitate, we can maintain the same formal denotation for metaphysical *should* as we had for the epistemic denotation of *should*, changing only the ordering source and modal base. The ordering source changes from **highest-plausibility** (no other information intervenes, that is, all else is epistemically equal) to **highest-entropy** (no other events intervene, that is, all else is metaphysically equal). The modal base changes from “what is known” (yielding \mathcal{E}_C , the set of worlds epistemically accessible from C), to “what is/has been the case” (yielding \mathcal{M}_C , the set of worlds metaphysically accessible from C). Additionally, we use the relation “more restrictive than” rather than “more informative than.”

- (26) for all C, p:
 $\llbracket \text{should}_{meta} \rrbracket(C)(p)$ asserts that **highest-entropy** $_C(\mathcal{M}_C) \subseteq p$, and presupposes that $\forall \mathcal{M}$ more restrictive than \mathcal{M}_C :
 $\mathcal{M} \cap \text{highest-entropy}_C(\mathcal{M}_C) \neq \emptyset$

But since “more restrictive” means “later” for metaphysical states, we have a problem. If the non-deontic reading of (27) involves metaphysical *should*, it ought to mean that according to the current state of affairs, if nothing else happens, John is home *after* the moment of utterance. This is because any more restrictive metaphysical state must be a later one, because something has to have happened to make it a different metaphysical state. But (27) cannot have that meaning. Instead it conveys that the speaker thinks John is home now, not later than now.

- (27) John should be at home now.

We could circumvent this problem by replacing “more restrictive than” with “at least as restrictive as,” as in (28).

- (28) for all C, p:
 $\llbracket \text{should}_{meta} \rrbracket(C)(p)$ asserts that **highest-entropy** $_C(\mathcal{M}_C) \subseteq p$, and presupposes that $\forall \mathcal{M}$ **at least as restrictive as** \mathcal{M}_C :
 $\mathcal{M} \cap \text{highest-entropy}_C(\mathcal{M}_C) \neq \emptyset$

However, for the case where $\mathcal{M} = \mathcal{M}_C$, the presupposition would be trivially

true, which is not what we want.

So perhaps (27) is epistemic after all. Even with *should have*, which gave us the original metaphysical example, it is possible to get (28b) to be better in the context where the ground is wet, with a marked intonation:

- (29) I don't understand it. The ground is wet, even the leaves on the trees are wet, as far as the eye can see. It SHOULD have RAINED. But you're telling me it's just a very sophisticated sprinkler system.

I find the marked intonation to be obligatory here.

So epistemic readings of *should have* do exist, albeit with marked intonation. Present-oriented *should* seems to require an epistemic reading. And there is indirect evidence that epistemic readings exist for future-oriented *should*, with the same marked intonation as for *should have*. As Leech points out, *should* is a bit weird out of the blue if the eventuality is not one the speaker desires.

- (30) (adapted slightly from Leech 1971, p. 102)
- a. Our candidate should win the election.
 - b. Roses should grow pretty well in this soil.
 - c. ?*Our candidate should lose the election.
 - d. ?*Roses should grow pretty badly in this soil.

The weirdness of (30c) and (30d), whatever its cause, goes away when the marked intonation is used:

- (31) a. She raised less money than the other candidate, she had a lousy campaign manager, and what's more, she's actually a convicted felon. Our candidate SHOULD LOSE the election. Unless somehow all the other candidate's supporters stay home.
- b. Roses SHOULD GROW pretty badly in this soil, but knowing how serious Jenny is about using the blue stuff, I bet her roses will do just fine.

Since the marked intonation correlates with the epistemic reading of *should have*, we may conclude here that it is the epistemic reading of *should* that escapes Leech's constraint.¹⁰

¹⁰Additional evidence for this claim comes from future-oriented epistemic *may*, which also needs the marked intonation:

So: while not every non-deontic use of *should* is epistemic, some of them are.¹¹ This denotation captures what we need for metaphysical *should*. Before going on to deontic or practical necessity *should*, we must confirm that metaphysical *should* is weaker than epistemic *must*, as desired.

To summarize the discussion so far: We have seen that the “traditional” analysis of non-deontic *should*, based on the idea that *should* has less-than-universal quantification than epistemic *must*, erroneously predicts that *should* should trigger Moore’s paradox. This led us to develop an alternative theory, in which *must*, but not *should*, requires that any later epistemic state turn out not to make p true. This distinction between *must* and *should* was called “efficacy,” since the issue is whether what is now most plausible actually

- | | | | |
|-----|----|-----------------|---|
| (i) | a. | John MAY LEAVE. | √epistemic (contrastive deontic also √) |
| | b. | John may LEAVE. | *epistemic, √deontic |

¹¹*Should* behaves like *must* and *may* with respect to “epistemic containment” (von Fintel and Iatridou, 2003). On the epistemic reading of (i), von Fintel and Iatridou note, *his* cannot be a bound variable.

- (i) #Every student must be home if his light is on.
- a. epistemic reading: *
- b. deontic reading: fine

Epistemic *should* and *may*, as diagnosed by the marked intonation, also block the variable binding:

- (ii) #Every student MAY be HOME tomorrow if his light is on.
- (iii) #Every student SHOULD be HOME tomorrow if his light is on.

But even without the marked intonation, the bound variable reading is not possible, as shown by the unacceptability of (iv).

- (iv) #Every student should be home tomorrow if his light is on.

This means that our metaphysical *should* also exhibits epistemic containment. This fact is not necessarily a problem for the current analysis, however. Even epistemic conditionals without overt epistemic modals have the epistemic containment property:

- (v) #Every student is at home if his light is on.

Since the sentence in (v) has an epistemic flavor in some sense, but no explicit modal, we have to wonder how exactly the conditional semantics, as separate from the modal semantics, plays a role; it may play the same role in (iv). But we will have to leave this question here.

manages to be true.

This theory proved not to be adequate, as some instances of *should* seemed not to involve epistemic inference at all. In addition to epistemic modality, we introduced metaphysical modality to account for these cases. The table below sets out the analogous terms between epistemic and metaphysical modality.¹²

	epistemic	metaphysical
modal base	\mathcal{E} : set of worlds that agree with what is known	\mathcal{M} : set of worlds that agree with what is the case
ordering source	highest-plausibility	highest-entropy

5 Deontic *should*

A similar issue of strength and weakness arises for deontic *should* and *must*:

- (33) a. #You must pay the rent, in fact you should.
 b. You should pay the rent, in fact, you must.

Does the analysis given above for epistemic and metaphysical *should* transfer to deontic *should*? Deontics are supposed to have a metaphysical modal base and an ordering source for the best worlds based on ideals. Suppose we entertain the denotations in (34) (just like the metaphysical denotations, but **highest-contentment** is a function that takes a set of worlds and returns the subset that best agrees with a salient set of ideals).

- (34) for all p, C:
- a. $\llbracket \text{must} \rrbracket(p)(C)$ asserts that **highest-contentment**(\mathcal{M}_C) \subseteq p,
 and presupposes that $\forall \mathcal{M}$ more restrictive than \mathcal{M}_C :
 $(\mathcal{M}) \subseteq \mathbf{highest-contentment}(\mathcal{M}_C)$
 - b. $\llbracket \text{should} \rrbracket(p)(C)$ asserts that **highest-contentment**(\mathcal{M}_C) \subseteq p,
 and presupposes that $\forall \mathcal{M}$ more restrictive than \mathcal{M}_C :
 $\mathcal{M} \cap \mathbf{highest-contentment}(\mathcal{M}_C) \neq \emptyset$

¹²It is worth asking whether the difference between *may* and *might* is that *may* has a positive efficacy presupposition, like *must*, and *might* has a zero efficacy presupposition, like *should*. My intuition is that *may* seems to express possibility given a certain state of affairs, while *might* seems to accept that outside events could intervene.

The presupposition of (34a) at first looks a bit odd, but it is exactly as Werner (2005) argues: “you do what you gotta do,” i.e., *must* implies *will*.¹³

The real problem is with the presupposition in (34b). It says that the set of metaphysically possible worlds at some later time has a non-empty intersection with the set of worlds that (as of this moment) would bring the highest-contentment. That is, it’s still possible that you’ll actually end up with a highest-contentment world. This is fine as far as it goes. Leech (1971) for instance says that:

[I]f someone says *You must buy some new shoes*, it is assumed that the purchase will be carried out; the tone of *must* tolerates little argument. But *You should buy some new shoes* is a different matter – the speaker here could well add in an undertone, ‘but I don’t know whether you will or not.’

It seems to me, however, that the presupposition in (34b) is more applicable to *be supposed to* than to *should*. Consider the minimal pair in (35):

- (35) a. You should buy some new shoes.
b. You’re supposed to buy some new shoes.

The *be supposed to* (“*bst*”) sentence in (35b) corresponds quite nicely to (34b). There is a set of ideals,¹⁴ which we use to construct a function **highest-contentment**. A *bst p* sentence says that all highest-contentment worlds are *p* worlds, and presupposes only that it’s possible for the actual world to turn into a *p* world, not that it will.

The *should* sentence, however, is a little different. Rather than presupposing that the salient set of ideals is a set that the interlocutor accepts, the speaker, if anything, is presupposing only that the interlocutor *might* accept this set of ideals.¹⁵ A paraphrase of a deontic *should p* sentence would thus be as in (36):

¹³Note that you don’t necessarily do what you should do; Werner specifically excludes *should* from his discussion.

¹⁴Of course *bst* requires that the ideals be provided by some authority. I’m betting that that lexical difference is irrelevant to the current discussion.

¹⁵Actually, not necessarily the interlocutor. Suppose you are talking to your friend about Tasha and you say (i).

- (i) Tasha should go to bed now.

- (36) Deontic *should p*: If you accept exactly these ideals, then, assuming certain laws of rational behavior, p

The rational behavior clause is needed because of course you could behave irrationally and not do what you want to do. But assuming that you are rational, you will do what you want to do.¹⁶

At first look, this paraphrase seems to ruin the parallelism between deontic *should* and the other *shoulds*. But note that we can restate the other *shoulds* in a similar way.

- (37) a. Metaphysical *should*: If exactly these facts are true, then, assuming certain laws of physics, p
b. Epistemic *should*: If exactly this information is known, then, assuming certain laws of probability, p

These paraphrases have three components:

- (38) a. The beginning ideals/facts/information
b. The ceteris paribus condition (exactly these, nothing else intervenes)
c. The laws of rational behavior/physics/plausibility

This is a slightly different picture from Kratzer's theory, which has two components:

- (39) a. Modal base: The beginning facts/information
b. Ordering source: ideals, ceteris paribus, possibly laws of physics, plausibility?

On Kratzer's theory, the role the ideals play is in the ordering source. That's what we were assuming when we created **highest-plausibility**, **highest-entropy**, and **highest-contentment**, functions that take a set of worlds handed them by the modal base and return the set of worlds that best agrees

Presumably you mean to say that whoever's in charge of Tasha's bedtime, if they accept the relevant ideals, will make her go to bed. Your interlocutor may or may not be in charge of Tasha, therefore it is not the interlocutor whose ideals matter, but whoever is in charge.

¹⁶One issue not resolved by this paraphrase is the fact that deontic *should p* sentences are strange if it is impossible for p to happen. So the laws of physics may need to be included here; it is not immediately clear to me how to include them.

with the ideals. So it should be analogous to the laws of physics or probability. Here, however, it seems as though the ideals are playing the same role the initial facts are, instead; for Kratzer, that would be the modal base. But on this theory, the analogies between the three kinds of *should* would be very surprising indeed; the ideals need to play the same role as the things that Kratzer put in the modal base.

One good reason to think that this is right is that ideals, like facts and unlike the laws of physics and probability, are particular to particular situations. What's constant are the laws of rational behavior. That, then, is the true analogue to the (naïve) laws of physics and the laws of probability. Note that all of these kinds of laws are deterministic, as long as we know what ideals, facts, or information we started with.

If ideals are in the same basket with things like facts and information, that basket must be able to have a set of propositions that yields no consistent set of worlds (this was the initial observation which led to the use of ideals in the ordering source in the first place). Can this be reconciled with epistemic and metaphysical cases? I think yes. There is no reason to believe that epistemic modals are not doxastic, with weighted beliefs, just as you might have differently-weighted ideals. As for the metaphysical cases, we can think of them as including differently-weighted physical forces.

If we accept this kind of story, we will have to jettison all the denotations we have given so far for modals, since they are predicated on the modal base/ordering source picture. More details need to be worked out, and there is no space to do it here, but here is a sketch of the formal elements of the meaning of *should* and *must*, to be treated as a starting point for future work.

Instead of a modal base, there is a situation argument, a set of propositions: everything that goes into a modal base, and ideals as well, all the things that are particular to the actual time and place. If f is a function that takes a situation and applies the relevant natural laws to it, the general assertion of all flavors of *must* and *should* is simply:

$$(40) \quad \lambda p \lambda s . p(f(s))$$

By choosing the situation argument appropriately, and choosing the natural laws appropriately, the different flavors are generated.¹⁷

¹⁷While it is not immediately clear how (40) might be extended to deal with *may* and *might*, since we are no longer quantifying over possible worlds, it could be done. One might take a subpart of the situation argument to be the argument for f , for example.

The presuppositions are harder to state, but an informal shot at them is given in (41):

- (41) a. presupposition of *must*: the salient situation s^* is exactly what matters (no other facts/ideals enter in)
b. presupposition of *should*: the salient situation s^* may or may not be exactly what matters (other facts/ideals may enter in)

With this admittedly unfinished suggestion we must stop and leave the details for future work.

6 What we should conclude

This discussion has generated three heuristics for further research on the semantics of modals.

Relative strength isn't always relative quantificational force; sometimes it's efficacy. Recall that the traditional theory of *should*, where it had a Most quantifier, was unsatisfying. I suspect that quantifiers like *most* play no role in modal semantics, and would hypothesize that relative weakness can always be explained by (a) the difference between existential and universal quantification, (b) whether the actual world is presupposed to be inertial or not (efficacy), and/or (c) how the quantificational set is selected. Thus, if a speaker of a language tells you that one modal sentence makes a “stronger” claim than another modal sentence, all you can conclude is that there is more work to be done.

Inference from evidence isn't always epistemic modality; sometimes it's metaphysical modality. While it seemed obvious that non-deontic uses of *should* were epistemic in nature, that turned out not always to be the case. There are two different ways of making an inference. Epistemic inference can go backwards in time; metaphysical inference does not. Therefore, if a speaker of a language tells you that a modal sentence involves some evidence for a claim, you still need to find out if it is epistemic or metaphysical inference (or conceivably something else).

We should reconsider “modal bases” and “ordering sources,” at least as they are currently understood. In trying to extend the analysis to deontic *should*, we saw that we needed to reconsider the role played by ideals, and thereby to reconsider the basic framework of modal semantics. Kratzer's theory of modality has brought us a very long way indeed, but it

would benefit from a re-division of labor along the lines sketched above.

References

- Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4(2), 159–210.
- Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. Oxford: CSLI and folli.
- Condoravdi, C. (2001). Temporal interpretation of modals. In D. Beaver, S. Kaufmann, B. Clark, and L. Casillas (Eds.), *Stanford Papers on Semantics*. Palo Alto: CSLI Publications.
- Copley, B. (2002). *The Semantics of the Future*. Ph. D. thesis, MIT. To appear, *Outstanding Dissertations in Linguistics*, Routledge.
- Demirdache, H. and M. Uribe-Etxebarria (2005). On the temporal syntax of non-root modals: Covaluation, binding, and scope reversal. Talk presented at the Paris Time and Modality International Roundtable.
- von Stechow, K. and S. Iatridou (2003). Epistemic containment. *Linguistic Inquiry* 34(2), 173–198.
- Horn, L. R. (1989). *A Natural History of Negation*. University of Chicago Press.
- Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research*. Berlin: de Gruyter.
- Leech, G. N. (1971). *Meaning and the English Verb*. Essex, U.K.: Longman Group Limited.
- Lewis, D. (1968). *On the Plurality of Worlds*. Oxford: Basil Blackwell.
- Lewis, D. (1975). *Counterfactuals*. Cambridge, Mass.: Harvard University Press.
- Stowell, T. (2004). Tense and modals. In J. Guéron and J. Lecarme (Eds.), *The Syntax of Time*, pp. 621–636. Cambridge, Mass.: MIT Press.
- Thomason, R. (1970). Indeterminist time and truth-value gaps. *Theoria* 36, 264–281.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic* 25(3), 221–261.
- Werner, T. (2005). You do what you gotta do, or why *must* implies *will*. In *Proceedings of SALT 15*.