

# Modalities of Normality\*

Seth Yalcin

yalcin@berkeley.edu

December 28, 2014

## 1

The modals *ought* and *should* are widely thought to take both deontic and epistemic readings. At some informal level, this claim is hardly questionable. Considering a sentence like

- (1) Noam ought to be in his office.

It is evident that there are at least two ways to take this sentence. One of these ways has a deontic ring, and seems to concern what is normatively called for in some sense; the other has an epistemic ring, and seems more tied to what it is reasonable to expect to be true in some sense.

But more could be, and usually is, meant by the idea that these modals can take both deontic and epistemic readings. On the stronger interpretation, what is meant is that *ought* and *should* each exhibit two kinds of reading (two “flavors” of modality) which are common to other modals, such that they stand in certain nontrivial logical relations with those other modals (relative to a choice of flavor). If we say in this sense that *ought* and *should* can be epistemic, what we mean is that they share a certain epistemically-flavored reading with other modals like *must*, *have to*, *may* and *might*, such that holding fixed this flavor (and relevant features of context), the corresponding modal sentences exhibit some nontrivial logical interaction or pattern of entailment—in particular, they can be ordered by logical strength:

$$\textit{must } \phi, \textit{ have to } \phi \vDash \textit{ought } \phi, \textit{ should } \phi \vDash \textit{may } \phi, \textit{ might } \phi$$

---

\*Thanks to Nate Charlow, Matthew Chrisman and Sophie Dandelet for detailed comments on an earlier draft. Thanks also to Fabrizio Cariani, Wes Holliday, Daniel Rothschild, Aynat Rubinstein, Judy Thomson and to participants in my minicourse on modality and dynamics at the Hebrew University of Jerusalem in May of 2014 for helpful discussion.

Indeed, *ought* and *should* are often called “weak” necessity modals because they are thought to be logically weaker than the corresponding “strong” necessity modals like *must* and *have to*, when we hold fixed a flavor of these modals as (e.g.) epistemic or deontic. When I talk about the thesis that *ought* and *should* possess both deontic and epistemic readings, I have this stronger sense of “reading” in mind.

This paper has two main objectives. First, I make a case against the idea that *ought* and *should* are capable of true epistemic readings. Anticipated in certain respects by Copley [2004, 2006], we will see that these modals admit of a certain reading distinct from, but easily confused with, the epistemic flavor of modality, which I will provisionally call the *pseudo-epistemic* reading. Once we recognize the existence of the pseudo-epistemic reading, we face the question whether there is any reason to think that, in addition to it, *ought* and *should* can *also* take a true epistemic reading, and function as true epistemic modals. I ultimately find no reason to postulate a true epistemic reading: the pseudo-epistemic reading suffices to explain the data that have prompted theorists to suppose that there is an epistemic reading.

My second objective is to explore the nature of the pseudo-epistemic reading. I will suggest that this reading has to do with what is normally the case, or with what it would be reasonable to expect to be true. I ultimately end up calling the pseudo-epistemic reading the *normality* reading of these modals. Inspired by Veltman [1996], I consider some connections between these modals and default reasoning. I close by noting some connections between the ground covered here and the default logic-based approach to deontic modals explored by Horty [2012].

To be explicit about it: this paper is not directly about deontic *should* and *ought*. It is about their close pseudo-epistemic cousins. But part of what animates the project is the hope to shed some indirect light on deontic *should* and *ought* by first trying to get clearer about the most common non-deontic reading of these modals. A relatively unified theoretical understanding of these two sorts of reading would be desirable if we could get it. Seeking a unified understanding, we expect there to be some structural parallels between the pseudo-epistemic modals and their deontic counterparts.

## 2

Consider a case which many would, at least initially, take as drawing out the putative epistemic reading of the English modals *ought* and *should*. Suppose Jones is in a crowded office building when a severe earthquake hits. The building topples. By sheer accident, nothing falls upon Jones; the building just happens

to crumble in such a way so as not to touch the place where he is standing. He emerges from the rubble as the only survivor. Talking to the media, Jones says in wonderment one or other of the following:

(2) I should be dead right now.

(3) I ought to be dead right now.

In a similar vein, Jones's sister says things like this:

(4) It's incredible! That quake was massive. He should be/ought to be dead. We're so lucky he survived.

Obviously Jones is not saying that in deontically or bouletically preferred situations, he is dead; nor is his sister saying that. That is, we don't have here any (normal) deontic reading of the weak necessity modals. Instead it seems like we have a paradigm case of the putative epistemic reading. But the modals are not functioning here like epistemic modals, for observe that no modal which is uncontroversially epistemic works in these sentences. Notably it would be bizarre for Jones to say:

(5) # I am probably dead right now.

Further, observe that if (2) and (3) were true epistemics, they would on any account obviously entail the epistemic readings of:

(6) # I might be dead right now.

(7) # I may be dead right now.

Generally, sentences that entail defective sentences are defective themselves; (2) and (3) are not defective; so plausibly (2) and (3) do not entail (6)/(7); so plausibly they do not contain true epistemic modals. To have a temporary label, let's call the reading that the weak necessity modals are taking in these sentences the *pseudo-epistemic* reading. The first point is that it is easy to mistake the pseudo-epistemic reading for a true epistemic reading.

Does there exist, in addition to the pseudo-epistemic reading of the weak necessity modals, a true epistemic modal reading of these modals? We investigate this possibility below.

An unpublished paper by Bridget Copley (Copley [2006]) anticipates these observations. She calls attention to the following contrast:

(8) # The beer must be cold by now, but it isn't.

(9) # The beer may be cold by now, but it isn't.

(10) The beer should be cold by now, but it isn't.

Copley correctly notes that these data present a problem for the idea that the *should* in (10) is epistemic. Of course, again, this does not yet show that *should* cannot be epistemic. It shows that there is reading that looks very epistemic but is not.<sup>12</sup>

### 3

Closely related data are also noted by Thomson [2008]. Considering the idea that there is a reading of *ought* tantamount to the epistemic modal *probably*, she raises the following objection:

Consider Rasputin. He was hard to kill. First his assassins poisoned him, then they shot him, then they finally drowned him. Let us imagine that we were there. Let us suppose that the assassins fed him pastries dosed with a powerful, fast-acting poison, and then left him alone for a while, telling him they would be back in half an hour. Half an hour later, one of the assassins said to the others, confidently, “He ought to be dead by now.” The others agreed, and

---

<sup>1</sup>On her ultimate analysis, Copley allows for an epistemic reading of *should*. As noted, I will push against this view.

<sup>2</sup>Swanson [2008], building on Copley, gives another example of the contrast:

(i) They left an hour ago, and there isn't any traffic. So they should be here by now. But they're not.

(ii) # They left an hour ago, and there isn't any traffic. So they must be here by now. But they're not.

In explanation of the contrast, Swanson writes that with “They must be here by now”, the speaker aims to add to the common ground the proposition that they are here—a discourse move in tension with the subsequent remark “But they're not”—whereas with “They should be here by now”, the speaker does not attempt to change the common ground in this way [Swanson, 2008, 1204]. But this difference cannot be the full story, since (as effectively noted by Copley [2006]) it would not account for the contrast between (i) and the following:

(iii) # They left an hour ago, and there isn't any traffic. So they might be here by now. But they're not.

(iv) # They left an hour ago, and there isn't any traffic. So they probably are here by now. But they're not.

Nobody thinks that “They might be here by now” or “They are probably here by now” serve as a means of adding the proposition that they are here to the common ground; yet (iii) and (iv) are marked. A natural conclusion to draw is that the *should* in (i) is not a true epistemic; hence it does not entail the corresponding epistemic *might* claim.

they went to look. Rasputin opened his eyes and glared at them. “He ought to be dead by now!” they said, astonished. It might be thought that when they first said the words, they meant that it was then probable that he was dead. Not so when they second said the words. By the time they second said the words, they knew perfectly well that he wasn’t dead. (202-3)

(The assassins could not have replaced their second remark with “He is probably dead right now!”) The first use of *ought* by the assassins would be reckoned by most everyone as the putative epistemic reading. The second use corresponds to what I have called the pseudo-epistemic reading. But is there any good reason to insist that there are two distinct readings of the modal here, rather than a single (pseudo-epistemic) reading?

In response to her case, Thomson recommends the following analysis:

... what it calls for is simply that we distinguish: if I say “The car keys ought to be on the hall table,” then I assert different propositions, according as my state of knowledge is different. If (i) I don’t know that the car keys are, or that they aren’t, on the hall table, then if I say “They ought to be on the hall table,” what I mean is that it is probable that they are there. If (ii) I know that they aren’t there, then if I say “They ought to be on the hall table,” what I mean is that it was probable that they would be there. (203)

The thought is that on what I have called the quasi-epistemic reading, *ought* (and, I take it, *should*), unlike *probably*, can optionally be evaluated relative to a past state of information, even when superficially appearing in a present tense construction. Roughly, the thought is that *ought* and *should* can mean *probably*, but they can also mean *was probable that it would be*.<sup>3</sup>

Thomson’s idea is a natural one. However, it is subject to the following counterexamples. Suppose an urn has five marbles, one black, four white. A marble is selected at random. We observe it is black. We can say:

(11) It was probable that the marble selected would be white.

But it would be odd to complain:

(12) ? The marble selected ought to be white.

(13) ? The marble selected should be white.

---

<sup>3</sup>Thomson later hints that we might need to change this to: *ought* and *should* can mean *very probably*, but they can also mean *was very probable that it would be*.

Indeed, note it is odd to say these even *before* the marble is revealed.

Here is a tentative step towards diagnosis (not far from some of the suggestions in Copley [2004, 2006]; see also [von Fintel and Iatridou, 2008, 115]): to say (12) or (13), it seems one has to take oneself to be in a situation where the *normal* way for the things to unfold is for the marble to be white. And that obviously is in tension with the stipulations of the case. The probability of white was of course notably higher than black; but this doesn't imply that the selection of the black marble was abnormal. In normal situations of this type—marble selection from the urn described—the marble selected sometimes is black.<sup>4</sup>

For another example, suppose we are playing craps. Alas, you roll snake eyes. You would be correct to believe:

(14) It was likely that the dice would not come up snake eyes.

However, it would be odd to say, and a mistake to believe, either of the following:

(15) ? The dice should not have come up snake eyes.

(16) ? The dice ought not to have come up snake eyes.

A cheater might whisper one of these to her conspirator, right after loading the dice against snake eyes. Outside this kind of thing, however, it rings false.

Another example, switching to the future tense. You buy a lottery ticket in a fair lottery. Compare:

(17) You will very likely lose.

(18) ? You should lose.

(19) ? You ought to lose.

The weak necessity modals seem clearly dispreferred, at least when we try to read them in some epistemically flavored way. Some report getting more of a deontically-flavored reading than an epistemic-like reading here. Of course, that itself is puzzling, on the hypothesis that there is a true epistemic reading. Were there such a reading, we would naturally expect it to shine through in just this kind of example, where the putative epistemic reading would be far more context-appropriate than the deontic reading.<sup>5</sup>

---

<sup>4</sup>Another way to put it might be in terms of *what one would be entitled to expect*. If (12) or (13) were true, one would be entitled to expect, in some sense, that the marble selected would be white. But it is plainly not true in this case that one would be entitled to expect that. (I am indebted here to conversation with Judy Thomson.)

<sup>5</sup>Notice further that the judgments flip if the situation is that you bought every lottery ticket except one. In that case,

Games of chance supply easy cases where *ought* and *should* pull apart from what is or was epistemically likely or unlikely. But we can find such examples without them. Most Bostonians are Americans. Hannah just married Henry, a Bostonian. Is that enough to say: *Henry ought to be an American*? That feels comparatively worse than: *Henry is probably American*.<sup>6</sup> Or again, suppose I am struggling to summarize a colleague’s objection. In the course of doing so, I might say:

(20) I am probably misunderstanding you.

(21) It’s very likely I am misunderstanding you.

(22) I must be misunderstanding you.

By contrast, it would be queer to say:

(23) ? I ought to be misunderstanding you.

(24) ? I should be misunderstanding you.

Finlay [2014] adopts the view that what I am calling pseudo-epistemic *ought* is basically akin to *probably*, except:

Whereas ‘probably’ is tied total present evidence, ‘ought’ like other auxiliaries can be relativized to any background, such as evidence at some previous time *t*, some subset of evidence at *t*, or what some agent *s* believes, and so on. (73)

What we have been observing, however, is that this level of context-sensitivity is profligate. Such a view would suggest that *ought* should be fine pretty much everywhere *probably* is, and in more places besides. That is not the case. A view like Finlay’s overgenerates. Overgeneration is what usually happens when the level of contextual flexibility postulated is out of proportion with the facts.

The earthquake and Rasputin examples were cases where modals which are uncontroversially epistemic are marked, but where *should* and *ought* can take a perfectly acceptable reading. The above examples highlight cases where epistemics like *probably*, *very likely*, and even *must* are fine, but where *should* and

---

You should/ought to win.

seems true. Observe that the probability that you win in this scenario is the same as the probability you lose in the scenario where you bought only one ticket. Nevertheless, the truth values of the corresponding *should* and *ought* sentences diverge. This highlights the way in which *oughts* and *shoulds* come apart from what is likely.

<sup>6</sup>I owe this example to Judy Thomson (p.c.).

*ought* are marked. There seem to be two possibilities. One is that there is no epistemic reading of *should* and *ought*; there is only what I have called the pseudo-epistemic reading, a reading whose nature it remains to explain. A second possibility is that *should* and *ought* are yet capable of a true epistemic reading (in addition to the pseudo-epistemic reading), but it is not one that is even loosely equivalent to something like *probably* or *very probable*.

Either way, these results tell against the idea, in Kratzer [1991], of classifying *probably/likely* semantically with *ought* and *should* as a “weak necessity” modal. It is unlikely these operators have structurally analogous semantics. While the kind of model structure appropriate to probability operators is currently a matter of active investigation (Yalcin [2010, 2012b], Lassiter [2011], Holliday and Icard [2014]), the examples above make it appear doubtful that *ought* and *should* possess a reading that is to be articulated in terms of this kind of structure, as would naturally be expected if they possessed true epistemic readings. The pseudo-epistemic reading of weak necessity modals appears to involve some different kind of structure.

I take it that we should not posit possible readings of modals beyond necessity; and I am unaware of data showing that, above and beyond the pseudo-epistemic reading of weak necessity modals, we must also recognize a true epistemic reading. I do not claim to have shown that this further reading does not exist. But the burden of proof is on those who would wish to recognize such a reading. Meanwhile it seems worthwhile to understand how far one can get without supposing there is such a reading.

Here I should briefly pause to defend against one kind of objection. Kratzer [1977] suggested that it is possible to explicitly control the restriction of a modal with an *in view of*-phrase. If one has this view, one will think it is trivial to show that there can be an epistemic reading of the weak necessity modals. One merely has to cite an example like this:

(25) In view of the evidence, Bob ought to be in his office.

But I deny that (25) settles anything, for I deny it is generally true that *in view of*-phrases systematically semantically control the interpretation of modals. The connection between the interpretation of modals and *in view of*-phrases is loose and indirect at best, as the following examples illustrate:

(26) In view of what we know, you can’t be parked here. (deontic reading of the modal clearly available)

(27) In view of what the tribal laws are, the guy performing the ceremony must be the chief. (epistemic reading of the modal clearly available)

- (28) In view of what the laws say, we should protest them. (deontic reading preferred, but not one relative to what the laws require)

Examples could be multiplied. It is not at all obvious that there is any canonical overt realization of the restriction of a modal. So there is no real obstacle here to the view that the modal in a construction like (25) is never a true epistemic.

Let me also defend against another sort of objection. Building on Horn [1989], Copley [2004] observes the following contrast:

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

One might think that this is evidence that *should* has a reading on which it is strictly weaker than epistemic *must*. But that would be much too fast. Observe:

- (30) a. #Xander must be there, in fact, he normally is.  
b. Xander is normally there, in fact, he must be there.
- (31) a. #Xander must be there, in fact, he usually is.  
b. Xander is usually there, in fact, he must be there.

From (31) it would obviously be a mistake to conclude that *normally* is on a scale with, and is strictly weaker than, epistemic *must*. Similarly for *usually*. A better explanation for the contrasts here would be to say that the *should/normally/usually*-claims serve to partially address some question under discussion, a question which is fully answered by the *must* claim. If we lead with a full answer (with the *must* claim), following this immediately with information only partially addressing the question is pragmatically odd.

Henceforth I proceed under the assumption that there is after all no epistemic reading of the weak necessity modals. The problem is now to further clarify the nature of the pseudo-epistemic reading. How should we characterize the examples so far reviewed? A rough step towards analysis, hinted at already above, would be to say this: they are cases where *normality* and *probability* come apart. The pseudo-epistemic use of weak necessity modals tracks, not probability, but something more like normality. Examples like the earthquake and Rasputin cases are ones where *what ought to be the case* lines up with *what would normally have been the case* rather than with *what probably is the case*. The probability that Jones is dead is zero, but as he observes, he ought to be dead. Examples like the marbles case are ones where *what ought to be the case* again lines up with *what would normally have been the case* rather than with either *what probably is the case* or with *what was likely to have been the case*.

## 4

I have recommended against Kratzer’s grouping of *probably* with the weak necessity modals, and against the idea of allowing that there is an epistemic reading of these modals. Still, some might suppose that that the apparatus of Kratzer [1981, 1991] yet has the resources to classify the pseudo-epistemic reading. Let me consider one possibility in this vein. First I briefly review the main features of Kratzer’s approach.

On Kratzer’s system, modals induce quantification over some partially ordered, restricted class of worlds. The relevant partially ordered, restricted class of worlds is generally fixed by two ingredients (what Kratzer calls *conversational backgrounds*): a *modal base* and an *ordering source*. Formally these are both functions from worlds to sets of propositions. The modal base fixes the restricted class of worlds quantified over, as a function of the evaluation world: this will just be the set of worlds making all the propositions delivered by the modal base true (the intersection of the propositions given by the modal base). The ordering source is then used to induce an order on the worlds in this domain.

Kratzer originally hypothesized that there were two principal kinds of modal base (corresponding roughly to the traditional linguistic distinction between epistemic and “root” modals): *epistemic* and *circumstantial*. Epistemic modals have epistemic modal bases. These deliver, relative to an evaluation world, a set of propositions known. Non-epistemic modals have circumstantial modal bases. As I discuss further below, these are somewhat harder to give a generic characterization of; but the basic idea is that given a world, they deliver a set of true propositions characterizing some contextually determined set of relevant circumstances at that world.

Kratzer further hypothesized that the two kinds of modal base are apt to be ordered by characteristically different kinds of ordering sources, and that these differences give rise to the various possible “flavors” of modal operators we observe. Circumstantial modal bases, when ordered at all, can combine with *deontic*, *teleological* (or *goal-oriented*), or *bouletic* ordering sources, and perhaps others. These ordering sources yield sets of propositions characterizing what is required, what is aimed for, what is wished for, and so on; and they induce a ranking of the circumstantial modal base worlds according to how well the worlds conform to what is required, aimed for, wished for, etc. Give or take some lexical idiosyncrasies, root modals are assumed to potentially take any of these flavors. Kratzer proposed that epistemic modal bases, by contrast, were apt to combine with either *doxastic* or *stereotypical* ordering sources, ranking the epistemically accessible worlds according to how well they conform to certain (contextually fixed) beliefs or to stereotypical situations, respectively. In

particular, a stereotypical conversational background yields, relative to a world, a set of propositions characterizing the normal course of events in that world.

I set aside the idea that there exists a specifically doxastic reading of epistemic modals, since there appears to be no evidence for this idea.<sup>7</sup> Stereotypical ordering sources play the more dominant role in the discussion of epistemic modals in Kratzer [1991]. On a natural reading of her account, stereotypicality is what primarily orders the epistemic domain (the worlds quantified over by epistemic modals): properties of stereotypical ordering sources are what grounds the differences between the grades of epistemic modality (*must* versus *probably*, and *probably* versus *might*).

Against this, I want to suggest that we work under the hypothesis that stereotypical ordering sources never play a role in the semantics of epistemic modals. Epistemic modals are not sensitive in any special way to facts of normality or stereotypicality. These modals are indeed sensitive to bodies of information, and a body of information may well be modeled by a set of worlds together possibly with some kind of ordering on those worlds. But the relevant ordering is not well-interpreted as anything like an ordering by stereotypicality or normality. The relevant notion is instead likelihood, broadly construed. I take it this suggestion is already motivated by the examples given above. The epistemic modality primarily concerns what might actually be the case, and what is likely or unlikely to be the case. It does not concern what is normally the case.

Of course, truths about what is normally the case may affect what you think might be the case. But truths about *anything* can affect what you think might be the case. The point is that epistemic modality does not have a special semantic connection to the notion of normality. Epistemic modals are not sensitive to normality orderings.<sup>8</sup>

---

<sup>7</sup>The idea that there exists a doxastic reading of epistemic modals is only mentioned once by Kratzer [1991], and in passing, without evidence. Portner [2009], following Kratzer, more explicitly maintains that there is such a reading of epistemic modals, but again, no evidence is supplied. Portner does supply a number of examples which, he suggests, illustrate modals interpreted with an epistemic modal base and a doxastic ordering source. However, this interpretation is never actually defended. For every example Portner supplies, we could hold that the relevant epistemic modal base is either not ordered, or that it orders itself, or that the order is a primitive feature of the relevant information state.

(Of course, epistemic modals under explicit belief operators arguably quantify over belief worlds (Stephenson [2007], Yalcin [2007]); but that is orthogonal to the present issue, which is one about what kinds of information can *order* the worlds quantified over. Moreover we are talking about what is possible for the interpretation of modals in abstraction from explicit semantic shifting.)

<sup>8</sup>Though there may be a connection between normality and epistemic modality at the level of default inference. See §7 below.

(To some, this will seem like hairsplitting about how to use the technical term ‘epistemic modal’. It is not. Epistemic modals are semantically distinctive in ways that set them apart from other modals in significant respects (see for instance Groenendijk et al. [1996], von Stechow and Iatridou [2003], Yalcin [2007, 2014], Anand and Hacquard [2013], Bledin [2014]). The evidence strongly favors the thesis that the epistemic modals form a natural class.<sup>9</sup> The thesis that *ought* and *should* do not admit of epistemic readings—readings that belong in this class—is substantive and nonterminological.)

Still, Kratzer was not wrong to suggest that some modals can be sensitive to facts of normality or stereotypicality. This was an important insight. Evidently, the weak necessity modals can be understood in this way: that is what I have already suggested about the pseudo-epistemic reading. But for reasons already reviewed, if we stay within the Kratzerian approach to modality, we do not want to interpret this reading of the modals as involving quantification over the epistemically accessible worlds (the worlds quantified over by epistemic modals); that is, we do not want to say that these modals take an epistemic modal base, and thus are epistemic modals. Again, this is because we do not want *ought*  $\phi$ /*should*  $\phi$  to have a reading on which it entails (epistemic) *might*  $\phi$ , or is entailed by (epistemic) *must*  $\phi$ . For that runs counter to our examples.

We have noted that Kratzer bifurcates the space of modal bases into epistemic and circumstantial varieties. So if we want to stay within the Kratzerian paradigm, we should say that the pseudo-epistemic reading of weak necessity modals corresponds to the combination of a circumstantial modal base with a stereotypical ordering source.

So for example, the quantification introduced by Jones’s modal operators is restricted to worlds that hold fixed certain circumstances obtaining in the actual world, circumstances like the following: a certain earthquake happened, Jones was in a building at the time, the building collapsed, etc. But we do not hold fixed all the circumstances known to obtain. In particular, we do not hold fixed the circumstance that Jones survived. Jones’s pseudo-epistemic modals must be allowed to quantify over a domain of worlds which include worlds where he was killed. By contrast, true epistemic modals do not allow this kind of thing. They do not allow this kind of selective deletion of items of information possessed in context. If it is a part of our information that Jones—against all appropriate expectations—survived, then it is just not the case that Jones might be dead.<sup>10</sup>

<sup>9</sup>I favor the thesis that the puzzles discussed in Yalcin [2007, 2014] are diagnostic of epistemic modals.

<sup>10</sup>Can’t we say something like, “In light of the information in the report, Jones must be dead. But he isn’t—he’s right there!”—thereby selectively deleting whatever items of information entailed that Jones is dead? But I take it that an operator like *In light of the information in the*

## 5

But the “circumstantial modal base/stereotypical ordering source”-analysis faces some difficulties. Relevant for assessing it is the question whether and what other modals can take the pseudo-epistemic reading. Consider first the strong necessity modals. Suppose Jones had instead said:

(32) ? I must be dead right now.

(33) ? I have to be dead right now.

The epistemic and deontic readings of these modals are false, owing to the facts of the context; so we might naturally expect the pseudo-epistemic readings of these modals to shine through here. Surely Jones *is* dead in all the relevant normal (circumstantially accessible) worlds—not just, say, those that are especially normal according to the relevant normality ordering (as Kratzerians might naturally analyze *ought* and *should*<sup>11</sup>). But these constructions are plainly marked. There are two possibilities: (a) *must* and *have to* cannot take the pseudo-epistemic reading; or (b) on the pseudo-epistemic reading, the strong necessity modals entail their prejacent.

Either proposal seems unexpected from the point of view of the Kratzerian analysis just described. From a theoretical point of view, (b) is not motivated. If the set of worlds that pseudo-epistemic *have to* and *must* quantify over is some set of worlds that are normal relative to the evaluation world, there is no conceptual reason that the evaluation world needs to be in this set. The world can be a very abnormal place, even relative to its own standards of normality.

One might propose that the strong necessity modals simply universally quantify over the circumstantial modal base worlds, and not over any further restriction of that modal base by normality or stereotypicality. Since the actual world must be in this set (the circumstantial worlds), this would indeed compel the

---

*report* is a (hyper)intensional environment which semantically shifts the state of information relevant for the epistemic modal under it (see, e.g., [Hacquard \[2006\]](#), [Yalcin \[2007\]](#), [Stephenson \[2007\]](#), [Anand and Hacquard \[2013\]](#)). The claim made in the main text concerns bare epistemic modals sentences where no compositionally induced semantic shifting takes place. There is no tension here.

<sup>11</sup>See for instance [von Stechow and Iatridou \[2008\]](#), who construe weak necessity modals as having a strictly smaller domain of quantification than the corresponding strong necessity modals: “strong necessity modals say that the prejacent is true in all of the favored worlds, while weak necessity modals say that the prejacent is true in all of the very best (by some additional measure) among the favored worlds” (119). They implement this idea in Kratzerian framework by hypothesizing that weak necessity modals generally involve two ordering sources rather than one (with the second serving to impose a further restriction not present in the case of the corresponding strong necessity modals). See also [Rubinstein \[2012\]](#).

strong necessity modals to entail their prejacent. Note this is, in effect, to accept (a): while weak necessity modals can get a special pseudo-epistemic normality reading, strong necessity modals would only get a pure circumstantial reading.

Option (a) is also *prima facie* unattractive from the point of view of the basic Kratzerian framework. In the context of this account, there is no principled reason why there should not be a pseudo-epistemic reading of the strong necessity modals—in the way that the system predicts that there are, for example, deontic readings of both weak and strong necessity modals. It is a familiar point that often the most important evidence for linguistic theory is negative evidence—the unavailability of certain readings or constructions. As we have been observing throughout, modals are in fact highly restricted in the readings that they can take. The absence of a pseudo-epistemic reading for the strong necessity modals requires explanation in the context of Kratzer’s theory.

There are further sources of discomfort with option (a). We are considering the hypothesis that the pseudo-epistemic reading corresponds to a circumstantial modal base and a stereotypical ordering source. Other things being equal, then, a pure circumstantial reading of a strong necessity modal should entail a sentence with a pseudo-epistemic weak necessity modal. But the intuitive evidence for this entailment is slim. Consider one of Kratzer’s examples of a case of a strong necessity modal getting a pure circumstantial reading:

(34) I must sneeze. (pure circumstantial)

(35) I have to sneeze. (pure circumstantial)

Neither of these intuitively entails either of:

(36) ? I ought to sneeze. (pseudo-epistemic)

(37) ? I should sneeze. (pseudo-epistemic)

One might reply that the pseudo-epistemic weak necessity sentences presuppose (or entail) that their prejacent are false in some of the relevant circumstantial worlds, and this blocks the entailment. Whether or not this suggestion is artificial depends on whether this kind of constraint applies also to other readings of weak necessity modals. It seems not to be widely assumed in the deontic case, where strong necessity modals are generally thought to entail the weak ones (indeed, that is where the jargon of “strong” versus “weak” comes from).

In any case, a further problem is that it is just strange to say these in the relevant kind of context. Why? Suppose you sense a sneeze coming on, but you are not convinced it is inevitable. Why don’t (36) and (37) seem like

natural words of warning, slightly weaker than the warning conveyed by (34) and (35)? (As deontic *ought* is thought to be weaker than deontic *must*.) We have the intuition that the flavor of modality is qualitatively different—not just weaker—when we move from the strong necessity modals here to the weak ones. This is surprising. Even if the pseudo-epistemic readings of (36) and (37) are marked without some additional setup in this kind of scenario, we naively might have thought that there should be a pure circumstantial reading of *ought* and *should* available, such that (36) and (37) can be appropriate when you feel a sneeze approaching, in the way (34) and (35) are. But such a reading seems not to be available. This requires explanation.<sup>12</sup>

The evidence so far reviewed suggests—tentatively—the following: strong necessity modals do not have a pseudo-epistemic reading, and weak necessity modals do not have a pure circumstantial reading. Further, it is not clear whether circumstantial strong necessity modals entail the pseudo-epistemic weak necessity modals.

We should wrap up this discussion of the way that circumstantial modality might interact with the pseudo-epistemic modality by briefly considering the situation with possibility modals. Kratzer [1991] suggests that *can* can take a pure circumstantial reading. If the pseudo-epistemics involve a circumstantial modal base, it would therefore be natural to expect pseudo-epistemic modal sentences to entail sentences with pure circumstantial possibility modals. So suppose Jones says:

(40) ? I can be dead right now.

This sounds like an strangely worded offer, from Jones, to kill himself.

Though odd owing to the facts of the particular example, (40) arguably does have a reading on which it is strictly speaking true (and on which the entailment from the corresponding pseudo-epistemic weak necessity claim goes through). The modal here takes an ability-like reading, which some would run together with, or closely tie to, the circumstantial reading. Interestingly, however, this

---

<sup>12</sup>The pseudo-epistemic readings of (36) and (37) also do not come very naturally, but they do at least seem clearly available. It is easier to hear if we fill in some special circumstances—for example, that the speaker has just inhaled some sneezing powder.

Seemingly also relevant to the availability of the pseudo-epistemic reading is the aspectual structure of these constructions. Notably, in the presence of a telic verb, the pseudo-epistemic reading is rather more natural with a progressive infinitive, or with some explicit temporal marking:

(38) I ought to be sneezing.

(39) I should sneeze any minute now.

possibility modal (like *may*) does not seem capable of being used to express the (true, contextually appropriate) thought that among the relevant *normal* worlds is a world where Jones is dead. Many speakers can achieve something like that reading with:

(41) I could be dead right now. (pseudo-epistemic?)

This tends as a default to get a pure epistemic reading (on which it is marked), but (especially with the right tone of amazement) it can receive the pseudo-epistemic reading, or something intuitively close to it.<sup>13</sup> This reading is heard as equivalent to something along the lines of:

(42) I could have been dead right now.

(43) I might have been dead right now.

—where these are not expressing the philosopher’s mere metaphysical possibility, but rather seem instead to effect existential quantification over some nearby normal worlds.<sup>14</sup>

Where does this leave us? The idea that the pseudo-epistemic reading involves a circumstantial modal base and a stereotypical ordering source might yet be massaged into working. I have raised some questions about it, but I have not shown that it cannot be done. A successful development would presumably say more about the character of circumstantial modal bases. It is fair to say that there is not a tremendous amount of clarity about the nature of circumstantial readings in the literature. Kratzer herself now seems ambivalent, if not skeptical, about clearly defining this category from a semantic point of view (see e.g. [Kratzer, 2012, 24]). Still, it may be possible to do, for all I have said.

Nevertheless, in the following sections, I wish to explore a different line of analysis.

## 6

In light of the examples discussed above, I want to consider the idea that on the pseudo-epistemic reading, *ought*  $\phi$  and *should*  $\phi$  are rather directly connected

---

<sup>13</sup>Select speakers can hear (6) in this way as well.

<sup>14</sup>Though it should be acknowledged that many speakers would tend to reach instead for a construction with a telic predicate, as in:

(44) I could have been killed.

(45) I might have been killed.

to something about the way things normally unfold. As a first attempt, let us consider the following idea:

It should/ought to be that  $\phi \approx$  Normally,  $\phi$

We could also put the idea in a way more aligned with the fact that these modals generally take infinitival clauses:

FIRST PASS

$\alpha$  should/ought to  $F \approx$  It is normal for  $\alpha$  to  $F$

(It is of course highly doubtful that we will be able to exactly paraphrase pseudo-epistemic *ought* and *should* somehow in terms of the word *normal*. Still, exploring how close we can get may teach us something.)

The following examples seems to illustrate a *prima facie* mutual entailment, and hence motivate FIRST PASS:

(46) Normally, Bob is in his office right now.

It is normal for Bob to be in his office right now.

(47) Bob should be/ought to be in his office right now.

A difficulty for the FIRST PASS is that Rasputin-like cases does not fit this paradigm as nicely:

(48) Rasputin should be/ought to be dead now.

(49) ? Normally, Rasputin is dead now.

(50) ? It is normal for Rasputin to be dead now.

*Normally*  $\phi$  seems to suggest that  $\phi$  corresponds to some kind of repeatable event, or a situation of a repeatable type; but, at least out of context, Rasputin's death does not strike us as such an event or situation. Yet the pseudo-epistemics are (as already noted) fine here, even without any special setup. Similarly for the earthquake case:

(51) ? It is normal for me to be dead now.

The sentences are improved if we combine a normality operator with the sort of morphology usual to counterfactuals. We can say:

(52) Normally, Rasputin would be dead now.

(53) It would be normal for Rasputin to be dead now.

(54) It would be normal for me to be dead now.

We might therefore adjust the proposal:  $\alpha$  *ought to F* and  $\alpha$  *should F*, on their pseudo-epistemic readings, mean something roughly equivalent to:

SECOND PASS

$\alpha$  should/ought to  $F \approx$  It would be normal for  $\alpha$  to  $F$

(This brings to mind [von Fintel and Iatridou \[2008\]](#)'s observation that crosslinguistically, *ought* is often expressed by combining a strong necessity modal with counterfactual morphology. It also brings to mind Swanson's suggestion that this kind of *should* concerns what "one would naturally expect" to be the case ([\[Swanson, 2008, 1024\]](#).)

The following seems to be a problem for this approach. On a Monday at the office, I can say:

(55) Bob normally doesn't come into the office on Mondays; he stays home. So it would be normal for Bob to be home now. But I see that his car parked outside. So he ought to be in his office.

But according to the second pass, there should be a problem here: the sentence *It would be normal for Bob to be home now* and the sentence *He ought to be in his office* should be in conflict. But evidently they are not. We can affirm them both.

What to say? Perhaps the following: in (55), the *normal* is talking about: *normal for a Monday*. By contrast, I tentatively conjecture that the pseudo-epistemic modality involves a notion of normality which is in some sense "all relevant things considered". The *ought* in (55) concerns what is normally true on a Monday where Bob's car is parked outside—and perhaps factoring in further relevant contextually supplied information, too. Although Bob is not normally in his office Mondays, it may yet be true that Bob is normally in his office on those Mondays in which his car is parked outside. Perhaps this latter fact, or something in the near vicinity, explains why we accept 'He ought to be in his office' in (55).

This leaves us with something like:

THIRD PASS

$\alpha$  should/ought to  $F \approx$  It would be normal, all relevant things considered, for  $\alpha$  to  $F$

Of course, we should like to say much more about what makes for normality in this sense. This seems to be a difficult problem, the surface of which I have

barely scratched. The only point I wish to make here is that there seems to be no reason why this kind of reading of *ought* and *should* could not be captured in terms of simple universal quantification over a domain of normal worlds—the worlds where, roughly, things in fact unfold as they (all relevant things considered) normally should. Tentatively embracing this possibility, I will now switch to calling the pseudo-epistemic reading the *normality* reading. I take it this reading has a normative quality. Roughly it expresses something about what one is, or would be, entitled to expect to be true.

## 7

Pseudo-epistemic modality now distinguished from epistemic modality, I will try to take some steps towards clarifying what exactly the relationship is between these modalities. In the next two sections, I explore two possible points of contact. First, it may be that the domain of normal worlds is fixed as a function of the set of worlds relevant for epistemic modals. Second, there may be important relations of default inference from the pseudo-epistemic modality to the epistemic modality.

Start with the first idea. Even if *ought* and *should*, on their normality readings, are taken to quantify universally over a domain of normal worlds, it may yet be that this domain of normal worlds is one that is somehow determined as a function of the set of worlds that epistemic modals quantify over. We could probe this question by looking to see whether semantically shifting the body of information relevant for epistemic modals affects judgments about the truth of *oughts* and *shoulds* of normality. Plausibly, the indicative conditional is a device for semantically shifting this body of information (for arguments, see Gillies [2004], Yalcin [2007, 2012c], Kolodny and MacFarlane [2010]; also Stalnaker [1975], Heim [1983]). Crudely, with an indicative  $\phi \rightarrow \psi$ , we consider temporarily updating our information with  $\phi$ , and check to see whether the information corresponding to  $\psi$  is incorporated therein. It appears that normality modals are indeed sensitive to such shifts, as it seems we can consistently affirm both of the following on the normality reading:

(56) Noam ought to be in his office.

(57) If Noam was arrested on the way to work, he ought to be in jail.

It would be a mistake to draw the modus tollens inference here, concluding that Noam was not arrested. Similarly with:

(58) My car should be parked on the street outside.

(59) If my car was stolen last night, it should be in a chop shop by now.

We find a structurally analogous situation with the deontic reading of *ought* and *should*—or so argue Kolodny and MacFarlane [2010], considering parallel examples for the deontic reading. Such examples suggest that *ought*  $\phi$  and *should*  $\phi$  are not *persistent* in the following sense: if *ought/should*  $\phi$  is accepted relative to some state of information, it does not necessarily follow that it is accepted relative to any strictly stronger state of information. You are of the view that Noam ought to be in his office. But if you strictly grew your state of information, keeping all your factual information and adding to it the proposition that Noam was arrested on the way to work, it may well no longer be the case that you accept that Noam ought to be in his office. In at least this way, normality *oughts* and *shoulds* seem not to be straightforwardly factual.

In thinking about the non-persistent character of modalities of normality, we do well to study Veltman [1996]. Veltman explores the possibility of a semantics for *normally* which licenses default inferences from claims about what is normally true to claims about (*inter alia*) what is presumably true. Following in his path, we could ask analogous questions concerning our target modals. For example, the following kind of inference seems to be licensed *as a default*:

Noam ought to be in his office. (normality reading)

$\Rightarrow$  Presumably, Noam is in his office.

That is: if the *only* relevant information you have is that Noam ought to be in his office, then you are in position to judge: *presumably, Noam is in his office*. But this conclusion is only licensed as a default. If one acquires further information of the appropriate sort, this can defeat the default, rendering the conclusion no longer derivable. Thus the following pattern is not valid, even though it strictly extends the premises of the above argument:

Noam ought to be in his office. (normality reading)

Noam is not in his office.

$\Rightarrow$  Presumably, Noam is in his office.

For this we would obviously require a nonmonotonic consequence relation. A default semantics for normality modals in the style of Veltman [1996] could enable us to predict these patterns of entailment. Let me illustrate, using the tools Veltman provides to extend the kind of semantic system developed in Yalcin [2007] for treating epistemic modals.<sup>15</sup>

<sup>15</sup>Veltman [1996] deploys his tools in the analysis of *normally*. I will apply many of his formal ideas to the analysis of normality *ought* and *should* instead. I will assume nothing about the correct semantics for *normally*.

We have an intensional semantics with points of evaluation which are pairs of worlds and what I will call *expectation-laden information states*. The extensions of expressions are relativized to such points. An expectation-laden information state is modeled as a pair of a set of worlds  $s$  (the information state) and an *expectation pattern*  $\succeq$ . The information state component reflects a possible stance on what the facts are. The expectation pattern component reflects a view on what is normal. It is a preorder (reflexive transitive order) on the domain  $\mathcal{W}$  of possible worlds, intuitively readable as *is as at least as normal as*. We will think of the preorder as induced by a finite set of propositions  $N$  (the *normality propositions*) taken to characterize a view about what is normally the case:  $w \succeq v$  iff every proposition in  $N$  true at  $v$  is true also at  $w$ .<sup>16</sup> (Thus strictly we should write  $\succeq^N$ , but we generally leave the reference to the normality propositions tacit.) We assume  $N$  always contains the logical truth  $\mathcal{W}$ . We will call a world *max normal* if it is in  $\bigcap N$ ; i.e.:

**Def 1.** A world  $w$  is **max normal** relative to an expectation pattern  $\succeq$ , or  **$\succeq$ -normal**, iff for all  $w' \in \mathcal{W}$ ,  $w \succeq w'$ .

We lay down the following definitions:

**Def 2.** The **normal set** relative to an expectation pattern  $\succeq$ ,  $n_\succeq$ , is the set of  $\succeq$ -normal worlds. (The normal set relative to an expectation-laden information state  $i$ ,  $n_i$ , is just the normal set relative to  $i$ 's expectation pattern.)

**Def 3.** An expectation pattern  $\succeq$  is **coherent** iff  $n_\succeq \neq \emptyset$ .

(We assume that in the non-defective case, expectation patterns are coherent.) Now let  $O$  symbolize *ought* and *should* on their normality readings. Then we could consider the idea that  $O\phi$  expresses a constraint on the expectation pattern of an information state, namely that the prejacent  $\phi$  be true throughout the normal set it determines:

$$\llbracket O\phi \rrbracket^{w,i} = 1 \text{ iff } \forall w' \in n_i : \llbracket \phi \rrbracket^{w',i} = 1$$

It is instructive to compare this to the clauses for the epistemic possibility ( $\diamond$ ) and necessity ( $\square$ ) modals (defended in Yalcin [2007]):

$$\llbracket \diamond\phi \rrbracket^{w,i} = 1 \text{ iff } \exists w' \in s_i : \llbracket \phi \rrbracket^{w',i} = 1$$

$$\llbracket \square\phi \rrbracket^{w,i} = 1 \text{ iff } \forall w' \in s_i : \llbracket \phi \rrbracket^{w',i} = 1$$

---

<sup>16</sup>According to von Fintel and Iatridou [2008], “epistemic *ought* differs from epistemic *must/have to* in being sensitive not just to the hard and fast evidence available in a situation but also to a set of propositions that describe what is *normally* the case” (119). While I agree that it is useful to appeal to a set of normality propositions, the idea of taking the domain of quantification for pseudo-epistemics to be a subset of the epistemically accessible worlds is problematic, for reasons already reviewed.

These modals are sensitive to the information state component, but not the expectation pattern. We can observe the failure of  $O\phi$  to entail  $\diamond\phi$  by defining an appropriate notion of consequence:<sup>17</sup>

**Def 4.**  $\phi$  is **incorporated** (or **accepted**, or **supported**) in an expectation-laden information state  $i$  iff for all  $w$  in  $s_i$ :  $\llbracket\phi\rrbracket^{w,i} = 1$ .

**Def 5.**  $\psi$  is an **informational consequence** of a set of sentences  $\Gamma$ ,  $\Gamma \models \psi$ , just in case: for all information states  $i$  that incorporate every  $\phi \in \Gamma$ ,  $\psi$  is incorporated in  $i$ .

It is not hard to see that  $O\phi \not\models \diamond\phi$ . For  $O\phi$  can be incorporated in an expectation-laden information state that rules out  $\phi$  as candidate for actuality (as an epistemic possibility). One can of course have the view that the world one inhabits is not among the most normal possibilities. An expectation-laden information state reflects such a view when the set of worlds left open by the information state is disjoint from the normal set determined by its expectation pattern. In this kind of case, the information one has rules out the possibilities that are max normal.

Still, within any given information state, there will generally be differences in normality among the worlds left open. It is useful to define the class of worlds that are most normal within the set of worlds an expectation-laden information state leaves open:

**Def 6.** A world  $w$  is **optimal** relative to an information state  $i$ , or  **$i$ -optimal**, iff  $w \in s_i$  and there is no  $w' \succeq_i w$  such that  $w' \succeq_i w$ .

**Def 7.** The **optimal set** relative to an information state  $i$ ,  $o_i$ , is the set of  $i$ -optimal worlds.

Using this notion of optimality, we can give a semantics for *presumably* (P) following in Veltman's footsteps:<sup>18</sup>

$$\llbracket P\phi \rrbracket^{w,i} = 1 \text{ iff } \forall w' \in o_i : \llbracket \phi \rrbracket^{w',i} = 1$$

Observe that  $P\phi \models \diamond\phi$ , since the optimal set will be a nonempty subset of any nonempty information state.<sup>19</sup> Since as recently noted,  $O\phi \not\models \diamond\phi$ , we know

<sup>17</sup>For further discussion of this notion of consequence, see Bledin [2014].

<sup>18</sup>This semantics is surely only an approximation. The restricted embedding potential of *presumably* suggests that it is semantically more like a discourse hedge or parenthetical than an ordinary sentential operator. We are following Veltman is abstracting from this.

<sup>19</sup>The right result, intuitively. For example:

(60) Presumably, Noam is in his office.  
 $\Rightarrow$  Noam might be in his office.

(61) ? Noam is not in his office, but presumably he is in his office.

$O\phi \neq P\phi$ . That is the desired result. We would not approve of Jones's reasoning like this:

I should be dead right now.  
 $\Rightarrow$  Presumably, I am dead now.

Still, as we noted above, the movement from a normality *should/ought*-claim to a claim about what is presumably true plausibly has default license. The intuitive problem with the above reasoning is that in Jones's context, the epistemic possibility that he is dead is eliminated; consequently the corresponding *presumably*-claim is not licensed. To consider what patterns of inference are default licensed, we should define a notion of consequence that restricts to states of information incorporating the premises and the premises only—states that are informationally minimal. We may do this as follows:

**Def 8.** The **minimal states**  $i$  incorporating a set of sentences  $\Gamma$  are the states  $i$  such that  $i$  incorporates every element of  $\Gamma$ , and there is no  $i'$  such that (a)  $i'$  incorporates every element of  $\Gamma$ , and (b)  $s_i \subset s_{i'}$ .

**Def 9.**  $\psi$  is an **default consequence** of a set of sentences  $\Gamma$ ,  $\Gamma \vDash_D \psi$ , just in case all the minimal states incorporating  $\Gamma$  incorporate  $\psi$ .

Now we can observe that although  $O\phi \neq P\phi$ ,  $O\phi \vDash_D P\phi$ . For example, return to:

Noam ought to be in his office.  
 $\Rightarrow$  Presumably, Noam is in his office.

The minimal states incorporating the premise are those where the information state component is just  $W$ . These are all the states  $i$  such that the proposition that Noam is in his office true throughout  $n_i$ . At these states  $n_i$  coincides with  $o_i$ , so the inference is licensed.

Observe next that although  $O\phi \vDash_D P\phi$ ,  $\{O\phi, \neg\phi\} \not\vDash_D P\phi$ . (So  $\vDash_D$  is non-monotonic.) For example, the following is not default licensed:

Noam ought to be in his office.  
 Noam is not in his office.  
 $\Rightarrow$  Presumably, Noam is in his office.

The minimal states incorporating the premises are those where the information state component is just the set of worlds where Noam is not in his office. This means at each of these states  $i$ ,  $o_i$  is some subset of these worlds. The conclusion

drawn is not incorporated by any of these states. (On the contrary, they all incorporate ‘Presumably, Noam is not in his office’—intuitively a correct result.)

In this paper we have seen many reasons to doubt that the following is a valid inference:

Noam ought to be in his office.  
⇒ Noam is probably in his office.

But now we can ask a different question. Is this inference *default* valid? That idea is considerably more plausible. If the only relevant information one has is the premise, the conclusion does seem like an acceptable one to (tentatively) draw. Vindicating this pattern as a default inference would, it seems, do much to accommodate the intuition that there is some nontrivial connection between what normally ought or should be and what is likely. (And correspondingly, it would further subtract from whatever motivation remains for recognizing a bonafide epistemic reading of *ought* and *should*.)

There are various ways one might vindicate this pattern as a default inference. It depends on what semantics for probability operators one assumes. One possibility is to think of information states as equipped with a probability measure  $Pr$  conditionalized on the state (as discussed for instance in Yalcin [2012c]). Then one could require that in minimal states  $i$  incorporating only information about normality,  $Pr_i(n_i) > .5$  (or whatever threshold is thought to be appropriate to license *probably*-claims in general).<sup>20</sup> This would be to say that in abstraction from factual information, it is likely that the world is normal.

## 8

We have yet to say what is going on with:

(56) Noam ought to be in his office.

(57) If Noam was arrested on the way to work, he ought to be in jail.

Above we noted some theories on which indicative if-clauses are understood to shift the information state relevant for evaluating the consequent of the conditional. Let me give a basic example of this kind of semantics. I will restrict attention to indicative conditionals with nonmodal antecedents.

**Def 10.** For nonmodal sentences  $\phi$ , the **proposition expressed by**  $\phi$ ,  $[\phi]$ , is  $\{w : \llbracket \phi \rrbracket^{w,i} = 1\}$  (for arbitrary choice of  $i$ ).

---

<sup>20</sup>For further discussion of the appropriate threshold, see Yalcin [2010].

**Def 11.** For nonmodal  $\phi$ , the **nearest state to**  $i = \langle s, \geq \rangle$  **incorporating**  $\phi$ ,  $i + \phi$ , is the pair  $\langle s_\phi, \geq \rangle$ , where  $s_\phi$  is  $s \cap [\phi]$ .

Then a simple information-sensitive semantics for the indicative conditional  $\rightarrow$  would be:

$$\llbracket \phi \rightarrow \psi \rrbracket^{w,i} = 1 \text{ iff } i + \phi \text{ incorporates } \psi$$

We can briefly note that this helps us to model the information-sensitivity of epistemic modals. For example:

(56) It is not the case that Noam must be in his office.

(57) If Noam is in the building, he must be in his office.

One can of course rationally accept both of these (and without also accepting the seeming modus tollens conclusion that Noam is not in his office). In the present context what we can say is that the if-clause shifts the sets of worlds that the epistemic necessity modal in (57) quantifies over to a strictly stronger state of information, one incorporating the antecedent information. See Yalcin [2007, 2012c], Gillies [2010] for further discussion.

Now the thing to observe is that on the semantics so far provided, normality *oughts* and *shoulds* are not information-state sensitive at all; rather, they are expectation-pattern sensitive. In the context of this account, the putative information-shifting feature of indicative if-clauses therefore does not help us explain the data in (56) and (57).

There are at least two possibilities. One possibility is that the expectation pattern of an expectation-laden information state is in fact determined as a function of the information state. On this view, if we strictly grow the information state, the expectation pattern systematically changes as result, and (hence) so do the *oughts* and *shoulds*. (Compare Kolodny and MacFarlane [2010], Cariani [2013], Charlow [2013], Silk [2013] on deontic modals.) Another possibility is that the if-clause shifts the the expectation pattern “directly”, and not indirectly via its effect on the information state. On this approach, we should change the semantics for the indicative by showing how to shift from a given  $\geq$  to a new preorder  $\geq_\phi$ , one which minimally adjusts  $\geq$  somehow as a function of the antecedent information  $[\phi]$ .

Let me give one abstract, nonlinguistic reason for preferring the former approach. In modeling modals of normality, we are indirectly modeling states of mind that consist partly in a view about what the normal course of events consists in, about what ought to be or should be so in roughly the sense of what one would be entitled to expect to be true. Plausibly (but not uncontroversially), a

view about what ought to be so in this sense does not reduce to a purely factual view about how things are: that is why we appeal to expectation patterns over and above information states. Nevertheless, our views about what (normally) ought to be or should be the case do not float free from what we take the facts to be. We take there to be a relation of supervenience between the pattern of expectation we endorse and the way we take the world to be. We take our views about what we are entitled to expect to be fixed by the facts. One's expectations, we take it, should not change unless there are corresponding changes in how one takes the world to be. (Cf. [Gibbard, 2003, Ch. 5], discussing normative discourse.)

We could put it like this: one component of one's state of mind determines an *expectation function*  $f$ , which maps information states to expectation patterns. This reflects one's conception of how the information one possesses governs the expectations it would be appropriate to have. The important feature of the expectation function is that it is a function: given an information state, we do not take there to be multiple, equally legitimate expectation patterns. Disagreements between agents about what (normally) ought or should be the case may trace to differences in the ways that each agent takes the world to be (their information); but they may also trace to a more fundamental difference in the expectation function each agent endorses.

Taking this approach, we would do well to model expectation-laden information states as pairs  $\langle s, f \rangle$  of an information state and an expectation function. The expectation pattern of an expectation-laden information state  $i$  is then just  $f(s)$ . This will determine the normal set for the state, and the semantics will be as above. In principle, this approach can capture the shiftiness apparent in (57), as the expectation pattern will now shift with the information.

It remains to explore what formal constraints (if any) govern  $f$ . The matter deserves separate exploration. But I note that one constraint we surely do not want is the requirement that the normal worlds determined by a given state  $\langle s, f \rangle$  be a subset of  $s$ . This would make it the case that  $O\phi \models \Diamond\phi$ —the incorrect result, as we have repeatedly observed.

## 9

It is about time to say something about the deontic readings of *ought* and *should*. A natural and simplifying thought would be that the deontic and pseudo-epistemic readings of *ought* and *should* have some nontrivial amount of underlying formal structure in common. That after all would help explain why the same linguistic expressions are used for both, and in many languages.

In that spirit, the preceding two sections, if on the right track, would lead us to suspect that the concepts of default logic would be of some use in theorizing about deontic modals and deontic modality.

Indeed they seem to be. Horty [2012] in particular provides some illuminating analyses of deontic *oughts* using the resources of default logic. On his theory (as I would describe it), deontic *oughts* express constraints on default theories. Structurally, default theories are akin to expectation-laden information states, except that (roughly) the expectation pattern component is replaced by a set of *default rules* together with a preordering of the rules by their relative priority. To a first approximation, a default rule  $X \rightarrow Y$  is the sort of thing that recommends the transition from proposition  $Y$  from proposition  $X$ , as a default. Default rules can be used to encode what conclusions are epistemically default licensed given some information. (Rough example: the default rule corresponding to *Birds fly* default recommends the conclusion that  $x$  flies, given the information that  $x$  is a bird.) Such rules can also be used to capture which actions are default licensed given some information. (Rough example: the default rule corresponding to *Keep your promises* recommends that I show up, given that I promised to.) Horty uses the latter application to clarify deontic *oughts*. To say  $\phi$  ought to be the case is to say that the realization of  $\phi$  is default recommended, in a sense formalized with the notion of a default theory.<sup>21</sup> Together with the work above, this would seem to point in the direction of a unified treatment of deontic and pseudo-epistemic *oughts* as devices for expressing defaults. Exploration of this possibility emerges as a promising avenue to pursue.

While Horty's discussion of *oughts* is trained chiefly on the deontic reading and on surrounding issues about practical reasons and deliberation, it is plain he means for his style of analysis to apply to the notion of an epistemic reason and to theoretical deliberation, and to what we have been calling pseudo-epistemic *oughts*. (Not surprising, since perhaps the standard application of default logic is to theoretical reasoning.) From a linguistic point of view, it would be worthwhile to investigate the connections between a semantics for pseudo-epistemic *oughts* of the sort sketched above and one based on default theories of Horty's variety—to see which differences are notional and which substantive, and where the data point in the cases where they come apart. I leave this for future work.

Stepping back, I would be inclined to take Horty's analysis, and the analysis of this paper, in what I would call an expressivistic direction, in the style of Yalcin [2012a]. If it is right that we need something like expectation-laden information states or default theories to model the language of *oughts* and *shoulds*, it

---

<sup>21</sup>Actually, he formalizes several possible senses, discussing the prospects for each.

needn't follow that this talk serves to literally describe expectation-laden states of mind, or the default we embrace, or our habits of belief or action. Unembedded, these kinds of sentences do not describe the defaults we embrace; rather, they serve to express those defaults. Yalcin [2012a] offers a recipe for making precise of this sense of "express".

Defaults, as formalized here or by Horty, might also sensibly be called *norms*. Characteristic of norms, after all, is their defeasibility. Switching terminology, then, the picture that emerges is that *ought* and *should* correspond, in a relatively technical sense, to modalities of norms, on both their deontic and pseudo-epistemic readings.

## References

- Pranav Anand and Valentine Hacquard. Epistemics and attitudes. *Semantics and Pragmatics*, 6:8–1, 2013.
- Justin Bledin. Logic informed. *Mind*, 130(490):277–316, 2014.
- Fabrizio Cariani. 'Ought' and resolution semantics. *Noûs*, 47(3):534–558, 2013.
- Nate Charlow. What we know and what to do. *Synthese*, 190(12):2291–2323, 2013.
- Bridget Copley. So-called epistemic *should*. *Snippets*, 9:7–8, 2004.
- Bridget Copley. What should 'should' mean? <http://halshs.archives-ouvertes.fr/docs/00/09/35/69/PDF/copley.should.pdf>, January 2006.
- Stephen Finlay. *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press, 2014.
- Kai von Fintel and Sabine Iatridou. Epistemic containment. *Linguistic Inquiry*, 34(2): 173–98, 2003.
- Kai von Fintel and Sabine Iatridou. How to say *ought* in Foreign: The composition of weak necessity modals. In *Time and Modality*, pages 115–141. Springer, 2008.
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, 2003.
- Anthony Gillies. Epistemic conditionals and conditional epistemics. *Noûs*, 38(4):585–616, 2004.
- Anthony Gillies. Iffiness. *Semantics and Pragmatics*, 3:1–42, 2010. ISSN 1937-8912.
- Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and modality. In S. Lappin, editor, *Handbook of Contemporary Semantic Theory*. Blackwell, 1996.
- Valentine Hacquard. *Aspects of Modality*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2006.

- Irene Heim. On the projection problem for presuppositions. In M. Barlow, D. Flickinger, and M. Wescoat, editors, *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*, pages 114–25. Stanford University, 1983.
- Wesley Holliday and Thomas Icard. Measure semantics and qualitative semantics for epistemic modals. In *Proceedings of SALT*, volume 23, pages 514–534, 2014.
- Laurence Horn. A natural history of negation. 1989.
- John Horty. *Reasons as Defaults*. Oxford University Press, 2012.
- Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, CVII(3): 115–43, March 2010.
- Angelika Kratzer. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1:337–55, 1977.
- Angelika Kratzer. The notional category of modality. In H.-J. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts, New Approaches to Word Semantics*, pages 38–74. Walter de Gruyter, Berlin, 1981.
- Angelika Kratzer. Modality. In A. von Stechow and D. Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*, pages 639–50. de Gruyter, 1991.
- Angelika Kratzer. *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press, 2012.
- Daniel Lassiter. *Measurement and Modality: The Scalar Basis of Modal Semantics*. PhD thesis, PhD dissertation, New York University, 2011.
- Paul Portner. *Modality*. Oxford University Press, USA, 2009.
- Aynat Rubinstein. *Roots of Modality*. PhD thesis, University of Massachusetts, Amherst, 2012.
- Alex Silk. Evidence sensitivity in weak necessity deontic modals. *Journal of Philosophical Logic*, pages 1–33, 2013.
- Robert Stalnaker. Indicative conditionals. *Philosophia*, 5, 1975.
- Tamina C. Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.
- Eric Swanson. Modality in language. *Philosophy Compass*, 3(6):1193–1207, 2008.
- Judith Jarvis Thomson. *Normativity*. Open Court Publishing Company, 2008.
- Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3): 221–61, 1996.

- Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–37, 2010.
- Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, CXII(2): 123–160, 2012a.
- Seth Yalcin. Context probabilism. In Maria Aloni, Vadim Kimmelman, Floris Roelofsen, Galit Weidman Sassoon, Katrin Schulz, and Matthijs Westera, editors, *Proceedings of the Eighteenth Amsterdam Colloquium*, volume 7218 of *Lecture Notes in Computer Science*, pages 12–21. Springer, 2012b.
- Seth Yalcin. A counterexample to Modus Tollens. *Journal of Philosophical Logic*, 41: 1001–1024, 2012c.
- Seth Yalcin. Epistemic modality de re. Unpublished manuscript, 2014.