NORTHWESTERN UNIVERSITY

Modals and their Arguments

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Philosophy

By

Daniel Skibra

EVANSTON, ILLINOIS

December 2020

Abstract

Modals and their Arguments

Daniel Skibra

The philosophical literature on modals is dominated by the following paradigm: modals are modeled as quantifiers over sets of possible worlds. The diversity of modal "flavors" (e.g., epistemic, deontic, teleological interpretations of modals) is accommodated within the paradigm by logical mechanisms that allow extralinguistic factors to restrict the quantificational domain of the modal. This way of modeling flavor diversity usually comes pre-packaged with the dominant paradigm. In the first part of the dissertation, I make a case for (1) rejecting the pre-packed assumption, and (2) adopting an alternative conception of the relation between extralinguistic context, the grammar of modals, and modal flavor.

The case against the assumption is easy to motivate with the contrast between (1) and (2).

(1) John must be at the store.

(2) John must go to the store.

(1) can have an epistemic or deontic interpretation, as the context might allow. I show that the conditions under which (2) can have an epistemic reading are highly constrained by grammatical factors – not simply by contextual ones. These constraints call for an explanation. However, as I show, the standard ways of implementing the paradigm fails to predict the constraints on modal interpretation. Drawing on the work of Hacquard (2006; 2010), I provide a way to extend the paradigm that both predicts these constraints and provides an adequate explanation of them. In the second part of the dissertation, I show that adopting my proposed framework pays unexpected explanatory dividends, for it provides the tools to explain some independent phenomena, among them i) the temporal interpretation of root modals, ii) so-called "hindsight" interpretation of modals, and even iii) the linguistic expression of the distinction between deliberative and evaluative *ought* familiar to practical philosophers.

Acknowledgements

I suspect I am not alone in enjoying reading a dissertation's acknowledgements section. It is like reading the liner notes to an album – but with the added bonus that the sense of relief expressed in the acknowledgements is often palpable, and the tone often festive. I had hoped to produce a similar kind of expression for this document, but I write this in the middle of a global pandemic the country of my residence has badly mismanaged. I presently occupy my time concerned with, among other things, the state of my country, the health and safety of my loved ones, the future and education of my children, to speak nothing of how we might be impacted by the economic fallout. So relief and celebration is currently in short supply. But I'll try. And I can certainly still express my appreciation to those that have helped me along the way. There are quite a lot of people I am appreciative of.

First of all, I thank my committee members; Fabrizio Cariani, Michael Glanzberg, Gregory Ward, Justin Khoo, and Alexis Wellwood. When I think of their work, I'd be hard-pressed to think of a group of philosophers and linguists I'd be more happy to have guiding me and assessing my work.

When I first came to Northwestern, I was the philosophy of language admit in my cohort. I think I selected this option, without much forethought, from a menu of potential research interests on the graduate application. (Northwestern had some philosophers of language on the faculty, so it seemed like a prudential way to represent my interests.) This haphazard choice turned out to be fortuitous in more ways than one. As it happened, I discovered that I did want to spend my time thinking about language. Also, Fabrizio was assigned to me as my first grad advisor. This means that Fabrizio has provided me with guidance and mentorship since the very beginning of my grad school career. So, his advice spanned quite a few topics, from the fulfilling of course requirements, to teaching, to presenting my work at conferences, right on down to navigating journal submissions and finishing up this dissertation. I can't imagine not having had Fabrizio as a sounding-board for my academic endeavors. From a standpoint of my own luckiness, the impressive and wide-ranging body of research he has produced in the meantime doesn't hurt either. Fabrizio has been a great academic role model.

Michael and I started at Northwestern in the same year, and he taught my first-year proseminar. As I progressed through the program it became apparent to me that I was interest in working on, well, "Michael-stuff". So it was natural that I would first ask

him and Fabrizio to co-chair my committee. My work on this dissertation progressed in fits and starts, and Michael was never fazed by this, providing help and guidance when I needed it. Moreover, he was always willing to help think through any argument or claim I brought to him, never discouraged me from any harebrained idea, and was quick to offer suggestions for how one might better make a line of argument work.

At one point I had mentioned to Michael that I was taking a pragmatics class with Gregory, and his succinct response was "Well, you'll be learning from the best." This was true. Gregory's classes were among my favorites in graduate school. I was excited to have him on board my committee, and even more excited that he seemed to take a genuine interest in the project. Unsurprisingly, perhaps, this interest resulted in his supplying me with a steady stream of interesting and difficult examples to consider and analyze. He may have signed off on the dissertation, but I'm still thinking of many of these "Gregory-examples".

Justin was kind enough to host me for part of a term as a visiting student at MIT, and as a result I was lucky to have gotten to know him, and have him on my committee. Aside from this warm welcome, I appreciate the astute suggestions Justin would make whenever I sent him my work. Though I'm aware that some of the claims in dissertation remain less-than-fully-baked, the most fully-baked sections are the ones that most bear Justin's mark. They are this way due to his helpful interventions.

Alexis started at Northwestern after I did, and I attended her job talk. She was my favored candidate and I was enthusiastic about her being hired. I knew quite quickly thereafter that I wanted to recruit her for my committee (well, before *that*, even). She provided important insight and comments at several crucial points in the dissertating stage and helped me navigate the terrain between linguistics and philosophy. It was a privilege for my time at Northwestern to overlap with hers; an even greater one to get to work with her on this.

I thank my cohort at Northwestern: Amy Flowerree, Morganna Lambeth, Nick Leonard, Raff Donelson, and Lauren Leydon-Hardy. I don't know how exceptional it is to like and respect the members of your graduate cohort quite so much, but if it's not, then I'm glad I lucked out with the lot of you. Also, I thank the rest of the current and former philosophy grads at Northwestern, but I'll mention specifically the people I interacted with most: David Johnson, Seth Mayer, Rebecca Mason, Ezra Cook, Gretchen Ellefson, Trevor Nyman, Mark Thomson, Jared Peterson, Nathan Weston, Carlos Pereira Di Salvo, and Cristina Carrillo Cañas. Then there are the folks in linguistics department, which came to feel like a second home: Ethan Myers, Haley Farkas, David Potter, and Masaya Yoshida (not a grad student, obviously, but getting a shout-out anyway). Peter Baumann deserves special acknowledgement and thanks for many conversations, ad hoc reading groups, and for input on quite a few papers, abstracts, and presentations.

This project got off to a start when I read Dilip Ninan's great "Two Puzzles About Deontic Necessity" for a seminar with Fabrizio. The seminar paper morphed into a conference paper and then this larger project, which no longer overtly address exactly the

same questions as its predecessor. Dilip provided comments for my 2014 APA talk, and met to discuss the evolving project several times, and was gracious and supportive in all our interactions.

Some of Bridget Copley's ideas loom fairly large in the dissertation. I was lucky to meet and discuss some of them with her at the conference on States in Natural Language at Northwestern in 2017. Her encouragement was motivating, and much appreciated.

As I mentioned above, I spent part of fall 2015 at MIT, when this project was still in its nascent stages. In addition to Justin, I had very helpful and formative conversations with Sabine Iatridou, Irene Heim, David Boylan, Matt Mandelkern, and Benjamin Storme. Sylvain Bromberger attended one of my talks and we met several times afterwards. I'm not ashamed to admit that I was a bit starstruck in meeting Sylvain, but his friendliness quickly put me at ease. Our conversations, on topics which ranged from modals, Quine, and the minimalist program, to, perhaps most surprisingly, the work of Levinas, were among the highlights of my time at MIT.

There are still others in the profession, who have discussed some of my dissertation work with me, and whose interventions I'd like to acknowledge. They include Camil Golub, Adrian Briçiu, Cem Şişkolar, Tadeusz Ciecierski, Paweł Grabarczyk, Rodanthi Christofaki, David Rey, Giulia Felappi, Natalia Karczewska, Andrei Moldovan, Paolo Santorio, Ryan Doran, Melissa Fusco, Ashley Atkins, Daniel Altshuler, and Giorgi Sbardolini. My apologies to those I've momentarily forgotten about, but who I will remember as soon as I submit this.

The work of several linguists was particularly influential to my thinking, and its influence will be apparent to anyone reading this who is familiar with these topics. Aside from Bridget, they include Angelika Kratzer, Valentine Hacquard, and Gillian Ramchand. I don't know any of them, but they deserve acknowledgment for teaching me through their work, so to speak.

One of the less congenial aspects of the last several years was balancing the time and financial constraints associated with graduate study with those imposed by parenthood. Both of my children were born during my (and my partner's) graduate studies. Once our second child arrived, it became clear that two graduate students' stipends was not going to be sufficient to allow us to cover necessities. Even more so once that graduate funding ran out. So, I had to seek work. This dissertation most certainly took me substantially longer to complete because of it, but I'm grateful for the people who saw fit to hire me to teach philosophy. Without a doubt, this allowed me *to* actually complete it! Thank you to Katrina Sifferd at Elmhurst College (and erstwhile colleagues at Elmhurst, William Hirstein and Tyler Fagan), James Murphy at Loyola University, Raja Halwani at the School of the Art Institute of Chicago, J.D. Trout at IIT, and David Hilbert at UIC. Additionally, I thank Shuji Otsuka with the N'Cat program at Northwestern, and my colleagues at the International Summer Institute at Northwestern–Julia Moore, Lisa, Brian Hampson, and Lisa Del Torto–for what turned out to be the most enjoyable time(s) I had earning a living over the summer semester.

I thank my family – first, my parents, Michael and Ingrid Skibra, who have always been supportive of my education and didn't bat an eye when I said I was going back to school for a PhD in philosophy. I'd probably have done all of this even without their approval, but actually having their vote of confidence made it a much more pleasant experience. But I still might not have been able to do it without their periodic help caring for their grandchildren when we most needed it. This has been a privilege and a blessing – both for the parents and their children. In fact, I owe my mother-in-law Marcia Will thanks for this as well. Come to think of it, *all* of my in-laws – Robert Will, and Stanley and Ana Haar – have made my life much better than I deserve to expect for the better part of the last two decades.

Finally, I've been thinking about this for a while, and still am coming up short in deciding how I can best express my gratitude to Deborah Haar. So I will keep it simple. Deborah; thank you and I love you. The rest I will tell you later. To Esther and Alexander – you've given me an education worth at least two PhDs. I often find myself in awe of the both of you. Thank you for keeping me on my toes. I don't think that you'll ever read this dissertation. But if you come across these acknowledgments, I hope they confirm what you already know – Papa ist sehr stolz auf Euch!

That about wraps up my liner notes. I think they turned out all right after all.

Table of Contents

CHAPTER 1

The Role of Context in Modal Semantics

1.1. Modal talk

This dissertation is about some issues in the semantics of modals – in particular, the semantics of modal auxiliaries; words like *must*, *may*, *might*, *should*, and so on.[1] In this introductory chapter, I will introduce the family of issues this dissertation will be concerned with, but I will do so in a slightly roundabout way – by giving a bit of a whig history of the semantics of modals. This whig history will be partial and tendentious; for example, possible worlds will loom large, and I will ignore Carnap [1947], Marcus [1993], and Quine [1969, 1953], whose interventions were consequential.[2] Really, I will just be describing a quick trajectory from modal logic to the kind of theorizing that categorizes much contemporary work in formal semantics on the semantics of modals. And I will be motivating this trajectory by a kind of flat-footed questioning about how we might characterize modal discourse. Charting the trajectory in this way will allow me to draw attention to the questions I will pose and attempt to answer through this work – in this chapter, what I'd particularly like to draw attention to is the role of *context* as it emerged in modal semantics in the development of a semantics based on modal logic.

[1]Semi-modals like *ought* and *have to* will also figure in the discussion as well, since they are thought to share the essential meaning properties with the modals, even though they are slight outliers in terms of their morphology.

[2]Some examples of a more thorough history would include, say, Copeland 2002 or Neale 2000.

Let's start with a basic question – what do modals mean? The tendency is to reach for a preliminary, and very rough, characterization like the following. The inventory of modal words just mentioned – *must*, *may*, *might*, *should*, *ought* – allow us to talk about potentially non-actual events and situations. It's hard to talk about precisely what modals mean or what they allow us to communicate without adopting modal language in our very explanation. In the gloss above, "non-actual" is a modal notion, as is (perhaps) "potential". One influential way of characterizing modal discourse is as one particular way of realizing the design feature of human language Hockett and Hockett [1960] calls *displacement*.[3] Displacement is the ability of language users to use various forms of discourse to talk about things beyond the here and now. That is to say, we can talk about people and things that are not present to hand, or about events that are not in our immediate present. We can also talk about events and states that are not at all actual. We wouldn't want to say that modal auxiliaries allow us to talk about *non-actual* events or states of affairs *simpliciter*, because we often use them to negotiate the live possibilities for the way things actually are. For example, if I remark to you an utterance of (1), my use of the modal *might* wouldn't indicate the rain is not actual.

(1) It might be raining.

The point of my utterance of (1) would be to indicate that the rain is in some sense a live possibility – at least to me, the utterer.[4] This live possibility could in fact turn out to hold. A second example. My wife asks me where our son is and I hear some commotion

[3]Cf. also the discussion in von Fintel and Heim 2002.

[4]In most cases, one wouldn't use this expression if it weren't a live possibility for oneself, though there are apparent exceptions. Cf. Egan et al. [2005]'s "exocentric" uses of epistemic modals or the Mastermind example in von Fintel and Gillies 2007.

coming from his room. I suspect that there's a good chance he's there, I respond with (2).

(2) Alexander might be in his room.

Suppose what I heard was in fact Alexander (he was banging some toys together), and that he is in fact in his room. Same story here – in spite of my using modal language to make this utterance, it is not in fact non-actual at all.

A different kind of example; if my daughter asks me for permission to have a second helping of ice cream, I can respond to her request with (3).

(3) You may have a second scoop of ice cream.

In making this utterance, I am giving her permission to have a second scoop. But her having this second scoop is not an *actual* event or state (not yet, at least). Insofar as I've given her permission, she will no doubt take the opportunity to have the second scoop. She could very well decide against it, in which case her having the second scope remains non-actual. If she's feeling particularly rebellious, she may have had a second scoop of ice cream even if I forbade it and told her to have fruit instead, by giving the utterance in (4).

(4) You should have fruit (instead of ice cream).

If she goes ahead with her insubordination and has the second scoop of ice cream, the event that I prescribe in (4) – that she have fruit instead of more ice cream – remains non-actual in spite of my prescription.

So it's not simply the case that when we use modals, we talk about non-actual states and events. Modals allow us to negotiate the potential actuality of these states and events

for our communicative purposes. But notice, again, that even in giving this gloss, we revert to modal notions in explaining what we are up to in using this language. Though these notions have strong intuitive purchase, if we want to describe what we are up to when we use this language, or what this language means, *without* adverting to modal notions in our explanation, perhaps the best strategy is to start invoking notions from the logic that has been used to model this kind of thought and talk.

1.2. Modal talk illustrated with modal logic

Much work in the semantics of modals was spurred on by developments in intensional logic in the 1960's – much of it influenced by the ground-breaking of Kripke [1959, 1963]. In this tradition, we treat modals as sentential operators. So, given a standard propositional logic, we can extend this language by introducing modal operators $\Box$ and $\Diamond$ with something like the following syntactic rules.

(5) a. If Φ is a wff, then so is $\ulcorner\Box\Phi\urcorner$

b. If Φ is a wff, then so is $\ulcorner\Diamond\Phi\urcorner$

$\ulcorner\Box\Phi\urcorner$ and $\ulcorner\Diamond\Phi\urcorner$ are read as "It is necessary that Φ" and "It is possible that Φ" respectively, and are interdefinable in the expected way.

(6) $\Box\Phi \iff \neg\Diamond\neg\Phi$

The truth of a modal formula is not simply determined in reference to an interpretation, as in propositional logic, but in relation to a model. A model for a modal sentence is a triple $\langle$W, R, $v\rangle$ where W is a non-empty set of possible worlds, R is an accessibility

relation, defined as a relation on W $\times$ W, and v is an interpretation function. We can then recursively define the truth of a formula in a model as follows:[5]

(7) a. If $v(w,\Phi) = 1$ then $w \models \Phi$

b. $w \models \neg\Phi$ iff $w \not\models \Phi$

c. $w \models \Box\Phi$ iff, for every $u \in$ W such that $w\mathrm{R}u$, $u \models \Phi$

d. $w \models \Diamond\Phi$ iff, for some $u \in$ W such that $w\mathrm{R}u$, $u \models \Phi$

This allows us to connect the intuitive glosses about our modal talk with the formal notions from logic. The set W is the set of possible worlds, which tracks the the space of possibilities. Each world is a complete, alternate history of the world. One of these is the actual world – it represents the whole history of the way everything plays out. Regardless of one's position on the status of possible worlds – whether they are concrete particulars as in Lewis 1986, or whatever – they are abstract entities that we use to model out modal notions as in Stalnaker 1976, 1984 – we can use this set of "alternate histories" to model our modal notions. To a first approximation, $\Diamond\Phi$, "it is possible that Φ", is to be understood as Φ's being true at some possible worlds. The worlds wherein Φ obtains may differ on a number of other facts – like, who won the NBA championships in 2019, whether I am currently sitting or standing right now, or whatever. But these worlds – this privileged subset of W – have at least one world wherein it true that Φ. Otherwise put, Φ is *true at* at one of them. Similarly for $\Box\Phi$; this is true just in case all the worlds (in some privileged subset) are worlds where Φ is true.

[5]I omit the familiar Boolean connectives (besides negation). I take it the reader will be familiar with these, in any event.

Consider the examples cited earlier: (1) – (4). A point that came out in discussing these examples was that the relation of the possibilities (read: *worlds*) we quantify over to the actual world differed, based on the kind of statement we were making. (1) and (2) were *epistemic* modal sentences – they negotiate possibilities based on our states of knowledge or evidence about the world. (3) and (4) were *deontic* modal sentences. The possibilities deontic modal sentences negotiate have to do with our putative obligations and permissions based on salient laws, morals, ideals or what have you. In English, as in many other languages, many of the same words can be used to express either epistemic or deontic modality.[6] The difference between *must* in the epistemic sense and deontic sense is typically glossed as a difference in modal "flavor".

What becomes clear rather quickly is that the modal flavor has a bearing on whether or not the actual world is among the possibilities quantified over by the modal. For example, in one of the deontic cases, (4), if in attempting to forbid Esther from having more ice cream and prescribing fruit for her instead, she could still disregard my demands and have herself some ice cream. In this case, (4) could be true – if it indeed is the case that the salient rules make eating fruit the best option – even if Esther fails to eat fruit and eats ice cream instead. This is not always the case. For epistemic modals, it seems that the actual world *is* among the worlds quantified over. Suppose I say, in a variant of one of our epistemic examples (2), that Alexander *must* be in his room. This is true just in case all the relevant worlds – the actual world among them – are worlds where Alexander is in his room.

[6]This point is gestured at in Kratzer 1977 and elaborated in Hacquard 2010, 2011. However, cf. Viebahn and Vetter 2016 and the citations therein for a dispute of this claim.

I've been glossing over a fact that will now be increasingly obvious; modals don't quantify over the whole set of worlds, but on privileged subsets of the set of worlds. So, modals are *restricted* quantificaters over sets of worlds. And, as the observation about the relation of the actual world to a modal's domain of quantification suggests, broad features of this domain seem to be pretty intimately related to the flavor of the modal. In the definitions given above, it's the accessibility relation R that restricts the modal domain – modals quantify over the set of worlds accessible to its evaluation world. In fact, it turns out, characterizing R in the right way determines the flavor of the modal.[7] For example, returning to example (4), when I tell Esther that she should have fruit, if I am indicating that her doing so (instead of having ice cream) is required by some the moral rules (i.e., if (4) gets a deontic reading), then R is characterized as follows:

(8) **Deontic Accessibility Relation**:

R is a deontic accessibility relation such that $w\mathrm{R}w'$ just in case all of the rules in w are satisfied in w'.

If, on the other hand, I were to use (4) to say something about the best way for Esther's to satisfy her *desires* (say that she desired to eat more fruit), then R would be characterized as a bouletic accessibility relation:

(9) **Bouletic Accessibility Relation**:

R is a bouletic accessibility relation such that $w\mathrm{R}w'$ for some individual i just in case all of i's desires in w are satisfied in w'.

[7]The point I've been pressing above, about whether or not the actual world is in the set of accessible worlds, shows that accessibility relations can also be classed based on the constraints they place on which worlds are accessible. For example, if some set of accessibility relations R_a is such that the world of evaluation is accessible to itself, then R_a is said to have the property of *reflexivity*.

Turning to (2), where the intended meaning is epistemic, R would be characterized as an epistemic accessibility relation.

(10) **Epistemic Accessibility Relation**:

R is an epistemic accessibility relation such that $w\mathrm{R}w'$ for some individual i and some time t just in case everything i knows at t in w is true in w'.[8]

The point I'm belaboring is not a terribly deep one, and will in fact seem fairly remedial to someone with more than a passing knowledge of modal semantics.[9] However, it's worth taking this circuitous path to make the following point vivid – possible worlds as utilized in the model theory for intensional and modal logics provide a serviceable way to model our modal notions.[10] The real heavy lifting done as far as modeling our modal notions is all due to the choice of R – the subset of possible worlds that are quantified over. If we are to apply this framework to natural language modals, we need some mechanism for an appropriate R to be chosen for a given expression on a particular occassion of use.

[8]This gloss seems to suggest that determining the relevant agent i whose information state is at stake is a relatively straight forward affair. But in fact, there's been considerable debate in the literature on epistemic modals *who's* information state is at stake, *when* the person (or group)'s information state is to be assessed, and whether characterizing the information state as a state of knowledge is appropriate (as opposed to, say, a belief state.

[9]One could also reasonably complain that I've also trained a disproportionate attention on but one property of R in discussion reflexivity, when there are other such properties, the joint obtaining of which results in a particular kind of modal logic. This is fair, but the point was to illustrate in a fairly theory-neutral way how intuitions about modal expressions are related to these more formal notions. The intuitions that underlie reflexivity are perhaps the easiest to evoke – certainly easier to evoke than transitivity. Cf. Kaufmann et al. 2006 for a a clear discussion on how various possible properties of R are related to the modal semantics of natural language, particularly in the kind of implementation I will call "the Standard View" a few sections hence.

[10]I originally introduced the idea by pointing out that it seems hard to explain what our modal talk means without in turn appealing to more modal notions. Ultimately, whether or not the explanation of modality in terms of possible world succeeds in avoiding the circularity I was trying to avoid is a controversial question in metaphysics. On a fairly superficial level, it seems to, and I'll rest content with this superficial impression for the present purposes.

This is particularly important for natural language modals for a number of reasons. One especially important reason is this: modal expressions typically admit of more than one flavor. This is not just true of English, but of modal expressions in most languages.[11] As we saw with the modals canvassed above, the same modal expression will admit of different flavors. Since, on this strategy for representing modal meanings, difference in modal flavor is due to a difference in R, we'd need some appropriate way to select the correct R to disambiguate the potential meanings a modal could have on a given occasion.

Having made this observation, I'll now turn to documenting and explaining the two most prominent ways natural language modals are modeled; what I will call the Context-Index View and the Standard View, respectively. The Standard View, especially, will loom especially large in the rest of this dissertation since it has a good claim to being foundational in modal semantics – if a contemporary theorist doesn't adopt the framework posed by the Standard View, they position themselves against it fairly explicitly. Prior to this, though, I will sharpen the rough and ready observations made above into some more robust generalizations about modals, which will serve as desiderata for what a theory of modals should both be able to represent and explain. We'll see how the two views meet these desiderata, as well as the kind of mechanism they posit for determining the value(s) of R. What should emerge from the discussion is an understanding of the role of context in modal semantics – in particular, how modal semantics as developed in these prominent frameworks appeals to context as the dominant (and often sole) mechanism in determining the value for R.

[11] Cf. Hacquard 2011 for discussion.

1.3. The Meaning Components of Modals

Before laying out two of the most common frameworks for the semantics of modals in natural language (including what I'll be calling the Standard Account of modals), I'll describe two features that an account of the modal auxiliaries should be able to account for. This will help show how we can apply the features of modal logic described above to the analysis of natural language modals. Also, the way such features are accounted for will prove important to the arguments of the next few chapters. First off, rather than expressing possibility or necessity absolutely, natural language modals express these with respect to a circumscribed domain. The routine way of describing this dimension of meaning is to say that modals admit of different flavors.

(11) a. Smith must be the murderer.

b. The students must finish their assignments by Friday.

c. You absolutely must try the fabulous dessert.

Many natural language modals are polyfunctional; they express multiple flavors of modality. In (11), *must* has an epistemic, a deontic, and a bouletic flavor, respectively. Epistemic modals describe the ways things could be, given some body of knowledge or evidence. Deontic modals express possibility or necessity given a certain body of laws or rules; bouletic modals express them with respect to a set of desires.[12] The categories designating modal flavor are themselves subject to finer-grained distinctions in meaning. Even within flavors of modality, there are subtle differences in meaning a given modal can have. Consider deontic *must*.

[12]Other flavors of modality include abilitive and teleological modality, expressing possibility/ necessity relative to a set of physical abilities and set of particular goals, respectively.

(12) a. The convict must serve time in prison for the crime.

b. Mary must do the dishes after dinner.

c. John must hand in his term paper by the due date.

If deontic *must* expresses an obligation, it is clear that the type of laws or ideals from whence the obligation issues is different in each of (12).[13] A little reflection suggests that all the flavors of modals admit of this kind of stratification of meaning within the flavor dimension. To fix some vocabulary, call the difference in meaning exhibited by (11) **inter-flavor**, and the kind of differences in meaning exhibited by (12) **intra-flavor**, since they concern differences in meaning within a particular flavor.

While modal flavor is but one feature of modals, any theory of modal semantics must account for differences among inter-flavor and intra-flavor distinctions in some way. It should be clear, to a first approximation, how inter and intra flavor diversity would be modeled according to the resources we have from modal logic. Both are accounted for via selection of an appropriate accessibility relation R, as discussed earlier. What we need to make this application tractable for natural language, is a mechanism by which an appropriate accessibility relation is picked out. On this matter, the prevailing view is contextualist; the interpretation of a modal is secured by context's supplying values to parameters introduced by the expression into the semantic representation of the sentence. On this view, inter- and intra-flavor distinctions share a source: the value context assigns

[13]To fill out the example a little bit, a natural context on which (12a) is intuitively true is one where, according to the laws of a country or state, the convict must serve time in prison, (12b) is one according to which the rules of the household are such that Mary must do the dishes after dinner, and (12c) is one according to which the norms of higher education are such that John must hand in his term paper by the due date.

to the contextual parameters relative to which the expression is interpreted. The parameters posited determine modal meaning on both inter- and intra-flavor axes, in virtue of the values context assigns to them.[14]

Secondly, and closely related to the topic of modal flavor, there's a further distinction amongst modals that will be important in what follows. It is common in the generative linguistics tradition to distinguish between 'root' and 'non-root', or epistemic, modals, since the latter exhibit different syntactic behavior than their root counterparts.[15] This is more of a grammatical distinction than is the classification of modals into flavors based on their interpretations. However, it is generally accepted that root modals correspond to the non-epistemic flavors.[16] I make use of this distinction in what follows, though I sometimes use these terms to refer to the grammatical type of modal word, and sometimes to the cluster of flavors aligned with the grammatical type, letting context disambiguate.

1.4. A Context-Index Semantics For Modals

A promising way to account for the flavor diversity of modal discourse is to treat them like a special case of indexical expressions. Indexicals are expressions in natural language that make essential reference to the context in which they are used. Examples include pronominal words like *I*, *now*, *here* – expressions that, if you were to understand what they were being used to refer to in some utterance, you'd have to know some details about the circumstances of that utterance – where it took place, who uttered it, and when, for

[14]Swanson 2008 puts the point as follows. "The striking ease with which a single modal can target different modalities is sometimes taken to suggest that context alone determines which modality is targeted – that there is no lexical difference between epistemic and deontic *must*, for example, but only a difference in some parameter or parameters supplied by context."
[15]Cf. especially Jackendoff 1972 and predecessors.
[16]Cf. Portner 2009 for a discussion of how the root/ non-root distinction applies to the flavor categories.

example. The context of utterance supplies this information to the interpretation. It would seem reasonable that modals had an element of this kind of indexicality as well – not simply insofar as the accessibility relation required for the interpretation of the utterance is to be determined contextually, but also because there are details within the accessibility relation which seem directly sensitive to the context as well. Let's consider (10), repeated below as (13).

(13) **Epistemic Accessibility Relation**:

R is an epistemic accessibility relation such that $w\mathrm{R}w'$ for some individual i and some time t just in case everything i knows at t in w is true in w'.

Recall that we not only have interflavor diversity, so we need to characterize different kinds of accessibility relations. But there's also intraflavor diversity, meaning that, holding fast the epistemic flavor, the accessible worlds will be different based (in this case) on which individual's information state is being considered, and at what time. Some have argued that the person is always the speaker and that the time is always the time of the utterance. If this is so, then there's a kind of indexicality built in to modal's meaning. (Similar considerations apply to the other flavors.) This is in addition to the other essential feature of modals – their (as it's been called) "shiftiness". Modals shift the world of evaluation to another world, leaving the other parameters or coordinates fixed, to the evaluate the sentence in its scope at that (or those) world(s). These two features motivate what are often called context-index theories, and owe largely to the work of Kaplan [1977] and Lewis [1980].[17]

[17]In his 2009, Portner describes what he calls "Modal Logic for Language", but this is essentially a form of context-index theory.

Kaplan [1977] famously distinguished between the *character* of an expression as its context-invariant conventional meaning, and the expression's *content*, or the component of meaning that the expression contributes to a larger expression containing it via a compositional semantic process. Kaplan's work rethought the relation between context and the meaning of indexicals. Since Kaplan's treatment of indexicals and related context-sensitive expressions provides a template for thinking about how context helps determine the content of these types of expressions, it makes sense to think about how an account would capture the facts about modals canvassed in earlier sections. Since I intend to point out problems for a contextualist semantics of modals more generally, prior to moving on to Kratzer's semantics, I will describe a toy semantics for modals in the style of Kaplan and Lewis.

Kaplan's logic of demonstratives makes use of two types of coordinates; the context coordinate, and the index coordinate.[18] Correspondingly, truth is defined relative to these two coordinates, yielding a notion of content like:

Content of ϕ at c: $\lambda i.[\![\phi]\!]^{c,i}$

The context coordinate is required for the context sensitive expressions in the sentence to have a semantic value. Its value is supplied by the extra-linguistic features of the utterance situation. (Of course, since we've defined truth relative to a context and an index, the context coordinate is always required, but for non-context sensitive expressions, its role

[18]This was an innovation on Kaplan's part. Earlier theories in the formal semantics tradition (Cf. Lewis1970a, Montague 1970, and the "advice" of Scott 1970) capture all manner of context sensitivity through the index, without reference to a separate content parameters here. Though I won't rehearse them here, Kaplan gives important arguments against modeling context sensitivity wholesale in this manner. (One important consideration is that pure indexical expressions like "I" don't shift their value under the scope of "shifty" operators.) Other salient arguments are given by Cresswell [1973]. (Cf. Glanzberg [2007] for additional discussion.)

is vacuous.) Consider a sentence like *I am tired.* The truth of this sentence depends not only on the way things stand in the world; who is tired and who is not, but who is speaking at the context at which this sentence was uttered. The context coordinate contributes the requisite information to semantic interpretation that allows the context-sensitive expressions contained in the sentence (i.e., the expression 'I') to have a content to contribute to the sentence that contains it– in this case, information about the speaker of the utterance. This illustrates the role of character in Kaplan's theory: it is a function from contexts to content (put otherwise, the character of an expression is a function that takes the context and yields as a value the extension of the expression relative to that context); $[\![\text{I}]\!]^c = \text{f}_I(c)$ = the speaker in context *c*. In slightly more formal terms, the context, c, is a sequence of coordinates, usually thought to be features of the discourse situation, like the speaker or agent, the addressee, the location, time, etc., of the utterance event. So *I* is a function from contexts to the value of the speaker/ agent coordinate of c.[19]

The index, by contrast, treats so-called "shifty" phenomena in natural language. These are expressions, represented as operators in the semantic representation of the sentence, whose sentences containing them are evaluated by shifting a feature of the index that is initially supplied by the context.[20] That is, the index coordinates are given the value from the relevant features of the context. The idea motivating this approach is that, while the

[19]By contrast, non-indexical expressions are also functions from c, but they are constant functions, picking out the same object/ set of objects in every context.

[20]At least, this has been the standard story. Recent work in the philosophy of language has complicated this picture, especially in light of MacFarlane [2003, 2009, 2014]'s suggestion that the initial values provided to the index parameters are not part of the semantic determination of the sentence's meaning, but to the "postsemantics", allowing for varieties of relativism to be articulated within the context-index framework.

sentence *It's raining* is true in a context c if *It's raining* is true at the index supplied by c, on the assumption that tense is such a shift operator, the sentence *It was raining* is true in a context c if the sentence *It's raining* is true is evaluated by keeping constant the other features of the index coordinates set by c, but shifting the time coordinate to that time specified by the past tense morphology. The index is itself an n-tuple consisting of precisely those features of the context that can be targeted by such index-shifting operators. Lewis [1980] takes sentences like the one just given as evidence for the existence of a temporal coordinate in the index; his construal of the index includes coordinates for modal, tense, location, and standard operators.[21] For illustrative purposes, let's follow Lewis by assuming that the index will contain a world coordinate (for modals), a time coordinate (for tense), and a location coordinate (for locative operators; I suppress Lewis's coordinate for standards). Then, the index is an ordered triple $< w, t, l >$, with world, time, and location coordinates. Index-shifting operators take the semantic value of the sentence containing the operator (which has its value relative to a context c and an index i), and evaluate the embedded sentence according to a new index, i', which is just like i except for the value of the coordinate in i that is shifted by the index-shifting operator.

Insofar as the index contains a world coordinate, this framework naturally extends to modal sentences. Take, as an example, (14).

[21]There is actually some controversy about whether language contains such expressions, and which expressions they are. For example, though Lewis [1980] assumes we should treat tense as such an operator, most contemporary theories of tense treat tense as contributing pronominal variables to logical form. Cf. Partee 1973, Enç 1986, 1987 for classic arguments for treating tense in this way. Stone [1999] makes parallel arguments for the modal domain. Cf. King [2003, 2007] for arguments that the index does not contain a time coordinate, and how this purported fact favors a theory of temporally specific propositions. Ninan [2010] argues, convincingly in my opinion, that the question of whether a particular kind of coordinate belongs in the index is not settled by the question of whether that coordinate is shifted by an index-shifting operator. These are thorny issues, and I don't intend my discussion to advocate an operator approach to any of these phenomena per se; just as a discussion of how such phenomena could be modeled, and how context contributes to this modeling.

(14) Esther must go to school.

Rendering *must* as a sentential necessity operator, we can analyze (14) as:

$[\![\text{MUST [Esther go to school]}]\!]^{c,i} = 1$ iff

(15) a. $[\![\text{MUST}]\!]^{c,i}$ ($\lambda i'.[\![\text{Esther go to school}]\!]^{c,i'}$) = 1 iff

b. All worlds w accessible from w_i are such that $\lambda i'.[\![\text{Esther go to school}]\!]^{c,i'} = 1$ iff

c. All worlds w accessible from w_i are such that $[\![\text{Esther go to school}]\!]^{c,\langle w,t_i,l_i\rangle} = 1$ iff

d. All worlds w accessible from w_i are such that Esther goes to school at t_i at l_i in w.[22]

This derivation allows us to see how a context-index theory models the apparent shiftiness of modals, and how it integrates this in a compositional semantic theory. But the derivation above abstracts from some important features that we'd need to give an accounting of for this to serve as an adequate account of modals. For example, we need some way to determine the flavor of the modal – whether the *must* in the above example has, say, a deontic or teleological flavor. One way of accounting for this diversity of modal meaning is to build it into the account by restricting the accessibility relations on the sets of worlds in the right way. In the derivation in (14), I appealed to a set of worlds accessible from w_i. If the worlds accessible to w_i are wherein all the relevant laws are obeyed, then *must* is deontic; if the worlds are those where all our goals are met, then it is teleological. We

[22]Note that since we're assuming that the index is "initialized" by the context, the values t_i, l_i, etc., are really t_c, l_c, etc. That is, they are the time, location, etc., of the context.

need some way to account for the inter-flavor variations of *must*. *Mutatis mutandis* for the intra-flavor variations.

Portner [2009] suggests a solution. With respect to the semantic interpretation of modals, we can consider the context not just to provide a world coordinate to the index, but also to set the flavor of the modal by selecting the appropriate accessibility relation. If M is a modal, then the set of all accessibility relations according to which M can be interpreted is the set A^M. We can follow the model of indexicals to yield a serviceable way for modals to be paired with the relevant accessibility relation. Recall that the character of a word like 'I' is the function f_I from contexts to individuals that are the speakers of the context, for any context c. Likewise, we can think of A as an accessibility relation function, such that it picks out the relevant accessibility relation given a context. So, A can be the function from contexts to A^M. This allows us to keep the formulation in the derivation above, but we are able to specify that the worlds w are accessible from w_i in virtue of being those worlds accessible via the contextually determined accessibility relation.

Also, as we saw, accessibility relations seem to exhibit their own kind of indexicality above the aforementioned context sensitivity. So, if the context is in the business of determining the modal flavor by indicating the accessibility relation, then we will also want context to provide the values for the parameters in the accessibility relation, much like context would "initialize" the coordinates in the index by giving them the values those coordinates have at the context of utterance.

This gives us a purchase on the flavor diversity exhibited by modals in a context-index theory. However, there are two difficulties with the formulation as I've just presented it.

First of all, the formulation just offered treats modals as having a unique accessibility relation for each context. Modals can often embed other modals, as in (16) (Portner [2009]'s (66)).

(16) It must be raining so we should take an umbrella.

Since we've aligned modal flavor with accessibility relations, this would amount to saying that there is but one flavor in (16), when a natural interpretation of (16) has it that the modals *must* and *should* have different flavors. There are two options for resolving this. One is to say that there are in fact separate contexts for the two coordinated clauses in (16). A second option is to allow a context to select an n-membered sequence of accessibility functions. Much like with the coordinates of the index, exactly how many members such a sequence would need to contain is an empirical matter, determined by how many modals we can embed in a given context. (16) suggests that n is at least greater than or equal to 2, but the framework provides no barrier for making the number arbitrarily large.[23] So, we can say there are at least n accessibility functions for a given context. For (16), $A_1(c)$ would be an epistemic accessibility relation $A_2(c)$ would be a deontic accessibility relation, and so on (that is, $A_1(c)$ picks out the first member of the contextually determined sequence of accessibility relations, $A_2(c)$ picks out the second member, etc.).

The second complication has to do with the following issue, which makes trouble especially for deontic modals. The account given above works quite well when we are discussing activities that are explicitly permitted, forbidden, or required. For we can say

[23]We could just make the sequence infinite, much like Tarski's assignment function, which we already make use of for the interpretation of quantifiers.

Esther should treat people with respect. If one of the rules is that you should treat people with respect, and then a world with a deontic accessibility relation will have it that all the accessible worlds are worlds where people are treated with respect. So the analysis delivers truth conditions which indicate that Esther treats people with respect is true in all the worlds where all the rules are met.

But consider a case where Harry commits a murder. It makes sense to say, in those conditions that *Harry should go to jail.* But our definitions have some trouble with this. For if we look to the accessible worlds, these are worlds where the rules are all satisfied – not murdering people being one of the rules satisfied. In this set of worlds, none of them are worlds where Harry goes to jail, since no murder will have occurred. What we would want, it would seem is for our interpretation of at least some modals to be able to encode the assumption that the murder actually happened; we want to be able to look at worlds that hold this fact constant, and the deontic accessibility relation by itself doesn't give it to us. So, it would seem that to take account of this feature of modal discourse, we'd need amend the semantics presented here. Rather than searching for a way to do this, I will turn to the final framework to discuss in this chapter – Kratzer's semantics for modals, what I'll be calling the Standard Account.

One observation to make about the context-index framework is that the mechanism for determining the set of worlds the modal quantifies over (and through that, for determining the modal flavor) is one that exploits the context sensitivity of modals and by making this a function of context. A second, and related observation is that inter-flavor and intra-flavor diversity are determined by essentially the same mechanism. For example, whether

a modal like *must* is interpreted epistemically or deontically depends on whether the function $A_1(c)$ picks out an epistemic or deontic R. And this R will have properties common to all epistemic or deontic relations. But obviously there are finer-grained distinctions between accessibility relations than this. For example, on the framework described above, the accessible worlds for an epistemic R will depend on a number of factors, like who's knowledge or information state is relevant, and which time. If these factors themselves are context dependent, as is suggested by saying that they are comparable to the context dependence of indexical expressions, then the mechanism just discussed will capture this just fine. But it means that *both* inter- and intra-flavor differences are accounted for by the same mechanism. Though there are reasons to make amendments to the framework laid out for the reasons pointed out, but it doesn't seem like these reasons will result in a fundamental change of this mechanism.

1.5. The Semantics of the Standard Account

What I will henceforth call the Standard Account is due a series of seminal papers by Angelika Kratzer.[24] The Standard Account construes modal sentences as having the the following underlying form: MOD(R)(ϕ). The modal operator, MOD, takes two arguments; R, the restrictor, which determines the domain the modal quantifies over, and ϕ, the nuclear scope, which is the sentence the modal scopes over (commonly known as the prejacent). The sentence *John must be the murderer* then has the form:

(17) MUST(R)(John be the murderer)

[24]Cf. Kratzer 1977, Kratzer 1981, Kratzer 1991, Kratzer 2012

According to Kratzer [1981], two "conversational backgrounds" comprise R and thereby jointly affect the restriction on the domain of possible worlds: the modal base, *f*, and the ordering source, *g*. Here is how these two features work. The modal base and the ordering source are both functions from worlds to sets of propositions. To accommodate the distinction between *epistemic* and *root* modals, Kratzer allows two modal bases, $f_{epistemic}$ and $f_{circumstantial}$ (henceforth abbreviated as f_{ep} and f_{circ}, respectively) an epistemic and a circumstantial one, with the circumstantial modal base corresponding to the class of root modals.[25] The nature of the set of propositions delivered by each modal base is commensurate with its role in the interpretation of the modal. Epistemic modals have to do with possibility or necessity given states of information, so f_{ep} will yield a set of propositions characterizing some information state – namely, the propositions that are true with respect to this information state.[26] Root modals have to do with possibilities or necessities given the obtaining of certain facts, so f_{circ} will yield a set of propositions characterizing some circumstances. Ultimately, we will want the modal to quantify over a set of possible worlds. If we treat propositions as sets of possible worlds, these functions give us a way to accomplish this. Intersecting the propositions given by the modal base gives us a single set of possible worlds, which we can call the modal domain. Then it makes sense to talk of a circumstantial domain or an epistemic domain, depending on whether the modal base is circumstantial or epistemic. The definitions of circumstantial and epistemic domains are given in (18).

[25]Recall from section (1.3) that root modals are traditionally considered to encompass all those flavors of modality that are non-epistemic.

[26]This description can characterize so-called epistemic modal bases, even if you think the information state that epistemic modals are evaluated with respect to a set of propositions *believed* true as opposed to being known to be true. I suspect that epistemic modals can be evaluated like this, but it won't be very important in what follows, and so I will put this aside.

(18) a. $\bigcap f_{ep}(w) = \{\text{w}' |\ \text{w}'$ is compatible with what is known by the relevant agent(s) in $w\}$

b. $\bigcap f_{circ}(w) = \{\text{w}' |\ \text{w}'$ is compatible with certain circumstances relevant in c in $w\}$

Note that we can define an accessibility relation in terms of Kratzer's conversational background. If v is accessible from w (so $w\text{R}v$), using Kratzer's apparatus, this means that v is true in every proposition in the set given by $f(w)$. Moreover, $v \in \bigcap f(w)$. The ordering source, g, also picks out a set of propositions, and the relation $\leq_g$ imposes a preorder on the worlds in $\bigcap f(w)$ according to which propositions in $g(w)$ hold at each world. Since this has the effect of restricting the set of worlds to those deemed "best" according to g, we can think of g as further restricting the domain as follows. BEST_g is a function that picks out the best worlds (according to g) in its domain. We are now in a position to see how the Standard Account implements the intuitive paradigm formally.[27]

(19) $[\![\text{MUST } \phi]\!]^{w,f,g} = \forall w' \in \text{BEST}_{g(w)}(\bigcap f(w))\text{: } \phi(w') = 1$[28]

Less formally: 'MUST ϕ' is true iff ϕ is true in all of the most highly ranked worlds (according to the standards of $g(w)$) compatible with $f(w)$. The contextual parameters f and g allow us to stay quite close indeed to the intuitive paradigm glossed in the introduction.

[27]Here I adopt the formulation from Portner 2009, which makes the Limit Assumption, though Kratzer 1981 does not. I adopt it here for ease of formalization, but this formulation is now standard.

[28]I've suppressed other parameters of the interpretation function besides those relevant to this presentation.

1.6. Key Features of the Standard Account

There are two key features of the Standard Account I wish to draw attention to. These features are perhaps more properly understood as ways the Standard Account makes good on some antecedent theoretical commitments, as opposed to entailments of the semantic theory. These are: the implementation of the uniformity hypothesis and a strong understanding of the role of context in determining modal flavor (both inter and intra).

1.6.1. Uniformity

On the Standard Account, both the inter- and intra-flavor differences in meaning are accounted for via the conversational backgrounds f and g. Accounting for them so allows the Standard Account to explain the differences in modal interpretations without thinking of modals as ambiguous. The uniformity hypothesis is a principled constraint, with the aim of avoiding a proliferation of modal senses for polyfunctional modals. Since (translations of) modal expressions apparently have a strikingly similar flavor profile across multiple languages, this makes it highly unlikely that polyfunctional modals like *must* are homophonous realizations of different words. The Standard Account can give content to this constraint, through something like:

> UNIFORMITY: Modal words in natural language have uniform lexical entries, in spite of their ability to be used to express different flavors of modality.

In spite of the fact that a given modal expression may be polyfunctional, what that expression introduces into the logical form of the sentence containing it is nonetheless univocal. By appealing to the lexical entry of the modal expression, UNIFORMITY makes

good on Kratzer [1977]'s contention that modals intuitively have a "common kernel" of meaning across flavors. In every sentence where *must* (for example) occurs, regardless of the flavor, it has the following lexical entry.[29]

(20) $[\![\text{must}]\!]^{w,f,g} = \lambda P \lambda f \lambda g \lambda w[\forall w' \in \text{BEST}_{g(w)}(\bigcap f(w)):\ P(w') = 1]$

The denotation of *must* combines with the denotation of a prejacent, P, and, given a modal base and an ordering source, the resulting proposition is true relative to a world according to which f and g get the appropriate values. The flavor the resulting modal sentence admits of is the result of the assignment of values to the parameters. On this picture, this is entirely a pragmatic affair, and requires no flavor-induced change in the lexical entry for the modal.

1.6.2. Pragmatic Resolution

The value-setting role of context on the Standard Account is rather powerful. Our motivating paradigm includes the insight that modals are quantifiers over contextually determined sets of worlds. As we saw in sections 1.5 and 1.6.1, the restriction is accomplished via parameters whose values are assigned via context. If this is right, then absent idiosyncratic lexical features, the values of the parameters are resolved pragmatically.[30] Short of such hard-wired restrictions, nothing in principle constrains context in its setting the value of the parameter. The model we have for this is essentially that of NP anaphora.[31] This feature is captured by PRAGMATIC RESOLUTION.

[29] Again, irrelevant parameters omitted.

[30] An example of such an idiosyncratic feature is that *might* is often said not to have a deontic reading. So, context could not conspire to yield a deontic reading for *might*. The advocate of the Standard Account allows this feature to be hard-wired into the lexical entry for *might*, limiting its possible interpretations.

[31] Cf. Stanley 2000 and Stanley and Szabó 2000.

PRAGMATIC RESOLUTION: The contextual parameters introduced by modals exhaustively determine the meaning of the modal along the flavor dimension through the pragmatic assignment of values to the parameters.

PRAGMATIC RESOLUTION encodes the intuition that modal expressions are context sensitive. It fact, it parallels the observation made about the context-index framework for the determination of modal flavor – that inter- and intra-flavor diversity is resolved through the same contextual mechanism. As we saw at the end of last chapter, accounting for the variety of meanings modals are capable of wound up demanding an increased complexity of the mechanism for determining the accessibility relation. In the Standard Account, the double-relativity of modals simplifies this somewhat, in that all manner of flavor diversity is determined by the interaction of the two parameters f and g. But since both f and g are contextual parameters, the observation made about the context-index framework applies here as well.

There's a noteworthy division of labor between the modal base f and the ordering source g – the modal base typically does not determine the modal favor on its own. (I say "typically", because it may very well be the case that an epistemic modal base is sufficient for an epistemic reading of the modal, though this claim may prove contentious.) It's the combination of modal base and ordering source which determines the flavor. For example, bouletic and deontic interpretations of a modal may share a modal base in f_{circ}, but differ in their ordering source.

CHAPTER 2

Introducing the Eventivity Constraint

This chapter will introduce what I call the "Eventivity Constraint". As the name suggests, the Eventivity Constraint constrains the kind of interpretation a modal can have. I argue in the chapter that the modal *must* exhibits this constraint, and set up the discussion to subsequently investigate how the Eventivity Constraint affects what I called the "Standard Account" of modals in the last chapter.

2.1. Diagnosing a Difference in the Prejacent

Here is a deceptively simple example that will serve as the basis for questioning some of the suppositions pointed out in the last chapter – in particular the assumptions pertaining to the relationship between context and modal flavor. I'll begin by pointing out this example, consider how it raises some puzzles, and then begin to investigate what is noteworthy about it. I hope thereby to get clearer on the nature of the puzzle, and ultimately lay some groundwork for explaining what is going on in the puzzle. Consider (21).

(21) John must go to the store.

At first blush, there is nothing particularly noteworthy about (21). It is a simple, unembedded modal sentence. It's not a conditional, it doesn't have any troublesome logical operators, and it doesn't contain any quantified NPs with scopal properties that impact

the interpretation of the sentence.[1] According to the way the Standard Account would have it, (21) would be analyzed as follows.

(22) MUST(R) John go to the store.

In (22), recall, MUST is a universal quantifier over a restricted set of possible worlds. R does the restricting. As the story goes, R is itself composed of two conversational backgrounds, which are determined contextually. Insofar as the conversational backgrounds determine the set of worlds quantified over, and these worlds determine the flavor of the modal, the flavor of the modal should then be determined contextually. Whether (21) is interpreted epistemically or deontically or teleologically or what-have-you, will be a matter of the contextual provision of values to the parameters that determine the modal flavor.

Here's the problem; (21) doesn't seem to be able to get (or at the very least, to easily get) an epistemic interpretation.[2] As a first observation, this is very surprising from the perspective of the Standard Account. (21) should be able to have an epistemic interpretation to the degree that context would allow it. But intuitions about the modal flavor of (21) – that a deontic or teleological flavor (or any one of the flavors associated with root modality which *must* admits of) is fine, but an epistemic flavor is out – are fairly strong, even prior to the consideration of *any* kinds of contexts in which (21) may be uttered.

[1]Depending on your analysis of definite descriptions, *the store* could be a quantified NP, but it doesn't cause any trouble for the interpretation of the sentence. Of course, the most natural reading of *the store* is as a weak definite, the proper analysis of which is difficult and contentious. What I have in mind, though, is that whatever one's account of this NP, or of descriptions in general or weak definites in particular, won't have much bearing on the interpretation, unlike some quantified NPs in subject position.

[2]I'm hedging this claim now, but a little bit more explanation will allow me to put claim in a bit more forceful terms.

This is very different from a sentence at least superficially similar to (21), but very different from the standpoint of modal interpretation, (23)

(23) John must be at the store.

Superficially, the difference between (21) and (23) is simply the difference between *go to* and *be at*. But from a standpoint of modal interpretation, and taking the Standard Account as given, (23) behaves as we'd expect it to. It admits of epistemic and root flavors of the modals just as easily. One can easily reverse engineer contexts that would make utterances of (23) appropriate given an epistemic or a deontic reading.

This comparison between (21) and (23) suggests a first place to look for what's special about (21) – at least, what's special about it such that it lacks an easily accessible epistemic reading. The predicate *be at the store* describes a state, whereas the *go to the store* describes an event. Perhaps the semantic properties of these event descriptions provide a clue as to why (21) apparently lacks an epistemic reading. If so, investigation into lexical aspect would be fruitful. In the literature on lexical aspect, predicates fall into one of at least four aspectual classes based the temporal structure of the eventualities the predicates denote; states, activities, accomplishments, and achievements. Of these four, the latter three are known as the 'eventive' classes.[3,4] Exactly how the eventualities differ,

[3]The classic classification of the aspectual classes (also called *Aktionsarten*) are due to Vendler 1957, Dowty 1979, and Mourelatos 1978, with important precursors in Kenny 1963 and Ryle 1949, though these distinctions trace back to Aristotle. Some authors add additional classifications, like semelfactives (Cf. Smith 1991), but others (e.g. Rothstein 2004) find this addition unnecessary.

[4]Two tendencies for variation in the literature are worth flagging. First, about what to call the class that encompasses all four aspectual classes. I follow Bach 1986 in referring to them as 'eventualities'. This avoids confusion the other oft-used term, 'situations,' (e.g., as in Smith 1991) can cause with respect to situation semantics. Second, some authors, like Mourelatos and Bach, classify activities as 'processes' and distinguish these from bona-fide events like achievements and accomplishments. I follow Vender, Dowty, and Rothstein in calling them events. While distinguishing between processes and events might make a difference in the mereology of events, for our purposes, the issue is merely terminological. Cf. Smith 1999 for a compelling case assimilating activities to the class of events.

and on what grounds they are to be classified, is a matter of debate, but Rothstein [2004] provides a taxonomy that will be useful for our purposes. Let's start by appealing to several examples. The following list is from Dowty [1979].

(24) a. **States**: know, believe, have, desire, love, understand, be happy

b. **Activities**: run, walk, swim, push a cart, drive a car

c. **Achievements**: recognize, spot/notice, find/lose, reach, die

d. **Accomplishments**: paint a picture, make a chair, deliver a sermon, draw a circle, recover from an illness, build a house

Rothstein suggests that this four-way classification of eventualities allows us to distinguish them on the basis of their status with respect to two properties, what she calls [±stages] and [±telic]. These two properties aim to capture intuitions about whether an eventuality has a natural stopping point, and whether it can be analyzed as progressing or developing (whether it is dynamic), respectively. The property [±telic] groups accomplishments and achievements together on the one hand, as being [+telic], and states and activities together on the other as being [−telic]. The intuition here is that states and activities can continue on indefinitely once they have started. Surely it's not physically possible for most, if not all, activities and probably most states to continue on indefinitely. However, there's nothing in the event description supplied by the predicate that would determine its end point. For example, an eventuality described by *know* or by *run* could continue on indefinitely, and still, just in virtue of continuing, be in the extension of *know* or *run*, respectively. By contrast, an eventuality falling under the denotation of *paint a picture* or *reach the summit* will not continue on indefinitely; there is a natural

endpoint to the eventuality – namely when the picture is completed, or when the summit is in fact reached. The eventuality will not extend beyond that moment and still fall under the denotation of *paint a picture* or *reach the summit*.[5] Krifka 1989 gives a precise characterization of this intuition. He identifies two further properties, **cumulativity** and **quantization**, possession of which leads to atelicity and telicity, respectively.

(25) A predicate X is **cumulative** iff:

$\exists x \exists y[X(x) \,\&\, X(y) \,\&\, \neg x \sqsubseteq y \,\&\, \forall x \forall y[X(x) \,\&\, X(y) \rightarrow X(x \sqcup y)]]$

(26) A predicate X is **quantized** iff:

$\forall x \forall y[X(x) \,\&\, X(y) \rightarrow [x \sqsubseteq y \rightarrow x{=}y]]$

According to Krifka, an event e is telic just in case the following holds: if e is in the denotation of X, then all subevents of e which are also in the denotation of X must have the same starting and stopping points. We see that if an event is quantized, it will be telic, on this construal. Namely, if e is in the denotation of X and e is quantized, then any other event e' which is a subpart of e yet in the denotation of X will ipso facto be identical to e. This is because, if it's quantized, it can't be the case that a *proper* subpart of e will still be in the denotation of X.[6] And, an event will be atelic if it is cumulative. As per

[5]In the linguistics literature, the "natural endpoint" locution seems to be a manner of speaking. To be sure, the distinction between telic and atelic events appears in Aristotle's *Metaphysics* Θ in his discussion of *energia* and *kinesis*, and Aristotle most certainly thought that the distinction was one between which events were teleological in nature and which were not. Perhaps there are such events. (Cf. Kroll [2015] for discussion.) But to think of things like this doesn't sit easily with Rothstein's insistence that lexical aspect merely provides descriptions of events (where the same real-world event can be described in different ways and therefore have different properties relative to that description), as opposed to tracking metaphysical distinctions in events themselves. So, by "natural endpoint", I presume that she is not advancing a metaphysical thesis about the natures of different events, but instead claiming that endpoints in some event descriptions are encoded by the lexical material in the event description, and some are unspecified by the lexical material.

[6]Rothstein points out that this cumulativity doesn't track telicity exactly because there are telic predicates that are non-quantized. For example, *run to Paris* is telic, but non-quantized, for the following reason.

the schematization above, a predicate will be cumulative just in case it has two distinct events in its denotation (so, neither of the events is a part of the other), and for any such events their mereological sum is *also* in the denotation of the predicate.[7] Essentially, the definition given in (25) holds for mass predicates generally, and as Bach [1986] points out, activities are like a kind of mass predicate in the event domain. There's a way of characterizing this property distinctively for the event domain with what Rothstein calls **S-cumulativity**.

(27) A predicate X is **S-cumulative** iff:

$\exists e \exists e[X(e) \,\&\, X(e') \,\&\, \neg e \sqsubseteq e' \,\&\, \forall e \forall e''[X(e) \,\&\, X(e') \,\&\, R(e,e') \rightarrow X^{S}(e \sqcup e')]]$

In (27), R is a relation, and ^{S} is the operation forming a singular entity out of a sum. The need for R and ^{S} is because two events can be summed to form a plural event, but given the right kind of contextually determined relation R (Rothstein suggests temporal adjacency as a common such relation), they can also be summed to form a new *singular* event. It is the summing to a new singular event that underlies the intuitions under discussion here. An event which is S-cumulative is also thereby atelic. As Rothstein explains, if e, e′ and $^{S}(e \sqcup e')$ are all in the denotation of X, and e is not a part of e′ (or vice versa), then one of e, e′ ends at an earlier point in time than $^{S}(e \sqcup e')$. Yet all of e, e′ and $^{S}(e \sqcup e')$ are in the denotation of the same predicate X, so unlike with quantized events, the stopping point is not lexically specified by the predicate. Let's say you run from Brussels to Paris. Now, consider the event of your run from Brussels to Paris, but not as an event falling under the predicate *run to Paris*, but instead just as an event

If you start in Brussels and run to Paris via Amsterdam, then your Brussels-Paris run will fall under the denotation of *run to Paris*, but so will the subpart of your run spanning Amsterdam-Paris.

[7]The sum operation is from Link 2002.

falling under the predicate *run*. Obviously, as just stated, your Brussels-Paris run falls in the denotation of *run*. But so do both the Brussels-Amsterdam and Amsterdam-Paris legs of your run, and neither of the latter two are proper parts of the other. So, unlike with the event description provided by the predicate *run to Paris*, the end-point is not specified by the predicate *run*. In summary, telics form a class of predicates that describe events, and there are a number of lexical properties associated with telicity.

The second property, [±stages], groups together activities and accomplishments on the one hand, and states and achievements on the other, according to the intuition that predicates that fall in these categories are dynamic, and their dynamicity consists in the fact that they progress or develop. A common test, then, for whether an event is [+stages] is the ability to appear in the progressive. For paradigmatic cases, activities and accomplishments can appear in the progressive, but states and achievements cannot.[8]

(28) a. *Alexander is believing in the afterlife/loving Esther. (state)

b. *Esther is recognizing Alexander/losing her pen. (achievement)

c. Esther is running/walking. (activity)

d. Alexander is reading a book. (accomplishment)

What exactly the progressive test shows ultimately depends on one's account of the progressive. But the intuition underlying this test is that the progressive allows one to assert, of an eventuality, that is on-going. When in the present tense, this amounts to saying that

[8]As with any such test, it admits of exceptions. There are cases where states and achievements can appear in the progressive. For states, there are what Dowty [1979] calls "interval states", like *This summer, we are living in Hannover*. For achievements, we have examples like *The guests are arriving now*. But even these exceptions are instructive because there is some evidence that what explains this behavior is a kind of coercion whereby a state gets something like an activity reading, and an achievement gets something like an accomplishment reading, in order to appear in the progressive. Cf. Rothstein 2004, Ch. 2 for an analysis of progressive achievements.

a stage of the eventuality is occurring at the present moment. So, states and achievements are bad in the progressive because they do not themselves have stages. Accomplishments and activities *do* (and, the extent to which states and achievements are OK in the progressive corresponds to the extent to which they can be coerced into a reading where they have stages themselves). To cash this out more formally, Rothstein [2004] appeals to a formulation from Landman [1992]'s account of the progressive. Landman argues that the meaning of a sentence in the progressive is that a stage of the eventuality described by a verb or predicate occurred, where e is a stage of e′ if e develops into e′. The way Rothstein sees it, whether or not an eventuality has stages is a proxy for the question as to whether or not it is a process or has a process as a part of it, and underlies the intuition that certain events "go on" or "progress". Take an example from the list of predicates above, *walk*. The eventuality described by *Alexander walked* goes on, as we'd say. Take the totality of the event described by this sentence; it certainly has a stage insofar as it has a proper part which develops into the entirety of the event described by *Alexander walked*. Or, take the accomplishment predicate *paint a picture*. The eventuality described by *Esther painted a picture* has a stage, insofar as there is an eventuality, say Esther's pressing the crayon to the paper making marks on it, which develops into the event described by *Esther paints a picture*.

In sum, we see the following distribution of properties for each category in the Vendler/Dowty categorization, as construed by Rothstein.

(29)

	±stages	**±telic**
States	−	−
Activities	+	−
Achievements	−	+
Accomplishments	+	+

Interestingly, when we cash out the differences between states and events in this way, there's no aspectual property that unites the eventive categories against the category of states. Typically, we think of events as happening, but if we think of "happening" as the development of an event through a series of stages, as Rothstein suggests, and construe this in terms of [+stages], then achievements don't have this property, since the change of state that characterizes them is instantaneous. Alternately, if we think of events as effecting some kind of change in the world, and construe this as [+telic], then activities do not have this.[9] This makes it at least prima facie puzzling what semantic property telic predicates share with atelic activities that could be leveraged to explain why you don't find them under epistemic *must*. At the very least, activities do not specify the relevant change that must take place in order for the event to be completed, the way achievements and accomplishments do. Because of this, it's difficult to think of a candidate property that all events have such that this property might explain why it seems to *force* an epistemic interpretation on the modal in examples like (21). So, the explanation for this is not obvious. Nontheless, let's press forward for now, to make sure that we're in fact dealing with a robust enough phenomenon that it even calls for an explanation. In particular, now that we have a better understanding of the various aspectual classes, we'll want to

[9]To be sure, an event falling under the predicate *run* involves a change of position of the person doing the running. But not all activities involve a change in position like this.

see the extent to which the interpretational profile I'm pointing out in (21) generalizes to examples from the other aspectual classes.

2.2. A putative generalization

So at first blush, it seems like when *must* has a prejacent with an eventive predicate, this seems to force a reading from the family of root flavors on the modal. So far, we're only noticing this behavior with *must*, but whether *must* is idiosyncratic in this regard, or whether other modals admit of the same kind of constraint is a question which I'll revisit later. Suffice it to say that whatever effect or phenomenon we're tracking in noticing this constraint is very prominent in the case of *must*. Let's hazard a tentative generalization like the following to describe what we're seeing.

(30) **Eventivity Constraint** (first version)

Must with an eventive prejacent has an obligatory root interpretation

Part of what seems to be going on, which we're aiming for (50) to capture is the fact that it seems like context doesn't seem play quite the role we were expecting it to.[10] I will later consider whether (50), or any of its amendments, creates any problems for the Standard Account. For now, we'll note that it's extremely hard to get an epistemic reading of (21), repeated below as (31), regardless of any context we can envision.

(31) John must go to the store.

(31) has an accomplishment predicate, but the pattern persists no matter which kind of eventive predicate is used.

[10]The behavior the EC describes is gestured at in Hacquard [2011]. EC behavior is noted by Lekakou and Nilsen [2008] in Greek, and Drubig [2001] discusses a related pattern, but to my knowledge the only extended semantic treatment of the EC is in Ramchand 2014, 2018. The earliest explicit acknowledgement of the pattern described by the EC seems to be Steedman 1977.

(32) **Activities**:
  a. John must run.
  b. John must swim.
  c. John must push a cart.
  d. John must drive a car.

(33) **Achievements**:
  a. John must recognize the answer.
  b. John must spot his car.
  c. John must reach the summit (by this afternoon).
  d. John must die.

(34) **Accomplishments**:
  a. John must paint a picture.
  b. John must make a chair.
  c. John must deliver a sermon.
  d. John must draw a circle.
  e. John must recover from an illness.
  f. John must build a house.

And, yet, *must* with stative predicates have no problem getting an epistemic reading, in addition to a root reading (in fact, for some of the examples in (35), the epistemic reading is much more prominent than any root reading).

(35) **States**:
  a. John must know the answer.
  b. John must believe the report.
  c. John must desire a vacation.
  d. John must understand French.
  e. John must be happy.

The generalization seems fairly robust across these various predicates, suggesting that whether or not the the predicate is stative or eventive is what underlies the generalization. However, when going through the examples cited above, one could begin to doubt the robustness of the generalization proposed so far. After all, you might notice that you

can conceive of a context where a modal sentence like *John must run* can indeed get an epistemic reading – simply think of a context where you hear that John can run a six minute mile. Knowing that such a time is incredibly difficult to get unless you are a seasoned runner – someone who runs regularly – you infer, and then utter, (32a). In fact, many of the activity predicates are susceptible to this kind of interpretation – an epistemic one, to be sure. From this, it's a short step to noticing that there's a something like a general recipe to "generating" these kinds of contexts, even for the other aspectual classes, as well as for our initial (21). Another example, (34c), *John must deliver a sermon* could also admit of a kind of epistemic reading, with a little effort. Envision a context where you are aware that John has some event scheduled tomorrow that involves public speaking, but we're not sure what it could be. You offer that you know he will be at church tomorrow, so I respond with (34c).[11]

My contention is that these examples fall into two categories: habituals and futurates. To be sure, the aspect of the lexical predicates in these expressions (*deliver a sermon* and *run*, etc.) are in fact eventive. But if we take a look at the semantics of these expressions, we come to see that they have a stative semantics in spite of the fact that they contain lexical eventive predicates. So, we're confronted with a claim advanced by Verkuyl [1972], that expressions larger than simply a VP (and so larger than the most deeply embedded predicate) will have an aspectual classification. In this case, what we find is that habituals and futurates are a kind of derived state, and so while these examples may have *seemed*

[11]The more natural way of putting it would be *John must be delivering a sermon*, and both this, and (34c) improves dramatically with the presence of a temporal adverbial, like *tomorrow*. But even with such an adverbial, it's clear that this may pose a challenge to our putative generalization, because the modal is epistemic while the predicate is eventive.

like an exception to the generalization proposed, they are actually quite consistent with it.

2.3. First Constraint: Habituality

Let's return to the original example, (21), repeated below as (36).

(36) John must go to the store.

On its most natural reading, (21/36) is perhaps teleological – so, one of the root flavors. That its most natural or prominent flavor is a root flavor comports with the suggestion that there is a generalization in the neighborhood here whereby eventive predicates pattern with root flavors. However, discussion of examples like (32a) might begin to tell against this generalization. Above, I suggested a particular kind of recipe for getting an epistemic reading for the likes of (32a). If we turn our attention to the likes of (21/36), the "recipe" alluded to above allows for the conditions that make an epistemic interpretation of this sentence accessible, in spite of the eventivity of *go to the store*.

The habitual reading which underlies the epistemic interpretation of the modal is most easily evoked with the help of a bit of contextual background and a frequency adverbial. (Think: *John must go to the store (Mondays/ often/ every day after work).* Or, even more evocatively; *Even though John's refrigerator is always empty over the weekends, it is invariably stocked full when I stop by his apartment Tuesday mornings. John must go to the store Mondays.*) In saying that the predicate is interpreted habitually, the contrast here is to the episodic reading of the predicate. To make this contrast vivid, we may first note that sentences can describe information about particular events, as in (37).

(37) John went to the store.

Abstracting away from features like tense, what makes a sentence like (37) true is the occurrence of a particular kind of episode or event, namely a going-to-the-store by John.[12] On the other hand, sentences such as (38) don't describe particular events, but rather regularities in the world which amount to generalizations over events.

(38) John goes to the store (Mondays/ often/ every day after work).

(38) has a habitual interpretation, and its truth isn't dependent on any particular event so much as a pattern of event-*types* (goings-to-the-store by John) that are asserted to hold sufficiently regularly.[13]

Two facts are important to the present discussion. First, habituals are often construed as a species of generic. More specifically, they are a variety of what Krifka et al. 1995 call *characterizing sentences*, one of the two basic varieties of genericity in natural language. Secondly, as the previous gloss on the truth conditions of episodic vs. habitual sentences makes clear, the distinction between habitual and episodic sentences reflects a semantic difference. That a robust semantic difference in the prejacents underlies the difference in the root and epistemic readings of (21) is easy to overlook due to the fact that, in English, habituality has the same linguistic form as the simple present. Linguists like Dahl 1995 point out that many languages opt to express habituality through the least marked tense-aspect choice available in the language. In other languages, it just so happens that this distinction is grammaticalized. That is to say, we might not notice that this

[12]To be sure, (37) *also* has a habitual reading, as in the following exchange. A: Throughout the past year, John had a relentlessly routine schedule. He dedicated each of his weeknights to a different chore. Mondays he did his laundry. Tuesdays he vacuumed his apartment. B: What did he do on Wednesdays? A: He went to the store.

[13]Cf. Carlson 2005 for discussion.

epistemic interpretation of (21/36) is different from the non-epistemic interpretations, simply because English doesn't grammaticalize this difference.[14]

Given the preceding discussion, we can meaningfully ask whether, in the epistemic reading we recognize in (36), the prejacent isn't just in fact the sentence (38). This would seem to be the case, given the consensus in semantics that habituality is a robust semantic phenomenon and is thought to warrant a mechanism representing this meaning at the level of logical form: The difference between a habitual and episodic interpretation of the prejacent in (36) is reflected in the semantic representation of (36) at a suitable level of abstraction. While the precise truth-conditions of generic sentences are a thorny matter, semanticists generally agree that genericity is represented in logical form by means of a phonologically null variable binding operator, GEN.[15] The epistemic interpretation of (36) all but vanishes on an episodic interpretation of the predicate.

The lesson here is that this way of getting an epistemic interpretation of (36) requires that (38) be the prejacent. According to standard proposals concerning generics, GEN is a dyadic operator, so generic sentences are also partitioned into a restrictor and a scope. The logical form of (38) is something like the following:[16]

(39) GEN(x, e)[x= John & e involves x] [Going(e, John) & To(e, the store)]

[14]There are ongoing debates about the semantics of generics, of GEN, and about whether the variable-binding GEN is even the best way to model generics. That said, the claim that habituals are a species of generic might be questioned. But the overall lesson of this section will be that habitual sentences are stative, not eventive – regardless of the aspectual class of the embedded predicate. This is easily explained (and most easily illustrated) if an operator, like GEN, is responsible for this aspectual shift. But if it turns out that habituals are not best modeled by means of GEN, the overall point – that habituals are stative – is not thereby affected.

[15]Cf. Krifka et al. 1995. Also cf. Leslie 2008 and Leslie 2013 for discussion. The operator was originally proposed by Heim 1982.

[16]Here I simply adopt the logical form suggested by Krifka et al. [1995].

If this is so, this means that the difference between the non-epistemic interpretation of (36) and this particular epistemic interpretation of (36) hinges on a difference in the prejacent. Holding fixed the event semantics employed in my gloss on the truth conditions of characterizing sentences, the difference between this epistemic and the teleological interpretation of (21) is made visible in (40).[17]

(40) a. EPISTEMIC: MUST(R)(GEN(x, e)[x= John & e involves x][Going(e, John) & To(e,the store)])

b. ROOT/ TELEOLOGICAL: MUST(R)(∃e(Going(e, John) & To(e, the store)))

All of this suggests that the epistemic interpretation of *must* will covary with the habitual reading of the predicate in the prejacent. When the predicate gets a habitual reading, the epistemic reading of *must* is available. On the episodic reading, a root reading is mandatory. Aside from forcing a root reading, there seems to be no further constraint on which of the root flavors are allowed when the predicate has an episodic reading. So, while I glossed the most likely reading of episodic (21) as teleological, an appropriate context may yield a deontic reading of (21) as well.

A final point to note. There is an aspectual difference between the episodic and the habitual readings of the predicates. Work on lexical aspect makes clear that habitual sentences have a stative semantics.[18] In fact, one of the very tests for characterizing sentences is that they seem to lose their characterizing sense when put in the progressive, unlike the episodic interpretation. Consider (41).

[17]As far as a representation of the possible logical forms of (36), this is clearly pretty coarse-grained. In fact, in attempting to gain insight into the logical form at a finer level of grain, it will turn out that (40b) will need to be amended. (Cf. Chapter 4.1 for discussion.) But for now, this coarse-grained representation will suffice to make the present point.

[18]Cf., e.g., Smith 1991, Carlson 2005

(41) a. Timo drinks wine with his dinner.

b. Timo is drinking wine with his dinner.

(41a) has a natural "characterizing" reading, expressing regularities and not reporting on a particular event. However, when (41a) is transformed into its progressive form, it tends to lose the characterizing interpretation. (41b) is not longer easily understood as a habitual, but as a report on a particular event of Timo drinking wine, understood to be on-going. This pattern is expected insofar as the progressive itself tends to exclude statives.[19] So, the thinking goes, this provides evidence for the claim that characterizing sentences are habitual. Like other statives, they tend not to appear under the progressive.

I take it as fairly obvious that this relevant reading, which clearly allows for *must* to get an epistemic interpretation, is a habitual or characterizing sentence embedded under the modal, simply in virtue of a reflection on its truth-conditions. But applying the progressive test to prejacents embedded under *must* corroborates this judgment. Consider what happens when we embed (41a) and (41b) under *must*. In (42a), the epistemic reading clearly involves a habitual interpretation of the prejacent. In (42b) we still get an epistemic reading (the progressive is itself thought to be a derived state), but the reading of the (42b) now differs from (42a).

(42) a. Timo must drink wine with his dinner.

b. Timo must be drinking wine with his dinner.

[19]As I've mentioned earlier, there are exceptions to both of these. In particular, there are lexical statives (so-called "interval statives" pointed out by Dowty [1979]) which can appear in the progressive. Likewise, the prohibition on habituals in the progressive has exceptions as well (Cf. Smith 1991, pp. 41 – 42). Still, even when these kinds of interval statives are in acceptable in the progressive, the meaning changes subtly, so with care it can still be used as a test for stativity. In particular, Deo [2009] shows that progressivizing lexical statives serves to give them a sense in which they are are understood to be more temporary.

This is precisely what we'd expect if a) in (42a), the habitual actually embeds under *must* – if the prejacent in *must* is a habitual with the characterizing reading of (41a), and b) the habitual is stative.

Since the inability for statives to appear in the progressive is more of a syntactic test for stativity, we can think of this as syntactic evidence for the stativity of habituals. Smith [1991] gives a semantic argument for their stativity by noting the following. They consist of an undifferentiated period rather than successive stages. If one were to investigate the truth of a habitual claim one would aim to uncover a pattern that held over the course of an interval rather than a particular situation. This makes the temporal schemata of habitual sentences stative.

The lesson to draw from this is that the sentences like (21/36) are are actually not evidence against the eventivity constraint – they're further confirming instances of it. Second, insofar as habituals are derived stative expressions, and the variable binder GEN is semantically responsible for this interpretation, we can then presume that the binding of the event variable by GEN is what renders the sentence with an otherwise lexically eventive predicate stative.

2.4. Second Constraint: Futurates

There is a second kind possible counterexample to the generalization posed the Eventivity Constraint, was raised by a consideration of (34c) (=*John must deliver a sermon*). The easiest interpretations to access for (34c) are surely root interpretations. However, certain contexts could make an epistemic interpretation of (34c) available. Since *deliver a sermon* is an eventive predicate, we would have an example of an eventive prejacent

with an epistemic *must*. In this case, the kind of context that brings out this reading is one where John was scheduled to be doing something or other involving public speaking tomorrow, and in weighing our evidence as to what it might be, we settle on (34c).

The interesting thing about this example is that having said this much gives us insight into what kind of sentence it would be *without* the modal. Suppose you asked me what John is supposed to be doing tomorrow and, remembering quite distinctly what's on his schedule, I utter: *John delivers a sermon (tomorrow)*.[20] The kinds of features we find in the context that allow us to access this reading – the fact that we are talking about the future, the requirement that the event under discussion needs to be planned, scheduled, or somehow *expected* to happen at a specific time – tell us something about the felicity conditions of this kind of utterance. And from these felicity conditions we can recognize clearly that we are talking here about a futurate sentence.

Futurates are sentences without future-referring verbal morphology that nonetheless refer to future events. So, while *John will leave tomorrow* is interpreted as concerning a future event of John's leaving in virtue of the auxiliary *will*, simple futurates (*John leaves tomorrow*) and present progressive futurates (*John is leaving tomorrow*) ostensibly talk about a future event without any verbal marking signaling future reference. However, as a condition on their assertability, they require a reading of the predicate such that

[20]The relevant reading here is *much* easier to get when the future adverbial is present. It's not impossible to get the reading without it, but it has to be clear from the context that the interval we're talking about is the one described by *tomorrow*.

the event it describes is planned or scheduled in advance.[21,22] So the idea here is that, in examples like (34c), we are actually dealing with a *futurate* that is in the scope of the modal.

As a survey of the eventive predicates we've considered so far will make clear, not *every* eventive predicate will be easily amenable to a futurate reading. This is unlike habituals – most eventive predicates admit of a fairly natural habitual interpretation fairly easily.[23] It took an example like (34c) to make the modal-over-futurate reading come into relief, and only because we could highlight the fact that the sermon delivery was planned to happen at a specific point in time. A futurate reading is not nearly as apparent for the prejacent in (21), but this comports well with the assertibility conditions of futurates. Had the prejacent contained an eventive predicate more easily construed as scheduled, then an epistemic reading of *must* would have been easy to come by, particularly if a temporal adverbial made a future reference time salient. Consider this next example in the following context. I have hosted John for a few days, see his packed luggage standing by the front door that evening, and surmise that his departure is imminent. I utter (43).

(43) John must leave tomorrow.

---

[21]Compare sentences like i. *My plane takes off this afternoon*, ii. *The Yankees play the Red Sox on Wednesday*, and iii. *I go to the doctor tomorrow* with i′. *#My plane crashes tomorrow*, ii′. *#The Yankees defeat the Red Sox on Wednesday*, and iii′. *#I accidentally meet the doctor tomorrow.* Cf. Copley 2008b, 2009 for discussion of the literature surrounding this kind of construction.

[22]Most, of the time, this requires a plan to have been made by some salient agent, where knowledge of the plan is contextually available. But there are futurates that cannot plausibly be the result of a plan by any agent, e.g. *The sun rises tomorrow at 6:34* AM. These types of futurates only seem to be available when it is contextually agreed upon that the event in question comes about through some expected, law-like process. Copley [2009] suggests that for such constructions, we can understand the law-like process as standing in for the plan, and the world as taking the role of the director of the plan.

[23]Unless, of course, the predicate describes an event that one cannot undergo on a regular basis.

With (43), the available readings include not only the deontic reading, according to which (roughly) the speaker claims that John is under an obligation to leave tomorrow, but also an epistemic reading, which was made salient by the envisioned context, in spite of the eventive predicate in the prejacent. The epistemic reading of (43) contains a futurate prejacent.

One would need to strain to access an analogous "scheduled" reading of (21). The relevant difference between (21)(=*John must go to the store*) and (43) is that leavings are easily thought of as scheduled (since they often *are* scheduled, as in the departures we associated with mass transit), so we don't require conversational machinations to accommodate this prerequisite. The difficulty in getting a non-habitual epistemic interpretation of (21) to gain traction is due to the difficulty in interpreting *go to the store* as plausibly scheduled. But even this isn't absolute – suitable cooperation by the context can over-ride this tendency, as evidenced in (44).

> CONTEXT: You and I are are discussing an upcoming party John is hosting the day after tomorrow. We know that John intends on buying a lot of supplies for the party. We see that his cupboards are still bare, so we attempt to determine when John will go to the store for supplies. You note that John has to work late today and wonder aloud what his plans are for tomorrow. I retort:
>
> (44) John must go to the store (tomorrow).[24]

Once the context allows for an interpretation of the predicate as a description of a planned event, we can have an epistemic reading of *must* with a non-habitual, eventive prejacent. Moreover, it is clear that the eventive predicates that more readily admit of epistemic

[24]Many speakers, myself included, would probably be more likely to utter a progressive futurate in the envisaged context: *John must be going to the store tomorrow.* But (44) is still acceptable as an epistemic modal sentence, which is what matters to the present discussion.

readings under *must* are precisely those that admit of scheduled readings rather easily. So, the evidence suggests that in those epistemic interpretations of *must* with non-habitual eventive prejacents, the prejacents are in fact futurates.[25] On the surface, we see little evidence of a similar constraint for the types of predicates that can embed under root *must*.[26] At the very least, it's clear that they do not need to be scheduled or planned. The significance of this asymmetry harks back to our discussion of habituals in the following two ways. First, the most prominent account of the semantics of simple futurates posits a special kind of futurate operator in logical form responsible for triggering the appropriate reading.[27] Assuming a non-habitual, epistemic reading of (21), we would then be justified in suspecting a difference in the logical form between the epistemic and the root sentences. Rendering this futurate operator as FUT, the logical form of the root reading would be as above in (40b), whereas the logical form of the epistemic reading under discussion would be something like (45).[28]

(45) MUST(R)(FUT (John go to the store))

[25]This suggestion is also pressed by Klecha 2016 and Ramchand 2014.

[26]Though cf. chapters 4 and 5 for considerations complicating this picture, when we turn our attention to root modals.

[27]Cf. Copley 2009 for a modal construal of FUT. Also, Kaufmann [2005] argues for a similar modal element. Copley [2009] suggests that FUT in simple futurates is actually a species of GEN. If this is so, the present discussion could in principle be assimilated to the previous section. Cf. Copley 2014 for a non-modal, "causal-chain" analysis of FUT. This analysis will become important later on in later chapters of the dissertation.

[28]In this rendering of logical form, I abstract away from the event semantics employed in (40). In future sections, we'll have the occasion to look at the logical forms in more detail. For now, I'm interested in representations at a rather coarse level of granularity, insofar as these are suggested by the apparent asymmetries in root and epistemic interpretations of *must*.

Second, the construction resulting from the application of FUT is a derived stative.[29] Some evidence for this is similar to the the evidence for the stativity of habituality. First, there is the conceptual evidence that, insofar as futurates seem to require a plan of some sort, and having a plan involves having a volition, and having a volition is a kind of state (cf. Copley 2014). Second, there is a direct analogy to habituals. We do in fact see futurates in the progressive, so we have to modify the test somewhat. (There is a (present) progressive form of futurates, often simply called "progressive futurates". By contrast futurates in the simple present tend to be called "simple futurates".) But we may note that simple futurates and progressive futurates exhibit a meaning contrast, exhibited in (46) and (47).

(46) a. John gets married tomorrow.

b. John is getting married tomorrow.

(47) a. Mary lives in Paris.

b. Mary is living in Paris.

(46) is a futurate whereas (47) is a habitual. The contrast between the (a) and (b) is such that the (a) sentences convey that the plan or the habit is rather permanent or long-standing whereas the (b) sentence conveys that the plan or habit is more temporary.

Further examples of the stativity of futurates comes from their temporal behavior in indicative conditionals, as pointed out by Copley 2009. The following tests are due to Copley [2008a], who notes that in conditionals, eventive antecedents cannot have a

[29]In fact, the stativity of futurates and the fact that they seem to require the presence of some kind of plan or schedule has led some linguists to suggest that the plan should be reified, receive its own Davidsonian argument, and that this newly introduced eventuality pronoun is a state. Cf. Dowty [1979], and Copley [2014, 2018] paper for a worked out proposal along these lines. For now, let's just note that, much like with habituals, futurates produce a derived state, even when the embedded predicates are eventive.

present orientation, and they force the consequent to be future-oriented.[30] So, you cannot modify the consequent with an adverbial, like *now*, that denotes the anchoring time (for simplicity's sake, we assume the anchoring time is the time of utterance).

(48) a. *If John goes to the store, he will be out of flour now.

b. If John is at the store, he will be hungry now.

c. If John leaves tomorrow, he will pack his bags now.

The reason (48a) is bad is that the antecedent, being eventive, will have a future orientation; the putative going-to-the-store event would follow the utterance time. Then, since the event described by the consequent would follow the going-to-the-store event, modification with *now* is incompatible with this temporal constraint. In (48b), since *be at the store* is stative, it is present oriented, and therefore the consequent can be modified by *now*. If (48c) is truly eventive, then we'd expect it to pattern with (48a). But unlike (48a), (48c) is actually fine, and patterns with the stative antecedent (48b), suggesting futurate antecedents are in fact stative.

Also due to Copley is the "It's true that..." test. "It's true that" indicates the time of evaluation of the antecedent, and since eventives force a future-orientation on the antecedent, combining this locution with *now* will be bad for eventives, but OK for statives. On the assumption that futurates are stative, we would predict that antecedents with a futurate would be compatible with "it's true now that...", like stative antecedents and unlike eventive antecedents. In fact, this is what we find.

[30]They are always future-oriented with respect to the anchoring time; which, let us assume is the time of utterance. Often, eventive consequents will force the temporal orientation of the consequent to be subsequent to the time of the antecedent, as is in *If John runs the race, he will win.* But as Copley points out, the consequent can be contemporaneous with the antecedent, as in *If you push the ball, it will move.*

(49) a. # If it's true now that John goes to the store,...

b. If it's true now that John is at the store,...

c. If it's true now that John leaves tomorrow,...

(49a) is bad, except, say, in philosophical English where we are discussing future contingents. In fact, we can borrow a page from the discussion of future contingents for an explanation of why (49a) is bad, but (49b) and (49c) is not. Since eventive antecedents in indicative conditions have a future orientation, so if the event being described by the antecedent were to happen at all, it would happen in the future. But the *it's true now that...* forces a present evaluation on the antecedent, resulting in anomaly.[31]

Statives, by contrast, can have a future- or present-orientation in the antecedents of indicative conditionals. So, stative antecedents can be evaluated with respect to the present – that's what it means for them to be present-oriented. As a consequence, *it's true now that...* is fully felicitous with stative antecedents. Another way of saying this is that *it's true now that...* seems to select a present-orientation of the antecedent. But eventive antecedents don't have a present-oriented interpretation available, so they are ruled out. What we see in this is that with a futurate antecedent like in (49c), this is absolutely fine. (49c) patterns with the stative (49b) in this respect and not with the eventive (49a) in spite of the eventivity of its embedded predicate, further suggesting that futurates are derived states. Just like with habituals, we see considerable evidence that futurates are in fact (derived) states, and so far from threatening the generalization underlying the Eventivity Constraint, examples like (34c) further confirm it.

[31]Or perhaps worse, ungrammaticality. Copley suggests that *it's true now that...* requires a proposition for evaluation and eventive antecedents aren't fully propositional. Such an explanation make the badness of (49a) look more like ungrammaticality.

In light of the discussion of these two constraints on how *must* with an eventive prejacent have an epistemic reading, it would make sense to revise our initial formulation, which seemed to take it for granted that it was the aspectual class of the lexical predicates that was at issue in whether or not *must* could have an epistemic reading. The two further constraints – concerning habituals and futurates – show that it has less to do with the aspectual category of the lexical predicate in the prejacent, but rather the aspectual properties of the expressions larger than the most deeply embedded verb. In particular, since habituals and futurates are both themselves aspectually stative, they both conform to the eventivity constraint, provided we consider the aspectual behavior of the entire expression.

That this is so should be evident from the fact that lexically eventive predicates can be converted into stative predicates by means of a habitual or futurate operator. So, we should reformulate the Eventivity Constraint to take account of this fact.

(50) **Eventivity Constraint** (second version)

Must with a bare eventive prejacent has an obligatory root interpretation

We capture this in the newly formulated version of the Eventivity Constraint in (50). By bare eventive prejacent, I mean one without an intervening expression or operator that produces an output that is a derived state. We'll hold on to this formulation of the Eventivity Constraint for now, as a description of the apparent phenomenon, adjusting as necessary when further insights become available.

CHAPTER 3

Diagnosing and Explaining the Eventivity Constraint

In this chapter, I will make a case for the claim that what I've glossed earlier as the Eventivity Constraint poses challenges to the most popular accounts of the semantics of modals. In particular, it serves as a constraint on a satisfactory theory of modals that it cannot predict or explain the Eventivity Constraint – and the Standard Account cannot meet this challenge. To start with, I will discuss the challenge the EC poses to the Standard Account already introduced in Chapter 2. Thereafter, I will discuss some alternative frameworks for the semantics of modals and discuss how the Eventivity Constraint comes to bear on these.

3.1. General Remarks on the Constraints

The motivation behind explaining the above constraints on an epistemic interpretation of the likes of (21; repeated below as 51) was to show that, to the extent that an epistemic reading was possible, the availability of the reading conformed to a predictable pattern.

(51) John must go to the store.

Epistemic interpretations of (51) and sentences that were similar in virtue of having eventive prejacents were admittedly not nearly as prominent as the root readings, and needed either considerable help from the context or the addition of certain kinds of temporal adverbials to make them more easily accessible. The above discussion was not intended as a fail-safe algorithm for yielding an epistemic interpretation for every *must*-sentence with

an eventive prejacent. Instead, it was intended to show that, to the extent that an epistemic interpretation was even available, it was only where the prejacent was an instance of a habitual or a futurate.[1] Tying an epistemic interpretation of (21) to the habituality or futurity of its prejacent has the following consequence. The epistemic reading of *must* is predicted to only be as prominent as the respective habitual or futurate readings of the prejacents. In contexts where these futurate readings are hard to get, the epistemic reading is correspondingly predicted to be hard to get as well, which is exactly what we see when we reflect on the data.

It also has further significant consequences. If context were directly responsible for the epistemic interpretation of the modal in (21/36), then we ought to be able to evoke this epistemic reading merely by manipulating features of the context. Recall that this is the model for context sensitivity suggested by the Standard Account. The parameters relevant to determining the modal domain are the modal base f, and its world argument w. Yet the value of w will simply be whatever the world of evaluation is – the actual world for matrix contexts, and a shifted world in cerrtain embedded contexts. And f is given a value contextually. So it would follow that f should be whatever the context allows, and whether a modal has an epistemic or deontic reading would then be a matter of context. Yet, evoking an epistemic reading for (21/36), as with all *must*-sentences whose prejacents have eventive predicates, seems to require interpreting the prejacents as habituals or futurates, and is difficult in precisely those contexts that don't lend themselves to habitual or futurate interpretations of the prejacents. Moreover, these two constraints

[1]In the last chapter, I suggested that there was a "general recipe" for making such readings accessible. It turns out that this recipe just involved envisioning contexts whereby the prejacents could have habitual or futurate readings. But I'd resist saying there's a general algorithm for generating such reading because certain predicates may resist attempts to give them such readings.

exhaust the options for epistemic readings of (21/36). Admittedly, obtaining a futurate or habitual interpretation of a predicate seems to require a good degree of cooperation by the context of utterance, as evidenced from the contextual gerrymandering I employed to evoke these readings. But this should not obscure the fact that, if our current accounts of habituals and futurates are on the right track, the habitual/ futurate interpretations of (21/36) trace to differences in the logical form from their episodic and non-futurate counterparts. Simply put, in spite of the fact that the string of words exemplified in (21/36) can have an epistemic reading, the nuclear scope of the modal is in fact different when the modal is interpreted epistemically. To the extent that context plays a role in securing these interpretations, it does so largely in a disambiguating role with respect to the prejacent, to use the terminology of Stanley and Szabó [2000]. The evidence we would need to attribute the epistemic reading of (21) to contextual factors is a case where the nuclear scope is unchanged with respect to its logical form for both the root as for the epistemic readings, and the *only* difference is some feature of the discourse context. So far, we haven't seen such evidence.

Since the habitual and futurate prejacents are stative and their episodic counterparts are eventive, it would appear that epistemic *must* does not tolerate eventive predicates in the nuclear scope of the modal. Whatever the cause of this restriction, it also persists under various embeddings, as exhibited in (52) where the salient reading throughout is a deontic/ root reading.

(52) a. Attitude verbs: Mary thinks John must go to the store.

b. Indirect discourse: Mary said that John must go to the store.

c. Antecedents of conditionals: If John must go to the store, then...[2]
d. Consequents of conditionals: If we are out of milk, then John must go to the store.

In light of the persistence of this effect, the null hypothesis is that the eventuality-type of the predicate in the prejacent has something to do with this restriction. The formulation settled on last chapter was given in (50), repeated here as (53).

(53) **Eventivity Constraint** (second version)

Must with a bare eventive prejacent has an obligatory root interpretation

Though the observation embodied by (53) may be subject to further reformulations and precisifications, we can use the EC to probe the empirical and explanatory adequacy of theories of modal semantics. In this chapter I will focus on what I've laid out as the Standard Account and some of its variants, but it should be fairly obvious how the criticisms I leverage apply to a standard context-index theory.

Though I will continue to focus on *must*, it is worth noting that *must* is not unique in conforming to the EVENTIVITY CONSTRAINT (henceforth "EC"). *Ought* and *should*, to the extent that they admit of epistemic readings, seem to conform to the EC as well.[3] I will discuss these auxiliaries in a subsequent chapter. *May* and *might* seem, at first blush, not to conform to it. The important point is that we have something of a robust generalization applying to a range of modals. The most immediate concern is that the

[2]It's possible that this observation is not very probative, since it has been observed that epistemic modals tend to be bad in the antecedents of conditionals anyway. However, cf. Hacquard and Wellwood 2012 for putative acceptable instances of epistemic modals in the antecedents of conditionals. Such instances seem to conform to the Eventivity Constraint.

[3]Cf. Yalcin [2016] for arguments that these modals do not in fact have epistemic readings. For Yalcin, these are pseudo-epistemic readings; he proposes a semantics where a normality ordering is imposed over a set of worlds characterizing an information state. I don't wish to quibble over whether, if this is an apt characterization, "normality" modals form a natural class with properly epistemic ones. Suffice it to say, that such an account is not inconsistent with anything I've said.

Standard Account fails to predict this pattern in the interpretation of these modals. In the next section, I will discuss why this is.

3.2. Diagnosing the Problem

The EC spells trouble for the Standard Account, since the latter predicts that context can conspire to produce values for f and g such that $[\![(21)]\!]^{w,f,g}$ is epistemic on an episodic reading of *go to the store*. Once we set aside habitual and futurate prejacents, the epistemic reading of (21/36)(=*John must go to the store*) is simply not attested. We lack an explanation of this apparent restriction on the interpretation of the modal.

Recalling the motivating paradigm, if we analyze modals as quantifiers over sets of possible worlds, we need some kind of restriction on the domain of quantification. To that end, the Standard Account posits a contextual parameter that restricts the modal's domain of quantification. One of these parameters, the modal base *f*, is responsible for determining whether the modal has an epistemic or a root flavor.[4] So a good place to look for an explanation of the EC is *f*. The story is that context contrives to secure values for *f*, yet there is no prima facie reason to suppose context is prevented from assigning a value for an epistemic modal base to it. After all, f is just a function that takes a world argument, and there is nothing in the recipe provided by the Standard Account that indicates when the value of the function should yield an epistemic as opposed to a circumstantial domain, aside from the needs of the conversational context. The EC makes it seem like the choice for the value of f is systematically restricted given certain

[4]Recall the observation that the determination of flavor for a modal utterance involved a division of labor between f and g, but that having a particular kind of modal base value (as either epistemic or circumstantial) was sufficient for determining whether the modal would have an epistemic or root interrpretation.

sub-sentential properties of the prejacent in the scope of the modal. It would be as if in MUST ϕ, the internal complexity of ϕ impacts the interpretation of MUST. However, as pointed out earlier, it was difficult to identify a particular semantic property responsible for this constraint, so it is yet unclear how the internal complexity could be responsible.

To see why this is a problem for the Standard Account, consider what it takes for f to get a particular value. Let's think of f_{ep} and f_{circ} as two different function-types, the definitions for which are given in (18a) and (18b) (repeated below as (54a) and (54b)).

(54) a. $\bigcap f_{ep}(w) = \{\text{w}'|\ \text{w}'$ is compatible with what is known by the relevant agent(s) in $w\}$

b. $\bigcap f_{circ}(w) = \{\text{w}'|\ \text{w}'$ is compatible with certain circumstances relevant in c in $w\}$

On the basis of the EC, we want to say that bare eventive prejacents somehow only allow the parameter f in (77c) to get a value of the type f_{circ}. The reason this is so difficult to accommodate on the Standard Account comes into relief in considering what would determine whether the modal base of a modal is f_{circ} or f_{ep} in a given context.

Exactly how this would work is partly a reflex of the details of the semantics' implementation. For completeness's sake, let's consider both implementations. On the first implementation, the one I adopted for the exposition of Kratzer's semantics, f and g were simply parameters of the interpretation function, $[\![\cdot]\!]$. That is to say, the interpretation of modal sentences is relative to these parameters (in non-modal sentences, they are vacuous); an interpretation is possible once values have been assigned to them. However, the idea is just that they are given a value pragmatically; this is the significance of Kratzer's

calling them "conversational backgrounds." Consequently, according to this implementation, whether the modal base is f_{circ} or f_{ep} is a matter of context. Of course, the EC shows us precisely that does not hold in full generality.

On the second implementation, f is treated as an object language variable in the level of syntactic representation of the sentence that serves as the input to semantic interpretation.[5] Like any free variable (think of unbound readings of pronouns, like deictic readings of *he* or *she*), it needs to be given a value in order for the modal sentence to be interpreted. This is done by the assignment function, a.[6] The assignment function determines whether f is assigned a value of type f_{circ} or f_{ep}. Treating f as an object language variable may then seem like an advance insofar as its value assignment is mediated by a, but this doesn't help with the EC. The appropriateness of a given assignment is determined by the context of utterance.[7] Absent a special condition on the appropriateness of variable assignments that manages to rule out assignments according to which f gets an f_{ep}-type value, this would again incorrectly predict that there are epistemic readings of EC modals with bare eventive prejacents. Of course there is no such reading. Moreover, what would such a special condition be, short of a mere specification of the prohibition? And what justification would we claim for the special condition, aside from a desire to render the theory empirically adequate? This maneuver really just recapitulates the data by building

[5]This is in fact the preferred approach in much of the current linguistics literature. Cf. von Fintel and Heim 2002.

[6]The assignment function is one of the coordinates of the interpretation function, $[\![\cdot]\!]$. So, the interpretation is defined relative to those coordinates. It is typically rendered 'g', but I've rendered it 'a' to avoid confusion with the ordering source.

[7]Cf. Heim and Kratzer [1998] (pp. 243 – 244), who place the following appropriateness condition on LFs with free pronouns:

> A context c is *appropriate* for an LF ϕ only if c determines a variable assignment $[a_c]$ whose domain includes every index which has a free occurrence in ϕ.

an ad hoc condition into the appropriateness conditions for assignment functions. Such a restriction has no explanatory value.[8]

3.3. Accommodating the Data with a Selectional Restriction

Though she is not addressing the data I'm concerned with here, Kratzer [1981] raises the possibility of the lexical entries for modals coming with selectional restrictions – that is, with restrictions on the kind of arguments they can take. Accordingly, one may think that selectional restrictions give one an easy way to explain the Eventivity Constraint. In this section, I will argue that the best explanation of the Eventivity Constraint ought not rely on selectional restrictions.

An early account in generative linguistics for how lexical items are inserted into syntactic structures had it that items appear in the mental lexicon as a structured bundle of features, and that some lexical items put restrictions on which sorts of words they can take as arguments, over and above their subcategorization requirements.[9] For a simple example, consider (55).

(55) John drank a brick.

---

[8]Here's a way to implement an explicit restriction on the assignment function that would have genuine explanatory value – if modal bases had Phi features the way pronouns do. A variable assignment could be ruled out if it assigned an object/ person to the pronoun *she* that didn't have the appropriate gender. Likewise, a variable assignment could be ruled out if it assigned to f an epistemic modal base when the prejacent was eventive. This would be genuinely explanatory *if* there were antecedent reasons to believe (1) that modal base pronouns had Phi-features, and (2) there were Phi-features like "having an eventive prejacent". I know of no such reasons (indeed, such a Phi-feature is far from the kind of features typically countenanced by theories of agreement, like person or gender).

[9]Cf. Chomsky 1965 for an early explanation of the role of selectional restrictions. Cf. Freidin 1992 for a clear, text-book introduction.

Since *drink* is a transitive verb, according to the grammar, it should receive a Noun Phrase argument as an object, and *a brick* is in fact a Noun Phrase.[10] But there's something strange about (55) nonetheless; a brick isn't a thing you can drink! (55) is syntactically well-formed, at least insofar as its verb takes a Noun Phrase complement, and transitive verbs subcategorize for this kind of argument. Nonetheless, it is unacceptable.[11] The idea here is that the verb *drink* takes an argument that must have something like [+liquid] as one of its lexical features. *Water* has this feature (among others, like [–inanimate]), and so the lexical item *water* can be inserted into the N slot of the object NP argument to *drink*. You could extend the idea to *must* by saying that epistemic *must* puts a similar kind of restriction on its arguments. In this case, epistemic *must* would be restricted to taking complements that have a [–eventive] predicate. However, only epistemic *must* would have this restriction. Deontic *must* would not have it, because it can have a [+eventive] complement. This should raise some concern since this approach would suggest that there are two *must*s with different restrictions. However, if we are respecting UNIFORMITY, we are assuming that there is only one *must*.

If the appeal to a selectional restriction follows the traditional model I've sketched above, where a word or lexical item selects another kind of item based on the presence or absence of a certain feature, then this appeal to a selectional restriction for *must* is just a

[10]I haven't said anything about what constitutes a grammar, and won't currently take a detour through such a discussion. For the sake of the example, we can presume something like a phrase structure grammar which includes the rule VP $\rightarrow$ V + NP.

[11]Chomsky's famous sentence *Curious green ideas sleep furiously* is a relevant example. Part of the import of this sentence was that, though it is perfectly syntactically well-formed, but it is purportedly uninterpretable. Selectional restrictions aim to do some work in explaining why this is. Note that selectional restrictions are not overridden on figurative interpretations of sentences. The theory would have it, rather, that in figurative interpretations of sentences, lexical items are required to interpret certain sentences like (55), in order to respect the selectional restriction.

capitulation of UNIFORMITY. While the contribution *must* makes to the logical form is the same for root as for epistemic *must*, one flavor of *must* has a feature the other lacks. The advocate of a selectional restriction is smuggling in a difference-maker surreptitiously.[12] Unlike *might*, whose idiosyncrasy does not result in multiple lexical entries, this proposal would split the lexical entry for *must* into one entry for $must_{epistemic}$, and one for $must_{root}$, with the former having a unique selectional restriction as one of its features.

Perhaps one could avoid this contravention of UNIFORMITY by positing a relational selectional restriction. The relational selectional restriction for *must* would presumably be something like (56).

(56) RELATIONAL RESTRICTION: if a modal sentence S has the form MOD(R) ϕ, where ϕ is the nuclear scope of S, MOD is *must*, and ϕ is [+eventive], then f in R is $f_{circumstantial}$.

Since (56) applies uniformly to *must* it conforms to UNIFORMITY. Nonetheless, as an explanation of the patterns of modal interpretation captured by the eventivity constraint, invoking (56) is deeply unsatisfying. For one thing, it simply restates the data. We also see no independent motivation for it. As a *place-holder* for a more robust explanation, I have no objection to (56). If, on the other hand, we simply stipulate that (56) is part of

[12]There is another worry, which I'll mention just to put aside. The features I allude to in my discussion of (55) concern features *lexical items* have as part of their entries in the mental lexicon, if you will. The kind of feature one would need to appeal to to leverage selectional restrictions in explaining the eventivity constraint are of a different sort. They would be features of *complex expressions*. The advocate of this response owes us a story about where these features come from, and if they come from the lexical items themselves, how complex expressions come to inherit them from the lexical items that compose them. The theory of lexical insertion I appealed to in my explanation of selectional restrictions is outdated and long superseded. More contemporary versions of Phi Theory or analyses dealing with agreement phenomena may be able to provide a story about the distribution of such features in the clause. (Cf. fn. 8.) However, as far as I am aware, none of these extant theories appeal to a feature like [+eventive], and the introduction of such a feature for the purpose of explaining the eventivity constraint would be *ad hoc*.

the lexical meaning of *must* in order to explain the eventivity constraint, the stipulation is *ad hoc*.

There is a further, methodological reason to be wary of an explanation by RELATIONAL RESTRICTION. Selectional restrictions were originally underwritten by a theory of lexical insertion which explained how lexical items are inserted into the syntactic structures they figure in, but this theory of lexical insertion is no longer in favor. So, it had better be that positing a restriction of this sort is independently motivated, since selectional restrictions themselves are no longer an integral part of some piece of grammatical theory. Otherwise it is utterly mysterious what just what this restriction is.

3.4. Uniformity and the Demand for an *Explanation* for the EC

I submit that the Standard Account over-generates; the EC shows how the Standard Account does so in a rather systematic way. I've shown that this is attributable to the mechanism in the semantics which yields the modal domain; the modal base parameter. I've also shown how the most apparent ways of remedying this over-generation problem are either unsatisfactory, or they contravene UNIFORMITY.

In response to this problem, one might think, "So much the worse for UNIFORMITY! The EC gives us a semantic reason to think modal auxiliaries do not have a uniform semantics, at least along the root/ epistemic dimension. Let's not hinder theorizing with such a commitment." It is not my goal to give a full-throated defense of UNIFORMITY in this dissertation – I take it for granted that being able to give a uniform semantics is a virtue, and investigate the compatibility of this desideratum with strategies for explaining the EC. But while the EC does pose complications for UNIFORMITY, this pessimistic

attitude is premature. The remaining sections of this chapter will consider how revisions of the Standard Account might be able to account for and explain the EC. Indeed, I settle on one such account as providing us with the tools to do so.

One might think, by way of a second response, that my claim that the EC poses a problem for the Standard Account is overblown. In fact, such an argument would go, the Standard Account can have a perfectly good response to the data posed by the EC, one which, moreover, would predict the data. This putative response is simply to build the restriction into the lexical entry of the modal, as it were. The idea here is similar to the one floated in the last section, where a condition is placed on acceptable variable assignments. Here, the idea is that a condition of this type is recommended by the lexical meaning of the word. The suggestion comes to this: we can build the Eventivity Constraint into the lexical meaning for *must* by stipulating that *must* is not well defined when its prejacent is eventive and the f is circumstantial. With such a lexical constraint, the rest of the lexical entry is exactly as we suggested in Chapter 1.

(57) a. $[\![\text{MUST } \phi]\!]^{w,f,g}$ is only defined if f is circumstantial or ϕ is stative

b. if defined, $[\![\text{MUST } \phi]\!]^{w,f,g} = \forall w' \in \text{BEST}_{g(w)}(\bigcap f(w))\text{: } \phi(w') = 1$

This suggestion is reasonable, but I think ultimately unsatisfying. For one thing, it buys us the ability to predict the EC at the cost of what appears to be another *ad hoc* stipulation. It appears *ad hoc* because it simply integrates the EC as a condition on the interpretation of the modal. Moreover, it does nothing to explain the EC. In response to my complaint, my interlocutor may object that semantics just needs to model the truth conditions of the target expressions in the fragment of natural language the semantic account is concerned with; demanding it *explain* the derivation of those truth conditions is misplaced. After

all, there may not be an explanation, properly speaking, for why an expression has the truth conditions it has. And building the restriction into the interpretation in the manner above signals that the meaning of the expression is messy and idiosyncratic.

As above, I think this theoretical posture is premature. To start with, the notion of explanation I'm invoking is not particularly deep. It would suffice for the semantic account to give some indication as to why the relevant property stipulated in the restriction is at all relevant to the interpretation of the modal, perhaps by appealing to the semantic or syntactic properties that figure in accounts of aspectual class or Aktionsart. A restriction of the kind in (57) falls short of that. Now, it might turn out that there is nothing more to be said about the aspectual class of the prejacent and its relation to the modal. But that would be strange, considering that the behavior captured by the EC turns out to be unexceptional as opposed to idiosyncratic – in the sense that it doesn't require special embedding conditions to exhibit itself. Moreover, it turns out that EC behavior is not a mere quirk of English. Though I won't be pursuing a rigorous cross-linguistic comparison of modals in service of this point, I can note some suggestive comparisons to other languages. Lekakou and Nilsen [2008] note a similar pattern with the Greek modal *prepi*. The German modal *müssen* exhibits similar behavior. There is also a similar pattern in Russian. The Russian necessity modal *dolžen* is a copula+participle construction roughly glossable as "obliged is"; effectively the translational equivalent of *must*. In (58), with an eventive prejacent, the modal can only have a root interpretation.[13]

(58) Vanja dolžen pojti v magazin
Vanja must.PRES go.pfv to store

[13]Thanks to Sophia Malamud for the Russian example and further discussion in her comments to my 2016 Pacific APA talk.

'Vanja must go to the store'

While these few examples by no account provide a systematic cross-linguistic picture of the behavior of modals, it does at least suggest that what is represented by the EC is nor merely a quirk of English. Were it a mere quirk, construing EC behavior as an idiosyncratic lexical restriction on *must* would be justified.[14] A consideration of the data doesn't warrant the suspicion that the EC or EC-type behavior are idiosyncrasies of English, so the account of it ought not treat it as such. Building the restriction into the lexical entry for the modal as a condition on its interpretation suggests we are dealing with lexically idiosyncratic behavior when the hunch that we have systematic behavior on our hands is more justified.

3.5. Condoravdi [2002] on the Modals' Sensitivity to Aspect

Condoravdi [2002] famously notes that the interpretation of modals can be affected by the aspectual properties of the predicates of their complements. In the semantics she gives, she builds on the Kratzerian account to account for both the temporal behavior of modals and the particular way in which the aspectual properties of their complements impact this behavior. Condoravdi doesn't address the EC directly, but given the circumstances described above, it's natural to ask whether her semantics could either predict or explain the EC. If so, we would have an off-the-shelf explanation of the EC that comports with Kratzer's paradigm in modal semantics. This section will evaluate the possibility of leveraging Condoravdi's semantics for such an explanation.

[14]Appeals similar to the one I am making have been lodged against ambiguity accounts of modals more generally. A methodological consideration against ambiguity accounts of modals is that they model what is better represented as systematic behavior of expressions in a way better suited to lexical accidents. (Cf. the work cited in fn. Hacquard 2011.) I think the same kind of methodological ambition ought to apply here.

To start with, Condoravdi makes two distinctions that are essential to her account. The first one concerns the difference between epistemic and metaphysical readings of modals. According to Condoravdi, they are both non-root modals in the sense discussed above, with epistemic modals having an epistemic modal base, as in our previous discussion, and metaphysical modals having a metaphysical modal base, the value of which is an equivalence class of worlds consisting of the set of worlds that are identical in history through some time *t*. (Condoravdi uses the term 'modal base' for what I've been calling 'modal domain', but this kind of ambiguity is unproblematic, and common in the literature.) In matrix contexts, a metaphysical modal base delivers the worlds which are historical alternatives to the world of utterance, at the time of utterance.[15,16]

The second distinction is between what Condoravdi calls the temporal *perspective* and temporal *orientation* of a modal.[17] The temporal perspective of a modal is the time at which the worlds in the modal base are calculated. The temporal orientation of a modal is the relation between the temporal perspective and the time of the described event. For root modals, since they have a circumstantial modal base, this would be the time at which the relevant facts hold which help determine the set of worlds against which the modal is evaluated.

Condoravdi is concerned to maintain a uniform account of the temporal semantics of non-root modals in light of the differing temporal behavior of certain constructions. Though her concern is different in detail from the concerns of the present chapter, to the extent that she aims to maintain a uniform semantics for modals in light of the temporal

[15]Cf. the branching time framework of Thomason 1970.
[16]One might have doubts about whether metaphysical modals are really not a variety of root modal after all, but I put this concern aside. Cf. Abusch 2007 for similar concerns.
[17]Cf. also [Matthewson, 2012], to which I owe my brief summary, for a concise formulation.

behavior of modals, and their aspectual sensitivity, it is in sympathy with the present aim of finding an explanation of the EC that may respect UNIFORMITY. What she ultimately argues for is that non-root modals uniformly have a present perspective and a future orientation. However, the temporal perspective of the modal can be shifted by operators of various sorts, changing the apparent orientation of the modal. Switching momentarily to the modal *might* to illustrate this point, in (59), the relevant reading is made especially prominent by the adverbial *still*.

(59) The team might (still) have won the game.

On an informal gloss, (59) is true on the relevant reading just in case, at some point in time prior to the time of utterance, it was still possible for the team to go on to win the game. As the game progresses, and the possible avenues for the team's victory were winnowed down, any historical alternative wherein they win was no longer open to them, and they consequently lost. This reading of *might* has a temporal orientation which is backshifted with respect to the time of utterance (the perspective is some time in the past). Yet, Condoravdi maintains her uniformity thesis by arguing that the backshifted reading is attributable to a scope reversal between the modal and the operator contributed by the Perfect, PERF. On her analysis, PERF scopes over the modal, and while the modal still has a present perspective and future orientation, its present perspective is shifted by PERF (so it is a kind of present in the past, not an indexical present) and it is future oriented with respect to that perspective.

This scope reversal involving PERF is the first way Condoravdi suggests that temporal behavior of modals can be affected by neighboring elements in the clause. The second way that the temporal behavior of modals is affected by other elements in the clause is

more central to my concerns. The precise nature of a modal's temporal orientation also depends on the aspectual type of its complement. Condoravdi specifies the translation of temporal operators in terms of the AT-relation, where this varies depending on the nature of its third argument (which, for our purposes, is the prejacent).

(60)

$$\mathrm{AT}(t,\ w,\ P) = \begin{cases} \exists \mathrm{e}\ [P(w,\ \mathrm{e})\ \&\ \tau(\mathrm{e},w) \subseteq t] & \text{if } P \text{ is eventive} \\ \exists \mathrm{e}\ [P(w,\ \mathrm{e})\ \&\ \tau(\mathrm{e},w) \circ t] & \text{if } P \text{ is stative} \\ P(w)(t) & \text{if } P \text{ is temporal} \end{cases}$$

That is, if P is eventive, then P standing in the AT-relation to t and w is a matter of the interval of P's holding in w to be included in t. If P is stative, it is a matter of P (in w)'s overlap with t. (P is temporal if it is a property of times.) We can then use the AT-relation to define the lexical entries for modals according to the following schema, where 'MB' stands for 'modal base'.[18]

(61) a. Possibility modal: $\lambda P \lambda w \lambda t\ \exists w'[w' \in \mathrm{MB}(w,\ t)\ \&\ \mathrm{AT}([t,\infty),\ w',\ P)]$

b. Necessity modal: $\lambda P \lambda w \lambda t\ \forall w'[w' \in \mathrm{MB}(w,\ t) \rightarrow \mathrm{AT}([t,\ \infty),\ w',\ P)]$

The present tense operator PRES identifies the free occurrences of t with the interval *now*, as in (62).[19]

(62) PRES: $\lambda P \lambda w[\mathrm{AT}(now,\ w,\ P)]$

[18]Condoravdi ignores the ordering source in these semantics – this would map onto the more classic Kratzer semantics by simply giving an empty ordering source, where $g(w)$ maps onto the empty set.

[19]The semantics for PERF is the following:

i. PERF: $\lambda P \lambda w \lambda t' \exists t'[t' \prec\ \&\ \mathrm{AT}(t',\ w,\ P)$

As can be seen, PERF shifts the time of evaluation of the expression in its scope to an interval before the reference interval. In the kind of scope reversal that characterized the reading of (59) I discussed earlier, the evaluation time for the modal is shifted back, giving it a past perspective. However, the modal is still future-oriented – uniform future-orientation being essential to Condoravdi's proposal here. The resulting semantics gives the kind of "future-in-the-past" interpretation that characterizes the 'metaphysical' reading we intuitively attribute to (59).

On Condoravdi's analysis, the arguments of modals are tenseless, with tense scoping over the modal, resulting in the following schema for modal sentences generally: PRES (MODAL(P)).

Taking into account the different instantiations of the AT-relation for eventive and stative prejacents, modal sentences result in the following kinds of derivations. Following (Krifka 1989), the function τ maps events (or events and their worlds of instantiation, as construed here) to their run-times. Let's use *be at the store* and *go to the store* as our paradigmatic eventive and stative predicates alike.

(63) a. He might be at the store.
b. he be at the store: $\lambda w \lambda$e [he be at the store](w)(e)
c. MIGHT$_{MB}$(he be at the store):
$\lambda w \lambda t \exists w'$ [$w' \in MB(w,t)$ & $\exists$e [[he be at the store]] (w')(e) & τ(e,w') $\circ$ [t, ∞)]]
d. PRES (MIGHT$_{MB}$(he be at the store)):
$\lambda w \exists w'$ [$w' \in MB(w,now)$ & $\exists$e [[he be at the store]] (w')(e) & τ(e,w') $\circ$ [now, ∞)]]

(64) a. He might go to the store
b. he go to the store: $\lambda w \lambda$e [he go to the store](w)(e)
c. MIGHT$_{MB}$(he go to the store):
$\lambda w \lambda$t$\exists w'$ [$w' \in MB(w,t)$ & $\exists$e [[he go to the store]] (w')(e) & τ(e,w') $\subseteq$ [t, ∞)]]
d. PRES (MIGHT$_{MB}$(he go to the store)):
$\lambda w \exists w'$ [$w' \in MB(w,now)$ & $\exists$e [[he go to the store]] (w')(e) & τ(e,w') $\subseteq$ [now, ∞)]]

(63) and (64) show how the truth conditions differ subtly for prejacents with stative and with eventive predicates. For stative prejacents, the temporal trace of the P-eventuality overlaps with the interval starting *now* and extends infinitely into the future. However, if the prejacent is eventive, it is included in that same interval. The stipulation that the temporal relation involved for events is one of inclusion rather than overlap has

important consequences. For one thing, since events are temporally extended, the requirement that the event be included in the interval [*now*, ∞) guarantees that the event be in the future of *now*.[20] For it to be properly included in the interval, it must start, at the earliest, at *now* and extend to some point in the future. Since stative eventualities merely *overlap* with the evaluation interval, they can start at some point before the interval and persist into the future, and still satisfy the overlap requirement. This allows for stative eventualities to be contemporaneous with the time of utterance, where no such possibility exists for events. This feature of the semantics thereby accounts for the obligatory future-orientation of eventive predicates compared to the merely optional future-orientation of stative predicates.

These are Condoravdi's semantics, in a nutshell. The question facing us is how the analysis would rule out epistemic readings of the likes of MUST ϕ, where ϕ is eventive. Though her analysis is sensitive to the effects of the prejacents' aspectual class on the interpretation of modals, nothing in her analysis rules out the readings targeted by the EC.[21] However, based on the resources afforded by her semantics we could surmise how an explanation of the lack of these readings would go. Since, by assumption, ϕ is eventive, the temporal requirement with respect to eventive prejacents has the ϕ-eventuality included in the interval starting at the utterance time and extending into the future. As discussed, this puts the ϕ-eventuality in the future with respect to the utterance time. The definition

[20]What about achievements? These are momentary changes of state, so properly speaking, they occur in an instant. Condoravdi doesn't consider these specifically, but I think it's proper to treat them as being temporally extended since, for an achievement to take place, a transition from one state to another has to occur. If we think of the resulting event as encompassing the minimal interval including the final instant of the beginning state and the initial instant of the target state, achievements will be temporally extended.

[21]To be clear, explaining or predicting the EC is not the task of the paper.

of necessity modals glossed above has it that all worlds in the modal base (i.e., all worlds consistent with what is known by the relevant agent(s)) are worlds where the ϕ-eventuality holds at some point in the future.

What might rule out epistemic *must*, given these truth conditions? Holding fast Condoravdi's semantics, a few possible explanations come to mind.

3.5.0.1. An Argument from Veridicality. We could hazard an explanation on the assumption that *must* is veridical. If *must* is veridical, as von Fintel and Gillies 2010 would have it, ⌜MUST ϕ⌝ entails ϕ. But if the future is open, then the kind of knowledge that would satisfy such truth conditions is perhaps exceedingly rare, since it would require knowledge of the inevitability of the eventuality described by ϕ.[22] In fact, the kind of circumstances that could underwrite such knowledge are circumstances where the ϕ-eventuality is treated as though its eventuating is settled (e.g. as though through a prior plan). But such circumstances are precisely those that support the felicity of futurate constructions, as outlined in section 2.4. So, this line of reasoning would go, the conditions under which an epistemic reading of MUST ϕ is true either simply does not exist for lack of appropriate foreknowledge, or else it cedes to the futurate construction. Absent acceptable truth conditions, MUST ϕ has an obligatory root reading, when ϕ is eventive.

Against this argument, we might note that the interpretation of *must* as veridical is controversial; in arguing against it von Fintel and Gillies position themselves as arguing against a "mantra". Indeed Kratzer 1991 accepts the mantra of non-veridical *must*. Though I'm partial to veridical *must*, one might worry about an explanation for the EC that hinges on a controversial analysis. Secondly, and more speculatively, to my ear,

[22]Or, at least, that ϕ's eventual occurring is entailed by one's current knowledge, which would amount to ϕ's inevitability, so far as one is aware.

modals like *should* and *ought* seem be EC modals.[23] If indeed *ought* and *should* are EC modals, then the explanation founders, because unlike *must*, *should* and *ought* are never thought to be veridical, even on their epistemic reading. But whether or not *ought* and *should* are EC modals is controversial – Ramchand [2014, 2018], for example, denies that they are – and we haven't discussed other modals in any kind of systematic detail, so let's put this question aside for now, coming back to it later. For now, we may note that *if ought* and *should* turn out to be EC modals, then an explanation which combines the veridicality hypothesis with Condoravdi's semantics will founder.

3.5.1. An argument from the unknowability of the future

There is another potential explanation one could leverage here, and is exhibited by Klecha [2016]'s claim that the future-oriented epistemic necessity modals are ruled out because the future is in some sense inherently unknowable. This claim allows one to provide a similar kind of explanation as the veridicality argument, though this argument is in some ways stronger. It's stronger because it's not simply the case that, on the basis of MUST's veridicality that we are unwarranted in making an utterance of the form MUST ϕ. Rather, any future-oriented necessity modals would seem to fall under this prohibition as well – even "weak" necessity modals like *ought* and *should*. The type of ordering source can ensure that the *should*-sentence is not veridical in the sense outlined above. For example, if the worlds in the modal domain are ranked by an ordering source characterized by a stereotypicality ranking – according to how things stereotypically unfold.[24] This would be unveridical since it can turn out that the world isn't among those ranked highest by

[23]That is, it seems like *John ought to go to the store* doesn't have an epistemic reading.
[24]Cf. Kratzer's stereotypical conversational background, 2012, p. 37.

the stereotypical worlds. But all that, since the modal under consideration is epistemic, it would still amount to an utterance of how the future will turn out, given a restriction a constraint that the future conform to stereotypes of events tend to unfold. But if the future is inherently unknowable, then any claim that links future outcomes to our knowledge of how things will turn out will be at the very least infelicitous, whether restricted via a "non-veridical" ordering source or not. So, this explanation rules out *must* with eventive prejacents not on the grounds of veridicality, but because they are future-oriented. On account of their future-orientation are epistemic readings of (21) ruled out.

Though there is something to this thought, this won't do as an explanation to the EC for two reasons. First, the explanation is too strong. Though it's a rather vague claim to say that the future is in some sense inherently unknowable, whatever sense this is meant is sufficient to explain that it would be infelicitous or otherwise unacceptable to utter a statement that claimed that it followed from what one knows (suitably restricted by an ordering source). But the same considerations would seem to rule out future-oriented epistemic *possibility* modals as well – if one can't say that it follows from what one knows that X, it is equally bad to say that any future event's occurring or future state's obtaining will be consistent with what one knows. If the future is inherently unknowable, then epistemic claims about the future will be out full stop. But the data doesn't suggest that this is the case – future-oriented *might* sentences seem to be quite OK, as evidenced by the *might* version of (21).

Secondly, the claim that the future is inherently unknowable is a philosophical claim – a claim that relies on premises about the metaphysics of time, determinism, and knowledge. For example, one way this claim has been motivated is through a commitment to the Open

Future – the view that statements about the the future are neither true nor false. If WILL ϕ is neither true nor false, then WILL ϕ can't be known. But if WILL ϕ follows from what one knows, then one could know that WILL ϕ, simply by reasoning from what one knows (provided the entailment isn't overly obscure). But since, by supposition, one can't, it must be the case that WILL ϕ doesn't follow from what one knows. Now, the kinds of indetermimist intuitions that motivate the Open Future are shared by many people, in particular about future contingents. And so this may seem to provide a promising avenue to ruling out future-oriented epistemic *must* claims. Nonetheless, if you ask "the folk" for their intuitions, it is likely that they would indeed impute bona fide knowledge to themselves of some future events, and not only for non-contingent claims. For example, most people will likely say that they *know* that the sun will rise tomorrow. Absent the metaphysical appeal to the open future, it is difficult to fault people on these intuitions. One of the observations about epistemic *must* common in the literature is that a *must*-sentence seems weaker than than the unmodalized version of that sentence. *It must be raining* is in a nebulous sense to assert something "weaker" than *It is raining.* This is often dubbed "Karttunnen's Problem" [25] Now, there are various attempts at explaining what accounts for Karttunnen's observation, but what matters for our purposes is that they all tend to share the idea that direct perception of an event (like a raining event)

[25]Cf. Karttunen 1972, von Fintel and Gillies 2010, Mandelkern 2018, 2019b, Goodhue 2017. I put "weaker" in scare quotes here because the putative strength or weakness of *must* is usually put in terms of veridicality: strong *must* is veridical *must*, and weak *must* is non-veridical *must*. However the kind of weakness observed here, often called, has more to do with the conditions under which an assertion of a *must* sentence is appropriate. Looking out the window and seeing the rain, it would be appropriate to say *It is raining*, but strange, and inappropriate, absent special conditions, to utter *It must be raining.* Proponents of strong/ veridical *must* recognize Kartunnen's Problem as a genuine phenomenon, and aim to explain it. In fact, proponents of veridical *must* will sometimes claim that considering *must* weak/ non-veridical is a misdiagnosing of Karttunnen's Problem.

make it ill-suited to be described by means of a *must*-sentence (like *It must be raining*). The *must*-sentence is felicitous only if one has some indirect evidence of the event.

But if one knows that the sun will rise tomorrow, in virtue of it occurring in the future, one *has* to know this through indirect means (or what have you). So *The sun must rise tomorrow* seems like a perfect candidate for a felicitous *must* sentence. However, for all that, it is still exceptionally hard to get an epistemic reading of *The sun must rise tomorrow*.

You might object that although the folk might think that they know that the sun will rise tomorrow, they are nonetheless mistaken, objectively speaking, and the semantics for epistemic modals should be based on this objective fact. Two comments in response to this; first of all, there is some controversy over whether the information states epistemic modals quantify over are really *epistemic*, making for truth conditions that track what an agent (for group of agents, etc.) knows, or whether so-called epistemic modals really quantify over information states that are more properly thought of as *doxastic*, in that they qualify over what a person believes. And if that is the case, then the folks' insistence that they know that the sun will rise tomorrow (moreover, by indirect means) should be sufficient for a felicitous utterance of *It must rain tomorrow*. But it is not. Secondly, supposing the first option is right, then it seems like the proper take-away from this insight is not that future-oriented *must* sentences are ruled-out, but that we should judge them to be incorrect. If people think that they know the sun will rise tomorrow, but really they do not, then the appropriate theoretical posture to take would be that these utterances are false.[26] But this doesn't seem to be the case – the data isn't that *It must rain tomorrow*

[26]Future oriented modal claims would then look like the example in Hacking 1967 where an agent makes a modal claim about the location of a sunken ship. In Hacking's scenario, not only is the ship not at

is always false on the epistemic reading, it's that there really isn't an epistemic reading to begin with. Caution is warranted about this last claim – given the right context and discourse conditions, you *can* get such a reading. However, if you pay attention to the context which allows for an epistemic reading of this, they are precisely the conditions which suggest that a futurate construction is in the scope of the modal.[27] This is itself an interesting constraint, in light of the discussion in section 2.4 which suggested that futurates involve some sort of stativizing operator. Why the necessity of this operator for the acceptability of future-oriented *must*-sentences; even ones where speakers would typically impute knowledge to themselves? In any event, the lesson here is fairly clear; we can't simply combine Condoravdi's semantics with a claim about the unknowability of the future in order to explain the Eventivity Constraint.

3.5.2. A final comment about Condoravdi's account of aspectual sensitivity in modals

There is another worry, which doesn't immediately bear on the explanation of the EC, but brings into question whether the way the aspectual dependencies are cast in Condoravdi's semantics captures the right truth-conditions. Specifically, the temporal relation that determines the nature of the temporal orientation for stative sentences under modals is too weak. The semantics allows the prejacent to be true at *any* point within an interval

---

the suggested location, but the speaker fails to take into account information that would rule out this possible that is readily available in a log the speaker had examined. (Cf. Teller 1972 for similar kinds of examples.) In this case, objective features of the scenario (the availability of the information in the log) make the sentence false in spite of the ship's suggested location being consistent with what the speaker knows. In our future-oriented case, the objective feature that would make the *must*-sentence false is the unknowability of the future.

[27]Copley [2009] calls these kinds of futurates "natural futurates", since they aren't scheduled by people, but there is something regular and schedule-like about their occurrences.

overlapping the time of utterance. That would mean, for a sentence like (63), *He might be at the store*, it would be true if there is a world in the modal base where the referent of *he* is at the store two weeks from now. The truth conditions are presumably set up in this way because we can easily evoke this meaning with the adverbial *two weeks from now*. Without such an adverbial, (63) is simply present-oriented; (63) would be false if *he* were not at the store at the time of utterance. So, for all its insights into the aspectual dependency of modals, I do not think Condoravdi 2002 will serve as a firm enough basis for an explanation of the EC.

3.6. The Event-relative Approach to Modals

So ends the critical portion of the chapter. I will now argue that the lack of epistemic readings of *must* with eventive prejacents in fact has quite a simple explanation, once we allow a small departure from the Standard Account. Seeing this involves taking a finer-grained look at the semantics of tense and aspect and their integration with modals. The result of this integration accounts for half of an explanation of the EC. The other half of the explanation, explaining why root *must* with eventive prejacents is acceptable, becomes a more delicate matter on the account I will advocate. But first the easy part. The way forward involves taking note of apparent difference in syntactic position between root and epistemic modals, what Hacquard 2010 calls 'Cinque's Puzzle'. Following the work of Cinque 1999 on the relative positions of functional heads, it is largely thought that the relative syntactic positions of modals is as follows (irrelevant projections omitted).[28]

(65) $\text{MODAL}_{epistemic}$>TENSE>ASPECT>MODAL_{root}>vP.

[28]Recall that Condoravdi takes non-root modals to scope under tense, so her account is at odds with Cinque's hierarchy.

These syntactic positions impose scope restrictions on modals. To illustrate a poignant example of this, epistemic modals are often said to scope higher than tense and aspect, whereas root modals scope below tense and aspect. Since *must* does not inflect for tense, this scope restriction is difficult to detect in *must*-sentences. It becomes apparent with the semi-modal *have to*, which does inflect for tense. (66) has both a root and an epistemic interpretation.

(66) Deborah had to be at the train station.

(Epistemic gloss) 'It is necessary, given what is known now, that Deborah was at the train station.

(Root/ teleological gloss) 'It was necessary, given Deborah's circumstances then, that she be at the train station.

As the glosses on (66) make clear, the evaluation time of the modal is not modified by tense on the epistemic reading, whereas it is on the root reading. On the epistemic reading, the modal is evaluated with respect to the salient information state as it is *now*, at the time of speech, amounting to a present temporal perspective, to use Condoravdi's terminology.[29] By contrast, on the root reading, the modal is evaluated with respect to the circumstances *then*, as though shifted by the past tense on the modal expression. This is consistent with the modal being in the scope of the past tense, as (65) puts it.

These syntactic differences are something of an embarrassment to the Standard Account, insofar as it aims at UNIFORMITY. The difference in syntactic location illustrated in (65) suggest that there might be at least some lexical differences between root and

[29]Discourse effects may make another perspective available, as can certain kinds of adverbials. I put these complications aside, since they do not affect present point.

epistemic modals, at least with regards to what licenses their different positions in the clausal architecture. To account for these data, Hacquard 2010 proposes a revision of the Standard Account, which I will adopt in what follows. Here the paradigm of treating modals as restricted quantifiers over sets of possible worlds is maintained, except that modals take an event argument as opposed to a world argument. Suggesting that modals take an event argument allows Hacquard to explain the two different syntactic positions for modals as follows. First, she assumes a Davidsonian event semantics where verbs are treated as predicates of events and introduce event variables into logical form.[30,31] Second, she posits an operator representing illocutionary force in the logical form of the sentence.[32] Hacquard proposes that this illocutionary operator is a predicate of events much like in the verbal domain (in this case, though, a speech event or illocutionary event). This makes at least two event variables available in the logical form of the sentence; a high event variable associated with the speech or utterance act, and a low event variable introduced by the verb. In making modals take event arguments, she is able to correlate the height of the modal with the availability of an event variable for the modal to take as argument. Hacquard derives the different flavors from the event argument the modals take. 'Low' modals are anchored to the event variable introduced by the vP. On Hacquard's picture, modal base projections from vP events yield a circumstantial domain, and correspond

[30]Cf. Davidson [1967]'s foundational arguments, and Higginbotham 1985 for the classic proposal to implement Davidson's event semantics compositionally.

[31]Although I made a commitment to using the cover term "eventuality" to talk about events and states (or, events, processes, and states, for those who make the further distinction), it is conventional to talk about "event semantics" and "event variables" even when it is understood that these may also stand for states. When following this convention would cause confusion, I will disambiguate.

[32]Cf. Krifka 2001 for an example of this. Theorists like Krifka self-consciously distinguish their approach from the 'explicit performative' approach to mood as in Lewis 1970, which is widely thought to be inadequate. Representing the illocutionary act in the semantics is meant to account for certain noteworthy embedding facts such as left-dislocation, as opposed to a semantic representation of mood or clause-type.

with root interpretations of the modal. 'High' modals are anchored to the speech event, and such modal base projections yield an epistemic domain.[33]

So far I've only noted that modals can sit in two syntactic locations and that on Hacquard's approach, this is attributable two different event variables in the clausal structure that the modal base can take as arguments. These facts alone do not justify the claim that high modals yield an epistemic domain and low modals a circumstantial one. In fact, there is a reason for this distribution. Certain predicates of events (namely, those events associated with speech or utterance events, and those associated with propositional attitudes) are thought to be associated with content, whether the content of the illocution or the attitude. We use CON as a function that is defined when e has propositional content. Then CON(e) denotes the content of e. This innovation allows us to reconstrue modal domains as follows.

(67) a. $\bigcap f_{ep}(e) = \{w' |\ w' \text{ is compatible with CON}(e)\}$

b. $\bigcap f_{circ}(e) = \{w' |\ w' \text{ is compatible with certain circumstances of } e\}$

According to Hacquard, for f to be epistemic, its event argument needs to be a 'contentful' event; CON(e) needs to be defined for f's argument if f is to be of type f_{ep}.[34] This explains why there are no 'low' epistemic epistemics.[35]

[33]Speech act operators are not the only way to motivate the presence of a speech event in the logical form of the sentence. Proponents of "Austinian Propositions" make use of a situation or event variable at the clause level to represent the speech act (Cf. e.g., Recanati 2007, Kratzer 2008). So, Hacquard's proposal to make modals take an event or situation argument is not married to the putative need for illocutionary operators in logical form. Indeed, Kratzer has adopted Hacquard's proposal on this count in more recent work (Cf. Kratzer 2013). It is for this reason that I labeled the target view "the Standard Account". Kratzer herself is no longer committed to the view in precisely its original formulation.

[34]Cf. Pietroski 2000 for independent arguments for such a function.

[35]Even though they are not anchored to an utterance event, modals directly under attitude verbs can also have epistemic interpretations, because attitude verbs like *think* or *belief* themselves introduce an event

Hacquard's revision of the Standard Account allows us to keep the semantics of modals uniform. It requires no major revision of the lexical entry for modals, so it preserves UNIFORMITY. My contention is that it allows us to integrate the semantics of tense and aspect more seamlessly into the modal semantics and thereby makes explaining the EC more tractable. Recall that the problem was that f needed to yield a circumstantial modal base if the prejacent was eventive, but that we couldn't make the selection of a value for f appropriately sensitive to the Aktionsart of the prejacent, and appeals to context didn't do the trick. If f gets an event argument, there is another way forward. On Hacquard's event-relative approach, the choice between f_{ep} and f_{circ} is formally constrained by the type of event argument available in the logical form. In turn, the type of event argument available in logical form is a matter of clausal architecture. This allows us to recast the analytical question as follows: why should a bare eventive prejacent prevent *must* from occupying the 'high' position? That is, why is (68a) good and (68b) bad?

(68) a.

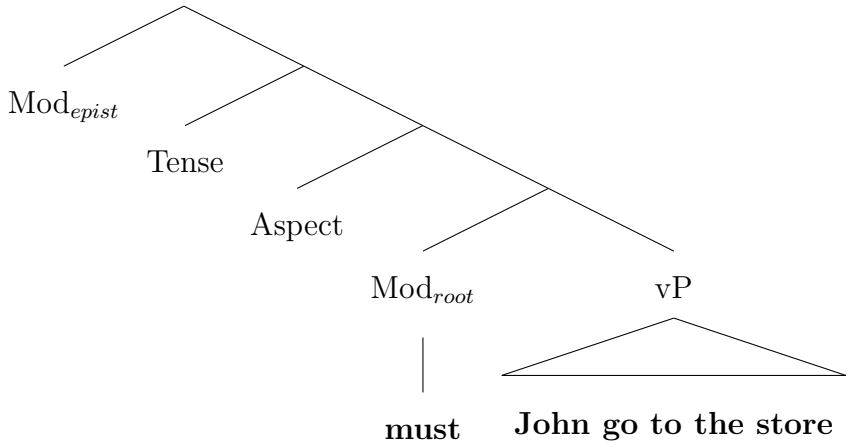

variable, and the believing-event or thinking-event are "content-ful" events. I omit further discussion of attitude verbs to avoid the added degree of complexity.

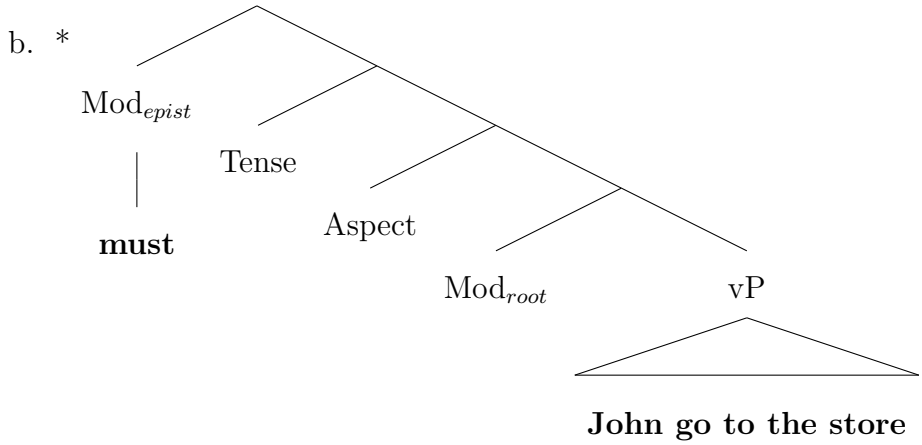

I will next show how recasting the question in this way gives us the means to explain the EC while preserving UNIFORMITY.

3.7. Explaining the EC

Our new way of putting the question essentially asks why (68b) is bad. Answering this requires saying more about what sits in the Tense and Aspect positions. Tense is fairly familiar; we can assume that (21)(= *John must go to the store*) has present tense. The Aspect head is where grammatical aspect is realized in the clause.[36] Determining what sits in the Aspect slot of (68b) requires some more background.

Aspect plays a role in many theories of tense. vPs build up event descriptions though the saturation of a verb's arguments and a series of optional modifiers, adding more predicative material onto the event variable introduced by the verb. On a popular conception of tense, associated with Reichenbach 1980 and developed by Klein 1994, tense locates these events on a time-line with respect to some privileged point of reference called the *reference time* in Reichenbach's terminology, or the *topic time* in Klein's. Note that our

[36]Not to be confused with "lexical" aspect or *Aktionsart*, which underwrites the distinction between the aspectual classes of eventive and stative predicates that the EC presumes.

use of event semantics runs us into trouble with logical types. Tense relates predicates of times to a deictically privileged time (the time of utterance), whereas vPs build predicates of events. So construed, Tense takes an argument that has a different logical type than the event description built up by the vP. The solution is for Aspect to take a vP denotation and turn it into something appropriate for Tense to modify; so to take a property of events and turn it into a property of times. Another way of putting it: the event description built by the vP has its own temporal properties (events are temporal entities, after all, and occur at a particular time). This has been called the *event time* (for Reichenbach), or the *situation time* (Klein). Aspect relates the situation time to the topic time. Tense relates the topic time to the utterance time.

Semantically, grammatical aspect turns on the perfective/ imperfective contrast, and refers to the presentation of the structure of the eventuality described by the predicate. To draw on descriptive metaphors in currency since Comrie 1976, perfective aspect describes *complete* eventualities, without regard for their internal structure. Conceptually, perfective aspect packages the eventuality described by the predicate as a bounded whole, whereas imperfective aspect presents the eventuality as in some way incomplete or ongoing. To deploy some more of the standard metaphors, imperfective aspect describes eventualities *from the inside*, as it were, allowing language to account for the internal structure of the event.[37] In English, the imperfective is often marked by progressive morphology on eventive verbs, but habitual interpretations of the predicate are also instances

[37]For example, in a sentence like *While I was writing a letter, Esther walked in*, the dependent clause is marked for imperfective aspect by way of the progressive, allowing for an interpretation whereby the event of Esther's walking in took place within the interval at which the event of my writing a letter occurred. That is to say, the writing of the letter was ongoing or incomplete at the point at which the walking-in event occurred.

of imperfective meaning, and often occur on unmarked forms of the verb. Perfective aspect occurs here in the unmarked case; the case where no imperfective marker is present, and habituality is ruled out.[38]

Having ruled out habituality in our characterization of bare eventive prejacents, the prejacents at issue in the EC have perfective aspect. A more complete representation of the structure of (21) would show that there is perfective aspect on the Aspect node, and the eventuality described by the vP is therefore represented as a completed event. (68b) is then more appropriately filled out as (69).[39]

(69) *

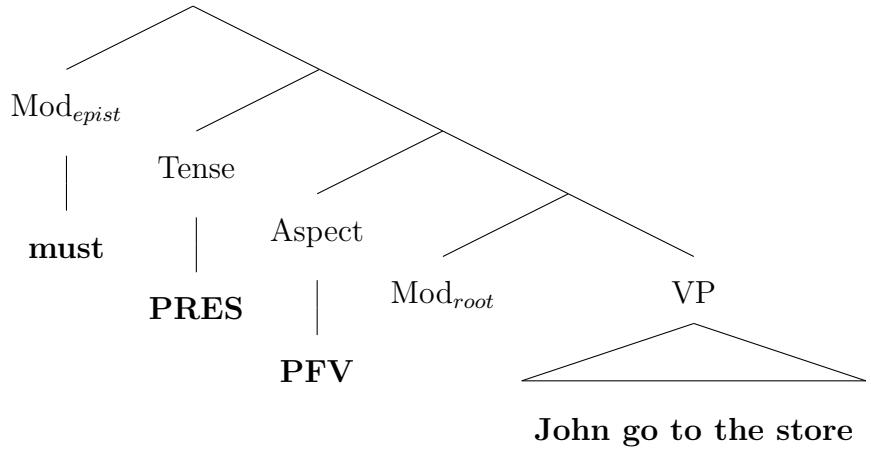

Upon adopting Hacquard's account, the question that came into relief was why (21) cannot have the structure in (68b), now filled out as (69). The "bare" eventive prejacents that constrained *must* from having an epistemic interpretation are actually *perfective* eventive prejacents in the present tense.

[38]Cf. Bohnemeyer and Swift 2004 for discussion. So-called "telic" languages differ from English on this. For these languages, unmarked telic verbs (accomplishments and achievements) get a default perfective reading.

[39]Here, as before, I include the syntactic positions (like Mod$_{root}$) that remain unoccupied. On some syntactic frameworks, such a representation is problematic, but I intend it as a harmless idealization to keep track of the relative positions of the projections we are interested in.

We know, independently of the present discussion, that perfective eventive sentences in the present tense are unacceptable. Consider (70).

(70) a. John goes to the store.

b. Mike eats an apple.

c. Mary pushes a cart up the hill.

d. Tim wins the race.

The sentences in (70) have an attested habitual reading (in addition to a possible futurate reading, on a scheduled interpretation of the predicate). There is an unattested reading of (70), which concerns our explanation, and which has the form $\ulcorner$PRES [PFV(ϕ)]$\urcorner$, where PRES is the present tense operator, and PFV is a perfective operator. Were the sentences in (70) in the past tense, they would have had a fine perfective interpretation.

(71) a. John went to the store.

b. Mike ate an apple.

c. Mary pushed a cart up the hill.

d. Tim won the race.

It is clear that the lack of an acceptable perfective reading is particular to the present tense. On the version of the Reichenbach/ Klein view of tense sketched above, tense locates the topic time in relation to utterance time. Present tense identifies the topic time with the utterance time. But perfective aspect constrains the situation time in such a way that it must be 'completed'. According to Klein [1994], a more precise way of capturing this gloss is to think of the perfective as putting a constraint on the situation time such that it be included in the topic time. (Imperfective aspect, by contrast, issues

the converse constraint; that the situation time be included in the topic time.) But then we have a conceptual problem. It is sometimes assumed, going back to Taylor 1977, that the utterance time is an instant. This means that the eventive predicates, which we've described as taking time to occur, are supposed to occur within the topic time, which is to be taken as an instant in virtue of its identification with the utterance time.[40]

We now have the makings of a conceptual explanation for why (69) is bad. The constraint perfective aspect places on the event description built up by the vP is incompatible with the function of present tense. It remains to be seen how this looks semantically, to which I turn now. First, let's specify the semantics of perfective and imperfective aspect and the present tense. We want Aspect to play two roles. First, it should take a predicate of events and turn it into a predicate of times, and second, it should relate the situation time with the topic time. Following Klein, we need the perfective to place a constraint on the predicate of times such that the situation time is contained in the topic time. We can make use of Kratzer [1998]'s formalization of Klein's conception of tense and aspect as below. For present tense, we just take the topic time and identify it with the utterance time, schematized as t_u.[41]

[40]This is thought to hold as a conceptual fact, even though producing an utterance of any kind is in fact an event that takes time. The idea is that the speech event picks out an instant that is to serve as the utterance time, which in turn constrains the temporal properties of the intervals that can be identified with it, and therefore the events that can be contained in it. For example, Hallman [2009] holds the view that the the utterance time is an instant (one that shares the denotation of *now*). Ogihara [2007] gives a similarly motivated view. Some authors endorse a variation of this idea, even if they don't endorse the view that the utterance time is an instant. Bach 1981 and Parsons 1990 are examples of the view where the utterance time is a constrained interval. Giorgi and Pianesi [1997] take the view that the interval characterizing the utterance time is homogenous, which places a mereo-topological constraint on the type of events that can be included in this interval– ruling out eventive verbs. I'll henceforth adopt the view that the utterance time is an instant. This may be a bit of an idealization, but if so it still comports with a body of literature on the subject.

[41]It's not uncommon to have it be a presupposition of present tense that the topic time is the utterance time, and have the semantics be undefined if this condition is not met, as in Kratzer 1998. Instead, I

(72) a. **PFV**: $\lambda P.\lambda t.\lambda w.\ \exists \mathrm{e}[\tau(\mathrm{e}) \subseteq t\ \&\ P(\mathrm{e})(w)=1]$

b. **IMPF**: $\lambda P.\lambda t.\lambda w.\ \exists \mathrm{e}[t \subseteq \tau(\mathrm{e})\ \&\ P(\mathrm{e})(w)=1]$

c. **PRES**: $\lambda P.\lambda w.\ \exists \mathrm{e}[t = t_u\ \&\ P(\mathrm{e})(w)=1]$

We see from the logical forms above that PRES requires some aspectual operator to relate the event time to the topic time, else PRES is otiose. For the representation of the vP, *John go to the store* will be rendered as '**John-go-to-the-store**$'(\mathrm{e})(w)$'. We also need to specify where the Davidsonian event variable is existentially closed. Since we assume that Aspect performs the role of turning a predicate of events into a predicate of times, Aspect is a good candidate for performing this function. So, we also stipulate that the event variable introduced by the verb is existentially closed by Aspect. Finally, I'll make use of the assertion operator used by Hacquard. This operator is introduced high in the clause, in the CP layer, and rendered as a predicate of events, like verbs. It is defined as follows, which relates the assertive act to its content via CON.

(73) **ASSERT** $\mathrm{e}_0 = \lambda P.\lambda w[\textsc{assert}'(\mathrm{e}_0,\ w)\ \&\ \forall w' \in \textsc{con}(\mathrm{e}_0):\ P(w)=1]$

Then our derivation proceeds as follows:

(74) a. $[_{vP}$ John go to the store $] = \lambda w.\lambda t.\lambda e$ **John-go-to-the-store**$'(\mathrm{e})(w)$

b. $[_{AspP}$ **PFV** $[_{vP}$ John go to the store $]]$
$= \lambda w.\lambda t.\ \exists(\mathrm{e}_1)\ [\tau(\mathrm{e}_1) \subseteq t\ \&$ **John-go-to-the-store**$'(\mathrm{e}_1)(w) = 1]$

c. $[_{TP}$ **PRES** $[_{AspP}$ **PFV** $[_{vP}$ John go to the store $]]]$
$= \lambda w.\ [t = t_u\ \&\ \exists(\mathrm{e}_1)\ [\tau(\mathrm{e}_1) \subseteq t\ \&$ **John-go-to-the-store**$'(\mathrm{e}_1)(w) = 1]]$

---

build this condition into the object language, but the other way would not change the diagnosis of the EC greatly.

d. $[_{Mod}$ **MUST** $[_{TP}$ **PRES** $[_{AspP}$ **PFV** $[_{vP}$ **John go to the store′**]]]]
$= \lambda w.\ \forall w' \in \text{BEST}_{g(e_2)}(\bigcap f(e_2))$: $[t = t_u\ \&\ \exists(\text{e}_1)\ [\tau(\text{e}_1) \subseteq t\ \&$ **John-go-to-the-store′**$(\text{e}_1)(w') = 1]\ ]$

e. $[_{CP}$ **ASSERT**$_{e2}$ $[_{Mod}$ **MUST** $[_{TP}$ **PRES** $[_{AspP}$ **PFV** $[_{vP}$ **John go to the store′**]]]]]
$= \lambda w.\ [\textsc{assert}'(\text{e}_2,\ w)\ \&\ \forall w' \in \textsc{con}(\text{e}_2)$: $\forall w'' \in \text{BEST}_{g(e_2)}(\bigcap f(e_2))$: $[t = t_u\ \&$
$\exists(\text{e}_1)\ [\tau(\text{e}_1) \subseteq t\ \&$ **John-go-to-the-store′**$(\text{e}_1)(w'') = 1]\ ]\ ]$

f. "In all worlds w' compatible with the content of the assertion e_2 in w, all of the best worlds w'' according to g in f are such that a John-going-to-the-store event, e_1, is included in the utterance time in w''."

The issue is the fact that the proposition which is the argument of the modal has impossible truth conditions.[42] If t_u is an instant and if e_1 is an eventive predicate, there is no way for $\ulcorner \tau(\text{e}_1) \subseteq t_u \urcorner$ to hold. We see why there is no interpretation of (21) where it has the structure in (69).

With this semantic analysis in hand, we might ask why epistemic modals with stative prejacents are acceptable. Recall our problematic sentences (70); they did not have perfective readings. By contrast, habitual interpretations of these sentences were fine. Moreover, if the predicates in these sentences were stative to begin with, they'd be fine as well. (Think: *John is at the store.*)

[42] As I've stressed before, this account is not beholden to the idea of illocutionary operators, though I have no qualms with them. An alternative which would work just as well is common in situation semantics where the clause is related to the utterance situation by means of some kind of anchoring relation. This topmost anchoring relation might not make use of CON, but otherwise the derivation would proceed similarly.

Do stative predicates, in virtue of their on-going readings in matrix contexts, therefore have default *imperfective* readings? This is the natural thing to say, especially given our notional characterization of the imperfective and the truth conditions I've assigned to perfective and imperfective aspect. Such a hypothesis has been prominently defended. Parsons [1990] suggests that stative sentences have a default 'Hold' predicate, which takes an eventuality and a time argument; ⌜Hold(e,t)⌝ means that e holds at time t. Meanwhile, events have a default 'Cul' predicate, which similarly takes eventuality and time arguments; ⌜Cul(e,t)⌝ means that e culminates at t.[43] In our framework, this amounts to saying that states have default imperfective readings (the 'Hold' predicate is one of the key ingredients to Parsons' account of the progressive) and events default perfective readings, since Cul encodes the intuition that the eventuality in question is completed. Keeping with our preferred semantics for imperfective aspect, a stative sentence like (75a) would then have the representation in (75b).

(75) a. John is at the store.

b. $[_{CP}$ **ASSERT**$_{e2}$ $[_{TP}$ **PRES** $[_{AspP}$ **IMPF** $[_{vP}$ **John be at the store′**]]]]
$= \lambda w.$ [ASSERT′(e_2, w) & $\forall w' \in$CON(e_2): [$t = t_u$ & $\exists(e_1)$ [$t \subseteq \tau(e_1)$ & **John-be-at-the-store′**(e_1)(w') = 1]]]

c. "In all worlds w' compatible with the content of the assertion e_2 in w, the utterance time is included in a John-being-at-the-store event, e_1 in w'."

[43] States can't instantiate a Cul-relation because a culminating state would indicate a change of state, making it a dynamic eventuality, and thereby eventive.

That said, let's assume that lexically stative predicates trigger default imperfective aspect.[44] In the modal case, then, we have the following.

(76) a. John must be at the store.

b. [$_{CP}$ **ASSERT**$_{e2}$ [$_{Mod}$ **MUST** [$_{TP}$ **PRES** [$_{AspP}$ **PFV** [$_{vP}$ John be to the store]]]]]

$= \lambda w.$ [ASSERT$'$(e_2, w) & $\forall w' \in$CON(e_2): $\forall w'' \in$BEST$_{g(e_2)}(\bigcap f(e_2))$: [$t = t_u$ & $\exists(e_1)$ [$\tau(e_1) \subseteq t$ & **John-be-at-the-store**$'(e_1)(w'') = 1$]]]

c. "In all worlds w' compatible with the content of the assertion e_2 in w, all of the best worlds w'' according to g in f are such that the utterance time is included in a John-being-at-the-store event, e_1 in w'."

Again, this is unproblematic, given our semantics.

When we consider eventive prejacents with root modals, our next question is why *these* are acceptable. My explanation of the acceptability of the root modals with eventive predicates presupposes that modal sentences are themselves derived statives. It makes sense why this would be so. If we recall that both GEN and FUT produce derived statives, and so have the subinterval property, the contention that a modal operator in root position does so as well seems reasonable. Moreover, if what root modals express has to do with possibility or necessity with respect to some salient ranking of circumstances, then the conclusion seems even more fitting. The holding of such a possibility/ necessity is

[44]It's worth noting that the issue about whether stative predicates trigger default perfective or imperfective aspect is not absolutely settled. Smith 1991 thinks that present tense statives do have perfective aspect; it's just that the initial and final endpoints of the eventuality are unspecified for states. Moreover, she says, the formal characterization of a stative perfective sentences is semantically like the imperfective. At the risk of idealizing somewhat, I find it more perspicuous to give an analysis of statives with imperfective aspect.

intuitively a state as opposed to an event. Let us marshal some evidence for this claim before proceeding.

CHAPTER 4

Extending the Event-Relative Account of Modals

4.1. Event relativity, root modality, and future orientation

Hacquard's event-relative modal semantics (Cf. Hacquard 2006, 2010) allows one to account for what she calls Cinque's Puzzle, the different positions of modal auxiliaries in the hierarchy of functional projections.[1] One of Hacquard's innovations in these works was to reconfigure modal bases and ordering sources to take an event as opposed to a world argument, which in turn constrain the modal's flavor. To repeat points summarized earlier, epistemic modals sit high in the clause (above Tense) and take the speech event as an argument, yielding an epistemic modal base as in (77a).[2] Root modals sit low (below Tense and Aspect), take the vP event as an argument, and yield circumstantial modal bases as in (77b). This permits modals' lexical entry to be uniform as in (77c), despite the difference in height.

(77) a. $\bigcap f_{ep}(e) = \{w'\text{: } w' \text{ is compatible with } \textsc{con}(e)\}$

b. $\bigcap f_{circ}(e) = \{w' : w' \text{ is compatible with the circumstances of } e\}$

c. $[\![\text{must}]\!]^{w,f,g}=\lambda P.\lambda f.\lambda g.\lambda w[\forall w' \in \text{BEST}_{g(e)}(\bigcap f(e)):P(e)(w')=1]$

[1] According to [Cinque, 1999], the relevant projections are ordered $\text{Modals}_{epistemic}$ > Tense > Aspect > Modals_{root}.

[2] CON is a function from a content-bearing event, like an illocutionary act, speech event, or attitude event, to the set of possible worlds characterizing its content. Cf. Pietroski 2000 for discussion, and section 3.6 for additional explanation.

The difficulty comes in balancing the event dependency of the modal with interpretive facts about the modal's temporal *perspective* and *orientation*. This terminology originates with Condoravdi 2002.[3] By way of review, the temporal perspective of a modal is the time at which the worlds in the modal base are calculated. The temporal orientation of a modal is the relation between the temporal perspective and the time of the described event. For root modals, since they have a circumstantial modal base, this would be the time at which the relevant facts hold which help determine the set of worlds against which the modal is evaluated. I'll focus on a single example that exemplifies the difficulty I have in mind: present tense root modals with eventive prejacents. In English, such sentences have a present perspective and a future orientation. But facts about the event-relative framework as put forward by Hacquard conspire to produce difficulties in yielding a representation which can capture this interpretation properly. This paper aims to lay out nature of the problem, diagnose its source, and propose a solution favorable to the event-relative framework.

4.1.1. Two problems for the event-relative account

In English, root modal sentences with eventive complements have a present perspective with a future orientation. Take (78a) as an example. I'll implement Kratzer 1998's semantics for **PRES** and **IMPF** as in (78b) and (78c). Ignoring the speech event, this gives one the truth conditions in (78d).[4] By way of context for our example, let's say that

[3]Cf. the discussion in section 3.5. Cf. also [Matthewson, 2012], to which I owe my brief summary, for a concise formulation.

[4]I'll largely ignore the role of the speech event, which is important the event-relative framework primarily for the interpretation of epistemics. Though cf. section 3.6 for deductions which include the speech event.

John's mother is due for a visit, and his refrigerator is empty. His goal is to have food in his fridge before his mother's visit.

(78) a. John must stock his refrigerator.

b. **IMPF**: $\lambda P.\lambda t.\lambda w.\ \exists e[t \subseteq \tau(e)\ \&\ P(e)(w)=1]$

c. **PRES**: $\lambda P.\lambda w.\ \exists e[t = t_u\ \&\ P(e)(w) = 1]$

d. $[_{TP}$ **PRES** $[_{AspP}$ **IMPF** $[_{Mod}$ **MUST** $[_{vP}$ John stock his refrigerator $]]]]$
$= \exists e[t \subseteq \tau(e)\ \&\ t = t_u\ \&\ \forall w' \in \text{BEST}_{g(e)}(\bigcap f(e))$:**John-stock-his-fridge**$'(e)(w')$
$= 1]$

The analysis in (78) simplifies Hacquard [2010]'s semantics in two ways. First, it glosses over the aspect movement which guarantees modals' arguments are of a uniform type. Second, it adopts a standard semantics for **IMPF**; Hacquard [2010] adds an extra layer of modality to the imperfective. Neither simplification affects the verdict of this section, but I'll return to the second point later. Notice, for now, that e in (78d) is both (i) the source of the modal parameters for *must*, and (ii) an event of John stocking his fridge in the teleologically ideal worlds picked out by BEST. These features conspire to yield two problems.

4.1.1.1. The Event Identification Problem (EIP). To get a purchase on the first problem, recall that f and g project from the vP's event argument. With Aspect sitting above the root modal, existential quantification of the vP's event argument takes wide scope over the modal. This commits us to an eventuality which exists in the actual world (or in the *generic worlds*, on the semantics of **IMPF** advocated by Hacquard), but which is a fridge-stocking event in the ideal worlds. Intuitively, actual-world e is a

state consisting of John's circumstances ("CIRCUMSTANCE$_e$"), which are held fixed in the worlds delivered by $\bigcap f(e)$ and where the stocking-the-fridge event occurs ("MODAL$_e$").

We can make this idea vivid by drawing on the context to flesh out the circumstances of the scenario a bit. John's mother is due for a visit and his refrigerator is empty. The relevant circumstances picked out by f and g involve John's empty refrigerator and his goals. Let's say his relevant sole goal is to have food in the fridge when his mother arrives. So $\bigcap f(e)$ yields a set of worlds where John's fridge is empty, and $g(e)$ is the set of John's goals. $[\![(78a)]\!]$ is true, relative to this f and g, just in case every world in $\bigcap f(e)$ in which John has food in his fridge by the time of his mother's visit is one in which he stocks his fridge. The problem is, we can't identify CIRCUMSTANCE$_e$ with MODAL$_e$, because any fridge-stocking event is *eo ipso* not an empty-fridge state. At best, we can perhaps say CIRCUMSTANCE$_e$ is a state which partially overlaps MODAL$_e$, but this is, properly speaking, a different event. Yet in (78d), both of these eventualities, an empty-fridge state and a fridge-stocking event, are implausibly picked out by e.

4.1.1.2. the Orientation Problem (OP). The second problem has to do with the temporal orientation of the complement clause in (78d). In (78d), **IMPF** introduces a reference time, and says that this reference time is included in the run-time of the event; the modal event of John-stocking-his-fridge (what I called MODAL$_e$ above). The reference time is identified with the utterance time via **PRES**. By the transitivity of identity, $\tau(e)$ is included in the utterance time, albeit where this time occurs in another world. This is the wrong prediction. The fridge-stocking event should be future oriented with respect to the circumstances. We want MODAL$_e$ to follow CIRCUMSTANCE$_e$ in the ideal worlds,

not be contemporaneous with it. (78d) fails capture the future orientation of the modal's complement.

An event-relative modal semantics is in need of some kind of response to the EIP and the OP. It might seem that the OP would be easy enough to deal with simply by adopting an off-the-shelf mechanism for securing future orientation. Part of the aim of this paper is to show that the matter is not so simple, because the EIP constrains the kinds of solutions one could hope to give to the OP. The next two sections describe some of these issues.

4.1.2. Perspectives on the source of future orientation in modals

There tend to be three types of proposals about the source of future orientation in modals. One position, advanced by Enç [1996] and Werner [2006], suggests that circumstantial modals have a future orientation as part of their lexical meaning. This might be due to the modal's meaning including a mechanism that explicitly extends the evaluation time into the future (as in the case of Enç 1996), or to some more general mechanism. In the case of Werner 2006, the structure of branching worlds combined with modal reasoning will ensure a future orientation for circumstantial modals more generally. A second kind of proposal is similar in that it ties future orientation to the lexical meaning of the modals. But this account, the classic example of which is Condoravdi 2002, also provides a mechanism whereby the Aktionsart of the embedded predicate can impact the orientation. This makes the orientation rely on the interaction of the lexical meaning of the modal and the Aktionsart of the complement with which it composes.[5] On Condoravdi's account, the temporal perspective is either given by tense (where present tense gives a

[5]For discussion of these points, cf. section 3.5.

present orientation), or by the interaction with a perfect operator PERF (where MOD > PERF gives a past orientation). The lexical entries are given in (79).[6]

(79) a. Possibility modal: $\lambda P.\lambda w.\lambda t.\ \exists w'[w' \in \mathrm{MB}(w,\ t)\ \&\ \mathrm{AT}([t,\infty),\ w',\ P)]$

b. Necessity modal: $\lambda P.\lambda w.\lambda t.\ \forall w'[w' \in \mathrm{MB}(w,\ t) \rightarrow \mathrm{AT}([t,\ \infty),\ w',\ P)]$

c. $$\mathrm{AT}(t,\ w,\ P) = \begin{cases} \exists e\ [P(w,\ e)\ \&\ \tau(e,w) \subseteq t] & \text{if } P \text{ is eventive} \\ \exists e\ [P(w,\ e)\ \&\ \tau(e,w) \circ t] & \text{if } P \text{ is stative} \\ P(w)(t) & \text{if } P \text{ is temporal} \end{cases}$$

The contribution to temporal orientation made by modals is due to the AT-predicate; this predicate specifies the interval at which the prejacent is evaluated. A possible future orientation is ensured by $[t,\infty)$, which describes an interval which extends infinitely into the future from t. The precise nature of the AT-predicate differs for eventive and stative predicates, as indicated in (79c). The run-time of eventive predicates is required to be included in this interval. To be included in this interval, the latest such an eventuality could start would be the start of the interval, t. Since events take time to occur, this constraint effectively gives eventives an obligatory future orientation. For stative predicates, by contrast, the run-time of the eventuality denoted by the predicate need only overlap with the interval $[t,\infty)$. In a nutshell, this difference provides for the fact that modals with complements with stative predicates can sometimes have a present orientation, but that modals with eventive complements are always future oriented.

A third approach, exemplified by Matthewson 2012 and Kratzer 2010 (cited in Matthewson), separates the element responsible for the future orientation from the modal itself.

[6]To be clear, Condoravdi is explicitly concerned with non-root modals. But her "metaphysical" modals are similar to circumstantial modals, and nothing in principle prevents us from considering the efficacy of extending her proposal to future orientation in root modals.

On this approach, the modal itself does not provide the future orientation, but the modal will be interpreted as future oriented when it occurs with some other element that triggers a future orientation. Both Matthewson and Kratzer suggest viewpoint aspect as the element thus responsible. Matthewson supports this claim with evidence from Gitskan, where the aspectual morpheme *dim* is both necessary and sufficient for securing future orientation. Her lexical entry for *dim* is below.

(80) a. $[\![\text{ASP}]\!] = \lambda P.\lambda t.\lambda w.\ \exists \text{e}\ [P(\text{e})(w)\ \&\ \tau(\text{e}) = t]$

b. $[\![dim]\!] = \lambda P.\lambda t.\lambda w.\ \exists t'\ [t \prec t'\ \&\ P(t')(w) = 1]$

Matthewson separates the existential closure over events into an ASP head, which acts primarily as a type-shifter, so that various aspectual operators like *dim* may stack above it. ASP existentially binds the event variable and introduces a reference time, which *dim* shifts into the future. As Matthewson points out, *dim* occurs obligatorily with circumstantial modals, intervening between the modal and the verb, so we'd expect present circumstantial modals with eventive predicates to have precisely same kind of interpretation in Gitskan as our main example (78a) in English. I think this expectation is largely borne out, except for the fact that *dim* occurs obligatorily, while there are present oriented circumstantial modals in English.[7]

In Matthewson's account, *dim* amounts to a prospective aspect marker, which echoes Kratzer's contention about the source of future orientation. The difference is that while such a marker is unpronounced in English, it is pronounced in Gitskan. The example

[7]Putting a finer point on the difference between Gitskan and English, in Gitskan, the Aktionsart of the verb has no effect on the temporal orientation, which is uniformly future oriented. In English, the Aktionsart makes a difference; stative predicates can have a present orientation for some modals. This fact is reflected in Condoravdi's analysis. Also, *dim* can occur *above* the circumstantial modal, allowing for a future temporal perspective in addition to a future orientation.

Kratzer uses to illustrate this is the ability modal *can*. I give Kratzer's lexical entries for both the modal *can* and the null prospective morpheme in (81).

(81) a. $[\![\textit{can}]\!] = \lambda R.\ \lambda x.\lambda t.\ \exists x'\ \exists t'\ [<x', t'> \in f(<x, t>)\ \&\ R(x')(t')\ ]$

b. $[\![$ [prospective] $]\!] = \lambda P.\ \lambda t.\exists e[P(e)\ \&\ e \preceq \text{future}_t]$

That Kratzer uses *can* as an example to illustrate the approach proves to be significant. Like Hacquard, Kratzer advocates conceiving of the arguments to the modal parameters as "anchors", which figure in the determining the value of the modal base. (Cf. Kratzer 2013) In (81a), the modal is anchored to a time-slice of an individual. So f takes an ordered pair of an individual and a time, $<x,t>$, and gives the set of modal alternatives to $<x,t>$, which are the other-worldly counterparts of x at t. This is already a departure from the event-relative framework because the modal anchor is no longer the vP's event as in (77b).

In separating out the element responsible for a modal's future orientation, the Matthewson/ Kratzer approach offers the most obvious way of augmenting the event-relative semantics envisioned in (78) so that it may deal with the OP. But there are a few complications. First of all, the ordering of the functional categories differs – at least between Hacquard's account and Matthewson. Hacquard [2010] relies on the hierarchy argued for by Cinque [1999], which has the following ordering: $\text{Modals}_{epistemic} > \text{Tense} > \text{Aspect} > \text{Modals}_{root}$. For Matthewson, both ASP and the additional aspectual operators she posits occur under (circumstantial) modals. With the event variable existentially closed before

composing with the modal, the modal parameters wouldn't have access to the event variable to anchor the modal.[8] Kratzer's account assumes the same ordering as Hacquard's but, as pointed out above, Kratzer's proposal for *can* doesn't have an event anchor.

Still, Kratzer's proposal for [prospective] gives us a clue as to how we might apply prospective aspect to (78), which I do in (82).[9]

(82) a. John must stock his refrigerator.

b. $[_{TP}$ **PRES** $[_{AspP}$ [prospective] $[_{Mod}$ **MUST** $[_{vP}$ John stock his refrigerator]]]]

c. [[82b]] = $\exists e$[e $\preceq$ future$_t$ & $t = t_u$ & $\forall w' \in \text{BEST}_{g(e)}(\bigcap f(e))$:**John-stock-his-fridge**$'(e)(w')$=1]

It might seem that adopting the Matthewson/ Kratzer account of future orientation, with Kratzer's [prospective] more specifically, solves the OP. However, this isn't so. In (82c), MODAL$_e$ is no longer erroneously predicted to be contemporaneous with the utterance time, and is now properly predicted to be future oriented. This is an improvement. However, closer inspection of (82c) shows that new problems now arise. The source of the new problem is that e is now included in a future t, and yet e is the argument to the modal's parameters. In short, the modal is now anchored to an event included in a future

[8]Not necessarily perhaps. Champollion [2015]'s quantificational event semantics allows the event variable to take lowest scope with respect to the quantifier in quantified NPs. On this semantics, verbs denote existential quantifiers over events, so verbal meaning directly introduces an existentially closed event variable to the semantic composition. The closed event description contains a second order variable f which ranges over event predicates, allowing the event variable to be modified even when it is closed. This allows for the possibility that the event variable could be "available" for modification even when it is closed. So, perhaps it could be "available" to the modal parameter to anchor the modal. But I put aside this possibility for now.

[9]Prospective is an aspectual operator, and would be added to the inventory of aspectual operators the that includes **PFV** and **IMPF**. Rather than indicating that the event time is either included in or includes the reference time, this prospective operator says of the event time succeeds the reference time. The way (82) employs this operator, it would be an alternative to either **PFV** or **IMPF**. Interestingly, though, the way Matthewson analyses *dim*, it can stack with other aspectual operators. So, it would be possible, at least logically, for *dim* to co-occur with **PFV**, for example.

time. Putting this in terms of the terminology used in section 4.1.1, MODAL$_e$ is now in the future, but CIRCUMSTANCE$_e$ is thereby shifted to the future as well. (82c) gives us a future perspective when we wanted a present perspective. While adopting this proposal does allow us to secure future orientation, it does so at the expense of the modal's present perspective.

Kratzer's use of a different modal anchor for *can* is precisely what allows her semantics for *can* to avoid this perspective problem. It allows one to set the temporal perspective of the modal directly via the t input to the modal base. But, as indicated above, to follow Kratzer in adopting this modal anchor would amount to giving up the event-relative framework, and the aim of this paper is to see whether we could have an *event*-relative semantics which gets the perspective and orientation facts right.

4.1.3. Revisiting Hacquard [2006]'s account

As I mentioned in section 4.1.1, the event-relative semantics for modals I presented in (78) was not entirely faithful to Hacquard's account. There were two additional innovations that (78) left out. First, Hacquard's semantics relativizes the events that aspect quantifies over to worlds. Second, she contends that imperfective meaning can involve an number of different aspectual operators which quantify over a verb's event argument; GEN, the progressive, and the counterfactual modal CF, which is construed as a universal modal with a metaphysical modal base and a future operator, making it look like a metaphysical modal in the sense of Condoravdi [2002].[10] Since viewpoint aspect relativizes the event quantified over to a world, the added layer of modality above aspect can shift its world

[10]In French, this is usually expressed by *conditionnel* morphology.

parameter. This is the source of Hacquard's account of actuality entailments. When *no* additional layer of modality sits above aspect, then the existential quantification of the event variable takes wide scope over the modal, and we have a configuration where the event occurs in the world of evaluation. Above, I glossed this such an event as CIRCUMSTANCE$_e$ and distinguished it from MODAL$_e$. But Hacquard 2009 proposes a much stronger principle for identifying events across different possible worlds, the PED.

(83) **Preservation of Event Description (PED)**: for all worlds w_1, w_2, if e_1 occurs in w_1 and in w_2, and e_1 is a P-event in w_1, then *ceteris paribus*, e_1 is a P-event in w_2 as well.

Perfective aspect involves no such added layer of modality. An application of the PED would then entail that the event also occurs in the actual world. (In fact, an application of the PED to the analysis in (78d) would entail that the fridge-stocking event occurs in the actual world as well.) Due to the additional layer of modality in the imperfective, the world the event is located in doesn't have to be the actual one. When, for example, the imperfective involves a GEN operator, the event is said to populate all "generic" worlds, which need not include the actual world.

The account of actuality entailments doesn't need to concern us farther, but the added layer of modality provides the opportunity for an additional element to secure the future orientation and address the OP. Though the kinds of examples offered by (78a) – present tense root modals with eventive complements– is not discussed in Hacquard 2010, Hacquard does discuss root modal sentences with eventive complements under PRES in her 2006 (cf. pp. 109 - 111). Upon surveying the inventory of aspectual operators that might be part of imperfective meaning in this case, she settles on the CF operator by process of

elimination. The progressive is out because it cannot scope over the modal, and GEN is dispreferred because because it doesn't accurately capture the natural reading of (78a), leaving only CF. So a more accurate representation of Hacquard's account would render the analysis of (78a) more like (84). For perspicuity, I will not complete the derivation past the CF modal $\Box_{meta}$. We just need to keep in mind that the time introduced by *fut* will be identified with t_u, as per **PRES**.

(84) a. John must stock his fridge.

b. [$_{TP}$ **PRES** [$\Box_{meta}$ [*fut* [$_{AspP}$ **IMPF** [$_{Mod}$ **MUST** [John stock his fridge]]]]]

c. [[84b]] = $\lambda w_3.\lambda t_2.$ $\forall w_2 \in$ META(w_3): $\exists t_1[t_1 \subseteq [t_2,\infty)$ & $\exists e_1[e_1$ in w_2 & $t_1 \subseteq \tau(e_1)$ & $\forall w_1 \in$ BEST$_{g(e1)}$ $\bigcap f(e_1)$: **John-stock-his-fridge**$'(e_1)(w_1) = 1]]]$

This seems to deal with the OP effectively, but not the EIP. To see why, let's take a systematic look at a few features of (84c); to start with, the underlined portion of (85).

(85) $\lambda w_3.\lambda t_2.$ $\forall w_2 \in$ META(w_3):$\underline{\exists t_1[t_1 \subseteq [t_2,\infty)}$& $\exists e_1[e_1$ in w_2 & $\underline{t_1 \subseteq \tau(e_1)}$ & $\forall w_1 \in$ BEST$_{g(e1)}$ $\bigcap f(e_1)$: **John-stock-his-fridge**$'(e_1)(w_1) = 1]]]$

The run-time of the e_1 now includes a time t_1, which itself is included in an interval that extends into the future (remember, t_2 will be bound by **PRES** and identified with t_u). We still have the event quantification scoping over the modal (which we have of necessity, since Asp $>$ Modal$_{root}$), but now the event that aspect quantifies over is relativized to a world, w_2, which is quantified over by the added layer of modality contributed by CF. Cf. the underlined portion in (86).

(86) $\lambda w_3.\lambda t_2.\underline{\forall w_2 \in \text{META}(w_3)}:\exists t_1[t_1 \subseteq [t_2,\infty)$ & $\underline{\exists e_1[e_1 \text{ in } w_2}$ & $t_1 \subseteq \tau(e_1)$ & $\forall w_1 \in$ BEST$_{g(e1)}$ $\bigcap f(e_1)$: **John-stock-his-fridge**$'(e_1)(w_1) = 1]]]$

e_1 may not occur in the actual world, but this difference is insufficient to solve the EIP. In fact, the combination of the analysis in (84c) and the PED presents us with a kind of dilemma. In order for $\bigcap f(e_1)$ to pick out the relevant circumstances against which the modal is evaluated, e_1 needs to be a circumstance characterized by John's empty fridge. In fact, the CF modal may give us this, since it will hold fast the circumstances of the world of evaluation up through t_2. But, since e_1 is an empty-fridge state in all metaphysical alternatives w_2, it is also an empty fridge state in all ideal worlds w_1. Then e_1 is *both* an empty fridge state and a fridge-stocking event. Alternately, e_1 is a fridge-stocking event in all ideal worlds w_1, so by the PED, e_1 is also a fridge-stocking world in all metaphysical alternatives w_2. But then e_1 in w_2 cannot characterize John's circumstances in w_3 where he has an empty fridge. Neither option is satisfactory.

4.1.4. An event-relative solution to both problems

The EIP seems to stymie otherwise plausible solutions to the OP. Part of the issue is that if we want the modal to have a present perspective but a future orientation, then it simply couldn't be that $\textsc{circumstance}_e = \textsc{modal}_e$. Moreover, as long as the argument for the modal parameters is the same e as that which is introduced by the verb, we risk running into the EIP. But taking a vP event as modal anchor was a central claim of the event-relative semantics for modals, and allowed the account to deal with Cinque's Puzzle. Can we maintain the event-relative framework, given these difficulties? I think so, provided we reify what I've been calling $\textsc{circumstance}_e$. Upon doing that, we can take a leaf from the Matthewson/Kratzer account of future orientation to secure the proper temporal orientation. Before detailing my proposed solution to both the OP and the

EIP along these lines, I will first present some independent evidence which suggests that CIRCUMSTANCE$_e$ can indeed have its own event variable.[11]

4.1.5. Reifying the circumstance state

The claim that root modals introduce a second event variable echoes a claim made by Homer [2011] that root modals are stative predicates of eventualities. Root modals then anchor to these stative eventualities ("the evaluation points of their accessibility relations"), which in turn are quantified over by viewpoint aspect. If the circumstance state had its own event variable, then we would expect this eventuality to have its own time and space coordinates independent of the eventuality that denotes the modal's complement. Homer provides examples from French indicating that this expectation is in fact met. In (87a), the temporal adverbial *hier* applies to the time of Pierre's obligation (the CIRCUMSTANCE$_e$), while *la semaine prochaine* applies to the run-time of an event of Pierre turning in his homework. In (87b), the spatial coordinate of the legal situation fixed by the adverbial *dan ce pays* contrasts with the adverbial which modifies the event of Jean getting his surgeon's degree, *a l'etranger*.

(87) a. Context: The rules have just changed: Pierre now has to turn in his homework tomorrow.

[11] Ramchand [2014, 2018]'s recent modal semantics also proposes a model along Homeric lines, with a kind of double situational semantics. Ramchand's account is similar to the present proposal in that it is not the vP event that anchors circumstantial modals, but a situation introduced above the vP. It is also similar in that it effects a future orientation without positing a prospective aspect head. But it also represents a more radical departure from the event-relative framework I aim to discuss here. Still, a comparison of the present proposal and Ramchand's semantics is worth undertaking, and I leave it to future work.

Hier encore, il devait$_{deon}$ rendre son devoir la semaine
yesterday still he must-PST turn in his homework the
prochaine.
next week

'Yesterday, he still had to turn in his homework next week.'

b. Context: Where he lives now, is Jean allowed to practice as a surgeon with his French degree?

Non, dans ce pays Jean ne peut$_{deon}$ pas avoir obtenu son diplme de
no in this country Jean NEG can-PRS NEG have gotten his degree of
chirurgien l'anger.
surgeon abroad

'No, in this country, Jean is not allowed to have gotten his surgeon degree abroad.'

Similar evidence can be found in English. They are a bit harder to get because English modals tend not to inflect for tense, making the contrast harder to draw out. But we can use the periphrastic modal *have to* to make the contrast vivid.[12]

(88) Context: On June 1st, John's fridge is empty. He knows his mother is coming on the 16th, so he marks on his calendar to go to the store on the 15th. In the meantime, he subsists on takeout. On the 14th, John's roommates decide to do him a favor and do his shopping for him to stock up his fridge for his mother's visit. After running John's errands, one of his roommates utter:

For two weeks, John had to go to the store tomorrow (but we wound up going for him).

[12]This example is structurally similar to example (44) considered earlier.

In this example, the lower adverbial modifies the run-time of the event described by the verb. The frame adverbial modifies the time of the obligation. Intuitively, (88) means that John had an obligation to go to the store at a certain time (where *tomorrow*= June 15th), and that this obligation lasted (at least) two weeks.[13] On the basis of the French data, Homer provides the following analysis of the root interpetation of *pouvoir*.

(89) a. $[\![pouvoir_{root}]\!]^{c,s} = \lambda\Phi.\lambda w.\lambda e.\ e \text{ in } w \ \&\ \exists w' \in \text{Acc}(e)[\Phi(w')]$

b. Jean peut$_{deon}$ travailler.

c. $[\![(89b)]\!]^{c,s} = 1$ iff at t_c there is an eventuality e in w_c such that for some world w' compatible with e, Jean works in w'.

Note that the eventuality quantified over by aspect is the argument of the modal's parameters (the function Acc), but is not the eventuality introduced by complement's verb. Put in the terms introduced above, the CIRCUMSTANCE$_e$ is no longer the MODAL$_e$.

4.1.5.1. Introducing the coercion operator Ω. In the informal gloss Homer gives on his analysis in 89c, CIRCUMSTANCE$_e$ is not temporally ordered with respect to MODAL$_e$, so we don't yet get the sense that the modal has a future orientation.[14] But we can combine Homer's reification of the circumstance state with the Matthewson/ Kratzer account of future orientation. To accomplish this, I propose a coercion operator, Ω, which introduces an eventuality, and relates this eventuality to the complement's eventuality.[15] As a first pass, consider 90a, and the resulting derivation in (90c).

[13]This is the same kind of adverbial evidence given for the syntactic reality of a dedicated projection for futurates. Copley [2009] offers this as evidence for the syntactic reality of a *plan* which is referred to by futurate statements. In Copley [2014], she amends the proposal slightly, such that futurates introduce a stative eventuality of their own.

[14]To be fair, temporal orientation is not the phenomenon Homer's paper is concerned with.

[15]For now, I just propose this coercion operator, observing that it solves the problem for us. I will later provide an independent motivation for this operator by considering what might trigger the coercion.

(90) a. $[\![\ [_{vP2}\Omega]\ ]\!]=\lambda e_2.\lambda P.\exists e_1[P(e_1)\ \&\ R(e_2,e_1)]]$

b. $[_{TP}$ **PRES** $[_{AspP}$ **IMPF** $[_{Mod}$ **MUST** $[_{vP2}\ \Omega\ [_{vP1}$ John go to the store]]]]]

c. [[90b]] = $\exists e_1\ [t \subseteq \tau(\mathrm{e}_2)\ \&\ t = t_u\ \&\ \forall w' \in \mathrm{BEST}_g(e_2)(\bigcap f(e_2))$: [$\exists(e_1)$[**John-stock-his-fridge**$'(e_1)\ \&\ R(e_2,e_1)](w') = 1$]

A few words on the proposal. The new eventuality introduced by Ω, e_2, is a stative eventuality, understood the way Homer suggests. Intuitively, it's the state characterizing the circumstances according to which the modal is evaluated. Ω existentially closes the event variable of the embedded verb, but it does not relate this event directly to a time, so it does not perform the function normally associated with viewpoint aspect. Moreover, it is not a function from a property of events to a property of times, which is why I've suggested that it is introduced in a vP-shell projection as opposed to an aspectual operator in AspP. The eventuality Ω introduces is in turn quantified over by aspect. However, we can use the predicate R to temporally order the eventualities e_1 and e_2 to get the orientation right.

So far, the proposal in 90a doesn't give us future orientation–we just have a dummy predicate R relating e_2 to e_1. A simple precedence relation like $\prec$ would be sufficient to give us future orientation, and identifying R with $\prec$ would be the most straight-forward way of implementing the Matthewson/ Kratzer insight about future orientation. However, I won't be advocating identifying R with $\prec$. Though I won't defend the claim at length here, I think $\prec$ is too weak a relation to relate e_2 to e_1. It doesn't capture the intuition that in sentences like 78a, $\textsc{circumstance}_e$ *calls for* $\textsc{modal}_e$. The reason John must go to the store is because his of his empty fridge and his mother's impending visit. It would be preferable, if possible, for the relation to encode the intuition that in order to

comply with demand on him 78a expresses, John's stocking his fridge must be the means by which he responds to his circumstances.[16]

Luckily, another idea is ready to hand which allows us to capture this intuition. In Copley 2014's causal chain analysis of futurates, she employs a coercion operator very similar to Ω, but for futurates. In Copley's semantics for futurates, R is simply the predicate CAUSE. Insofar as e_2's causing e_1 implies that e_2 precedes e_1, it gets the facts about temporal orientation right, and it also encodes the intuition that e_1 is brought about by the relevant circumstances. Adopting this proposal, we can update the first pass at Ω with (91b).

(91) a. $[[\ [_{vP2}\Omega]\ ]] = \lambda e_2.\lambda P.\exists e_1[P(e_1)\ \&\ \text{CAUSE}(e_2,e_1)]]$

b. [[90b]] = $\exists e_2\ [t \subseteq \tau(\text{e}_2)\ \&\ t = t_u\ \&\ \forall w' \in \text{BEST}_g(e_2)(\bigcap f(e_2))$: $[\exists(e_1)[$**John-stock-his-fridge**$'(e_1)\ \&\ \text{CAUSE}(e_2,e_1)](w') = 1]$

The way this addresses both the OP and EIP should now be fairly obvious. According to the analysis in 91b, 78a is predicted to be future oriented, due to the fact that MODAL_e is subsequent to CIRCUMSTANCE_e because $\text{CAUSE}(e_2,e_1)$. Yet, CIRCUMSTANCE_e is quantified over by aspect, and the temporal trace of the newly introduced event variable is bound by tense. So we have future orientation without sacrificing present perspective, and thereby a genuine solution to the OP. As for the EIP, there is no expectation that e_1 is the same event as e_2, whether identified across possible worlds or not. However, the relation

[16]The intuition recalls a debate pursued by a number of philosophers and deontic logicians who describe the difference between so-called "deliberative" and "evaluative" *oughts*. Cf. Harman 1973, Geach 1982 and more recently, Schroeder 2011, Chrisman 2015. I will pursue this intuition further in the following Chapter.

between e_1 and e_2 makes it appropriate to see one as the circumstances relative to which the necessity of the other is evaluated.

Finally, I will address a loose end for this proposal. Upon introducing Ω, I said that it was a coercion operator, but I neglected to say what triggers the operator, or what grammatical circumstances call for the coercion. Homer's account, it is worth noting, does not involve coercion; he claims that root modals always introduce their own event variable. In construing Ω as an operator introduced by coercion, I don't assume that root modals *always* introduce a new event argument. Part of the reason for this is that not all English root modals are future oriented; their perspective does not always differ from their orientation. Complements with stative predicates, like those in 92, can have present perspective and present orientation, and do not give rise to OP or EIP.

(92) a. There should be world peace.

b. Milton, you should be alive at this hour! (paraphrased from Wordsworth)

Instead, what seems to trigger the future orientation, and the ensuing need for Ω, is a complement with an eventive predicate, echoing the kind of Aktionsart effect pointed out in Condoravdi 2002. The coercion proposal then integrates the kind of aspectual sensitivity that Condoravdi builds directly into the lexical meaning of modals. My idea, instead, is that root modals themselves do not affect the underlying aktionsart of the eventuality they anchor to. So, *were* the modal to project from the complement's event variable, as in (78d), the eventuality thus described would still be eventive (since it is a fridge-stocking event). However, as pointed out by Bohnemeyer and Swift [2004], eventive predicates trigger default perfective readings in English. Eventive prejacents without Ω

would therefore trigger **PFV**, instead of **IMPF**. Of course, the **PFV** + **PRES** configuration is ruled out, since an event cannot be included in the utterance time. Ω therefore introduces a stative eventuality, and being stative, allows the expression to compose with **IMPF**. In a nutshell, Ω is triggered by the eventive predicate in the complement, so that the expression may compose with **IMPF** as opposed to **PFV**.

In conclusion, maintaining an event-relative semantics for modals would require root modals with eventive predicates to introduce a new eventuality. I've proposed just such an account by means of the coercion operator Ω, which maintains the essential ingredients of Hacquard's event-relative semantics. While Chapter 3 sought to explain the Eventivity Constraint by appealing to the event-relative framework for modals, this present section diagnosed and fixed a problem in the event-relative framework that arose in considering deontic modals.

4.2. Can Tense Scope Over Epistemic Modals?: the Case of Hindsight Readings of Epistemic Modals

The explanation of the Eventivity Constraint advanced in chapter 3 trades on the claim that functional projections like epistemic modals, tense, and root modals are ordered in a consistent hierarchy. In doing so, it follows long-standing claims made among others by the cartography movement in syntax, and follows Hacquard in assuming that this ordering is essentially as these syntacticians claim. However, this thesis is not universally accepted, and, in fact, it has been resisted by some philosophers and semanticists. A major reason for this resistance comes from so-called "hindsight" readings of epistemic

modals. Hindsight readings of epistemic modals[17] are typically motivated by the following kind of scenarios. von Fintel and Gillies [2007, 2008] give the following example.

(93) CONTEXT: Sophie is looking for some ice cream and checks the freezer. There is none in there. Asked why she opened the freezer, she replies:

a. There might have been ice cream in the freezer.

b. PAST(MIGHT(ice cream in freezer))

Von Fintel and Gillies claim that a proper interpretation of (93a) renders it as (93b), with the tense scoping over the modal. The need for the past tense to scope over the modal is motivated by the following considerations. It seems reasonable to interpret Sophie's utterance in (93a) as true and reasonable – she is not uttering something obviously false. But at the time when she utters (93a) she is well aware that there is no ice cream in the freezer; she just checked! So if the interpretation of (93a) had been (94) instead, it would mean that Sophie is claiming that it is consistent with her knowledge (now, at the time of utterance) that there was ice cream in the freezer at some time in the recent past, namely just before she checked.

(94) MIGHT(PAST(ice cream in freezer)

This would make (93a) false, because having just checked the freezer, she now knows that there is no ice cream in the fridge and it is not possible that there *was* ice cream in it when she looked. Nonetheless, it doesn't seem inappropriate for Sophie to have uttered (93a), nor does it seem false. And, in order for the utterance of (93a) to have been appropriate *and true*, the modal needs to to access her state of knowledge the moment before she

[17]In this section, I'll often refer to these as "hindsight cases", taking it as understood that I'm referring to a particular kind of interpretation of epistemic modal sentences.

opened the freezer door, when it was consistent with her knowledge then that there be ice cream in the freezer. If modals are assessed with respect to the speaker's information state of the moment of utterance, it doesn't seem that there is a way to capture this.[18] But if past tense can scope over the modal, then it could shift the time back to a past state of mind as needed.[19]

We can draw on Condoravdi [2002]'s terminology (Cf. also Matthewson 2012) to make the issue with these hindsight readings of epistemic modals more precise. Condoravdi distinguishes between the temporal perspective of a modal and its temporal orientation. The temporal perspective is the time at which the modal base is calculated. In less formal terms, and for epistemic modals, it's the time at which the information state is assessed.[20] The temporal orientation is the relation between the temporal perspective and the time of the event or state described by the prejacent. In other words, it is the time at which the event described by the prejacent is to hold if the the sentence turns out true. The modal has as a future temporal orientation if the event described by the prejacent is in the future with respect to the temporal perspective. On the semantics I've adopted, Hacquard's event-relative semantics, epistemic modals have a *present* temporal perspective in virtue of being keyed to the speech event. This is because the modal base is a function from this event to the content characterizing the information state at the time of the speech event. However, as von Fintel and Gillies would have it, the temporal

[18]In saying that the relevant information state is the speaker's, I bracket concerns about putative cases of epistemic modals that seem to be interpreted relative to information states characterizing distributed knowledge, or about uses of epistemic modals to describe another person's information state. What interests us here is the temporal behavior of the information state.

[19]Silk [2016] gives similar examples to motivate the idea that tense scopes over the modal in these hindsight cases. For examples such as these, cf. also Egan et al. 2005 and Dowell 2011.

[20]Insofar as circumstantial modals are interpreted relative to a set of circumstances, the temporal perspective for circumstantial modals is the time at which the relevant circumstances hold.

perspective of hindsight readings of epistemic modals is not present – it's past, and a plausible reason for it having a past perspective is that the past tense scopes over the modal, shifting the perspective from the present to the past.

There's a potential challenge to my analysis and explanation of the EC in the offing if we accept this explanation of these hindsight readings. Since my explanation of the EC trades on a view of clausal architecture that places epistemic modals above tense, and leverages this fact into an explanation of why the EC is in fact a constraint on obtaining an epistemic interpretation of modal sentences, if it turns out that tense can scope below epistemic modals, then this explanation might turn out on shaky ground. On the one hand, there's the mere fact that the scopal properties in the clausal architecture aren't as universal as I've been supposing, and as some in the syntax community would believe. But also, were epistemic modals to scope under tense, it would create problems for Hacquard's account of modals more generally. According to Hacquard, recall, if a modal merges above tense, then the event argument is bound by a speech or attitude event, and the modal base is epistemic. If it merges below tense, then it is bound by the VP event and the accessibility relation is circumstantial. If hindsight readings of modals require (past) tense to scope over the epistemic modal, this spells trouble for Hacquard's account because the relative scope of tense is supposed to provide direct evidence of where the modal merges. Insofar as tense scopes above the modal, this is evidence that the modal merges below tense.[21] And on this framework, this would suggest that it is bound by the VP event. But if the modal is bound by the VP event, the resulting modal base would be circumstantial. This is troubling, because the hindsight cases we are considering are clearly epistemic. So,

[21]For example, the fact that tense can shift the temporal perspective of root modals is typically taken as evidence that root modals scope under tense.

if tense scopes over the epistemic modal, and we hold fast Hacquard's account of modals, we have trouble accommodating some pretty clear facts we are attempting to explain.[22]

Hacquard [2006, pp. 158–159] has her own response to the claim that tense scopes over the modal, suggesting that hindsight readings of modals are actually bound by an elided *attitude* event. So, the semantics for these kinds of hindsight cases does not involve the scoping of past tense over the epistemic modal, but is rather like where epistemic modals are embedded under attitude verbs. The idea is that we can key the notion of epistemic possibility to a point in the past *without* the tense scoping over the modal yet *still* shifting back its point of evaluation by assuming that the elided embedding attitude verb is itself shifted back by past tense. This allows for the circumstances where the possibility no longer holds now with respect to the speaker's information state at the time of utterance, but it did with respect to the internal now of the past attitude. Hacquard gives the following examples to illustrate.

(95) a. (~~I thought that~~) my keys might have been in there.

b. We bought a ranch which (~~we thought~~) might have contained an oil reserve.

In (95) the embedding attitude is elided, though still present in logical form. And because it is still present in logical form, sequence of tense rules allow for the deletion of the embedded tense. So, the complement under the modal can be interpreted relative to the present time *with respect to the time of the the attitude* (or the internal "now" of the

[22]There is another possibility, which I will mention, but not pursue. It could be that somehow the epistemic modal is keyed to the speech event but that tense still scopes over it. This would allow Hacquard's semantics to sit comfortably with von Fintel and Gillies's proposal, but it feels a bit like having one's cake and eating it, too. I'm not sure how to motivate this possibility myself.

attitude), as opposed to a back-shifted time with respect to the utterance time. Such an explanation would explain von Fintel and Gillies's (93), repeated below as (96) as follows.

(96) CONTEXT: Sophie is looking for some ice cream and checks the freezer. There is none in there. Asked why she opened the freezer, she replies:

a. There might have been ice cream in the freezer.

b. (~~I thought that~~) there might have been ice cream in the freezer.

The apparent past perspective of (96a) is, again, the temporal present perspective of the elided past attitude represented by the elided *I thought that*. Just like von Fintel and Gillies wish, this analysis allows it to be the case that Sophie says something true in uttering (96a/96b), about a possibility which nonetheless no longer holds at the time of the utterance.

My proposal for how to deal with hindsight cases without needing tense to scope over the modal comes in two parts. The first part of the proposal is that in hindsight cases, epistemic modals are interpreted relative to a speech event that departs from the actual utterance event. So, rather than the highest event denoting the *actual speech event*, the event variable denotes a past event which serves as a *hypothetical* speech event. Comparing this to Hacquard's proposal will bring out some features of what I am suggesting. On Hacquard's proposal, the elided attitude event will indicate a past attitude event(/ state) at which the relevant information state the holds. The modal is then keyed to the past attitude event, allowing it determine the properties of the modal base worlds. My proposal suggests that we key the modal to such an event, though not as an elided

attitude event, but as a hypothetical speech event.[23] There is then no elision, and the epistemic modal would compose similarly in the hindsight cases as in the normal cases, as in (97). (The material after the existential closure of the modal is elided because we haven't yet discussed what operators are in the scope of the modal on this proposal.)

(97) $\lambda w.$ [ASSERT$'$(e_2, w) & $\forall w' \in$CON(e_2): $\forall w'' \in \text{BEST}_{g(e_2)}(\bigcap f(e_2))$: [[. . .]]],
where e_2 is a proposed hypothetical speech event.

So, e_2 is not the actual speech event, but a hypothetical speech event (hypothetical, because it doesn't require that this speech event actually occurred). But because the modal base recovers the relevant information (say, that of the speaker), at the time of the speech event, it doesn't much matter that no actual speech event occurs. It is sufficient that there *was* an information state at the time which the modal can quantify over.

Before proceeding to the second part of the proposal, I want to take a moment to elaborate on what I think justifies this move. The first thing to say in support of this move is that, although in most discourse in natural language, there is a strong presumption that the highest event would be the actual speech event, and there are many conventions to support this presumption – that indexicals and and tense is interpreted relative to its parameters, for example. But under certain circumstances, this presumption is in fact

[23] "Hypothetical speech event" is perhaps an unfortunate name for what I have in mind, but I'm not sure I can think of a better one. To be sure, the event is hypothetical because no actual speech need have occurred. But the kinds of hindsight readings of epistemic modals are noteworthy in part because of their limited distribution. We tend to see them only in action explanation contexts – for example, where an agent retroactively justifies an action based on her information state at the time of action. The event I therefore have in mind is an actual event – one where an agent has acted on the basis of her intentional states. And having intentional content at this time, CON(e) is defined, much as it is for an actual speech event. Yet, I will nonetheless call it a *hypothetical* speech event because we still want this agent to be the hypothetical "speaker" of the context to resolve any indexical pronouns in the right way. Importantly, because the event is an actually occurring one (modulo the speech part), this makes a truth-conditional difference to the interpretation of the sentence – there is an actual information state being queried by the modal, and the modal sentence is true based on the nature of this information state.

lifted, and discourse is interpreted relative to a distal event (with respect to the actual utterance). There are several examples of this: various forms of narrative, free indirect discourse (more on this later), or, to draw on an example that is much discussed in the literature – messages on an answering machine.

Most of the literature on the answering machine "paradox" works within the Kaplanian framework. The reason for this is that, in spite of Kaplan 1977 providing a useful logic and framework for dealing with demonstratives and indexicals, it seems to falter on answering machine cases. It does so in part because of its treatment of sentences like *I am here now* as true at all contexts. However, a typical answering machine greeting will often contain explicit denials of this sentence, as in "I am not here right now", which seems both appropriate, informative, and true, when heard on an answering machine. A common proposal for how to deal with these answering machine sentences on Kaplan's framework is to split the univocal context parameter into multiple parameters, one for the the context of production/ inscription (of the message) and one for the context of audition/ tokening (in other words, a context for when the message is consumed). In most discourse, these contexts are identical – at least, there is no reason to distinguish them. However, in certain cases, they will come apart, as they do in the answering machine cases. The idea is that, *I am not here now* can be true because certain of the indexicals are interpreted relative to the context of production ('I') and others ('here', 'now') are interpreted relative to the context of audition/ tokening.

If we compare the split-context proposal and the event-relative semantics for modals, where the highest event is an illocutionary/ speech event, we could note that utilizing two context parameters in this way doesn't in any necessitate that there are two speech

events that would somehow appear in the logical form of the sentence. Instead, there is one speaking event that took place, namely the speech event that resulted in the recorded message (what on the Kaplan-style proposal would be the context of production), but that the sentence is not interpreted relative to this speech event. Instead, it is interpreted relative to a proposed, *hypothetical* event where the actual speaker and actual place is still considered the speaker and place (and so, serves as the values of the indexicals "I" and "here"), but the time of the utterance is the time at which the message is heard (so *this* time serves as the value of the indexical "now").[24] What this shows is that, if we are already including speech or illocutionary acts in logical form, then certain kinds of discourse necessitate that we countenance non-actual speech events as the value of the speech or illocutionary event for these kinds of discourse. The first part of my proposal for how to deal with hindsight readings of epistemic modals is not *sui generis*, but is already required for these other constructions.[25]

In the hindsight cases under discussion in this section, resolving the hypothetical speech event is actually a much easier affair than in the case of the answering machine cases. This is because, in those hindsight cases, the hypothetical speech event is evoked

[24]This is appropriate for, say, the interpretation of an answering message. If I write a note in my office, but leave it for my spouse at home, then we could easily imagine that the value of the indexical "here" would be the place at which the message is read, not the place at which the message is produced.

[25]It's a difficult question why indexicals behave the way they do in the answering machine example, with some of them being interpreted relative to the context of production and others interpreted relative to the context of audition. Discussion of the metasemantics of answering machine sentences contains a number of proposals; for example, that the very character of the indexicals is modified in certain contexts (Sidelle 1991, Cohen 2013), that the speaker's intentions (Predelli 1998, 2005) determine the extension of indexicals, that an (idealized) audience that determines the extension of the indexicals (Romdenh-Romluc 2006), or that there are metarules governing the valuation indexicals which can predict the answering machine cases (Michaelson 2014). (Cf. also Cohen and Michaelson 2013 for discussion.) The way to understand these various metasemantic arguments in terms of my proposal is that they differ on how the resolution of the hypothetical speech event is to occur. I don't have much to add to this discussion here, accept to point out that, as we'll see, similar problems do not arise in the case of hindsight modals.

by the previous discourse usually through a conversational demand on one of the conversational participants for rationalizing their action. "Asked why she opened the freezer...", the von Fintel and Gillies set up asks, thereby making salient an earlier discourse situation, namely when the discussant was about the open the freezer. The hypothetical speech event is the event that would have occurred had she uttered there. In fact, we simply don't see hindsight interpretations of modals in a discourse that doesn't *first* make the hypothetical speech event salient in this manner.

The first step of the proposal explains the apparent past orientation of the modal – my proposal essentially says that the modal still has a present orientation, but it is present relative to a temporally distal, hypothetical speech event. There's also the second step of my proposal, which is to explain why the value of the prejacent isn't shifted back on account of the past tense. In other words, assuming the modal is keyed to the past speech event as evoked by the discourse, there's still the manner of the supposed past tense in the scope of the modal.

I believe that there are two possible responses to this. One is just to take a leaf from Hacquard and say that the past tense is not interpreted, though not because it is deleted by a sequence of tense rule. If there is no elided attitude verb, as I am supposing, then no sequence of tense rule could be initiated. Rather, we might suspect that the past tense is not interpreted because the utterances in these hindsight cases are actually a kind of free indirect discourse (FID). FID behaves a bit like a mixture of direct and indirect discourse. In an instance of FID, we typically see that tenses and pronouns behave as they would in an attitude report, but all the other indexicals (including temporal indexicals) behave as in direct discourse. As with the answering machine case, it's been said that

FID interprets these two kinds of variables relative to different kinds of contexts. The tense and pronouns are interpreted relative to the actual context of utterance, but the other variables are interpreted relative to (what Schlenker [2004] calls) the context of thought. (Cf. also Doron 1991) This is to say, the tense and pronouns are interpreted as we'd normally expect, but the other indexicals are interpreted relative to the recreated/ imagined context at which the original thought or expression (which is now being related by the FID) was had. The variables interpreted relative to the context of utterance can trigger reference failure or infelicity if the constraints they place on the context of utterance are not met (for example, if the the time of the described action does not occur before the context of utterance, when inflected for past tense), but otherwise is interpreted as though bleached of its tense features and relative to the context of thought. If we consider these generalizations about FID, if the sentences uttered in the hindsight circumstances are instances of FID, the past tense in the hindsight readings of epistemic modals ensures that the described possibility is past, but it otherwise semantically interpreted relative to the context of thought, with the past tense morphology not contributing further to the semantic interpretation.

This story seems plausible to me, but integrating it the Hacquardian framework would require amending the framework to deal with FID more generally, and I won't undertake this work here. However, the proposal would also make some predictions that are a bit more difficult to assess. Importantly, it would mean that, aside from tense and pronouns, other indexicals are interpreted relative to the context of thought. So, if it is an instance of FID, in principle the hindsight readings should include temporal indexicals that are interpreted not relative to the context of utterance but rather to the context in which the

original thought had occurred. So we should be able to see something like (98) in response to the query put to Sophie by her interlocutors in von Fintel and Gillies's scenario, where *now* is interpreted as referring to the actual time at which Sophie had the thought to look in the freezer for the ice cream.

(98) There might have been ice cream in the freezer now.

I find this difficult to assess because, although (98) doesn't sound terrible, it is certainly less natural than without it. It sounds a little bit better if we include an attribution indicating to whom the thought expressed by the FID originates with. In this case, it is with Sophie herself, though at an earlier time.

(99) There might have been ice cream in the freezer now, I thought.

(99) sounds better to me, though the most natural utterance still seems to be the bare (93a). So, while it is plausible that these hindsight readings of epistemic modals are indeed instances of FID, some of the data are more equivocal – for instance, this prediction about the behavior of indexicals in the discourse.

Another possibility, the one I will endorse, is that what we see in these hindsight examples isn't actually a tense, but an aspectual operator, the Perfect. Even when taking the embedded expression to be an aspect and not instance of past tense, this does not dissolve the issue – at least, not without some further remarks, because we'll need to get clear on some properties of the Perfect itself. The English Perfect is usually glossed as indicating that the event is completed before the reference time. For example, Parsons [1990] glosses the Perfect as denoting a kind of state – what he calls the consequent state. This is the state that obtains when an event is completed, as we see in (100).

(100) a. John has eaten an apple

b. ∃e∃x. (eat(e) & Agent(e, John) & Theme(e,x) & apple(x) & HOLD(CS(e), S)).

According to the logical form Parsons assigns (100a), it means that there is an event of eating whose agent is John, whose theme is an apple, and whose consequence state holds at S (where S is understood to be the utterance time). As far as the consequent state goes, Parsons offers the following characterization (Cf. p. 234).

(101) e's consequent state holds at t iff e culminates at some time at or before t.

The issue this raises for the present proposal has to do with Parsons's notion of culmination. This analysis of the Perfect has it that that an event or state described in the Perfect is *completed*; no longer ongoing. But in the hindsight cases we're considering, that would cause trouble. If the ice cream's being in the fridge has *culminated*, wouldn't that mean it was no longer in the fridge by the reference time? If that were the case, then Sophie's utterance would not be true and appropriate, contra the data we are trying to accommodate. If the embedded event were to be *over* by that point, that would mean that the ice cream's being-in-the-freezer is over. On the most straight-forward understanding of a completed event, this would seem to mean that the ice cream is no longer in the freezer. As a consequence, the statement would be saying something like: it's consistent with my (past) information state that the ice-cream's being in the freezer is over. But such an utterance would no longer make any sense – why would Sophie be looking in the freezer for ice cream, because it's consistent with her information state at the time that the ice cream is no longer there?

The way forward is to pay attention to the aspectual properties of the verbs embedded under the Perfect. In the example given cited from Parsons, the embedded verb *eat an apple* is eventive – an accomplishment verb. Parson's analysis of the Perfect as describing a completed verb seems perfectly appropriate for such verbs. But in the hindsight cases we are considering, the embedded predicates are all stative. And if we consider stative predicates, Parsons's characterization of the Perfect seems considerably less appropriate. For example, see the examples in (102).

(102) a. I have lived in Chicago since 2011.

b. Gerald has known Kim for 10 years.

A first thing to note from these examples is that the Present Perfect with embedded lexical statives (like *live* or *know*) tend to require an adverbial describing an interval at which the state holds. As discourse-initial utterances, the sentences in (102) would be decidedly odd without the adverbially specified intervals *since 2011* or *for to years*. However, this is not a categorial restriction – (102) without the specified intervals is fine in contexts in where my living status in Chicago or Gerald's friendship with Kim is already under discussion. This is a welcome observation because in the hindsight sentences we are considering, where, by hypothesis, the present perfect scopes under the epistemic modal, we rarely see an interval adverbial. But of course, the status of whatever is in the freezer is precisely under discussion.

Secondly, the Present Perfect with embedded statives emphatically does *not* mean that the embedded state is completed by the utterance time, the way this seems to be the case with embedded eventives. An utterance of (102a) does not mean that the state of

my living in Chicago is completed – that I no longer live there.[26] An utterance of (102b) does not mean that the state of Gerald's knowing Kim is over.

Since Parsons's proposal on how to analyze the Perfect seems too strong, giving truth conditions that are inappropriate for Perfect sentences with embedded stative predicates, we need a different account of the Perfect. Kratzer [1998] proposes for the following lexical entry for the Perfect.

(103) **PERF**: $\lambda \mathrm{P} \lambda t \lambda w \exists \mathrm{e}[\mathrm{time(e)} \prec t \ \& \ \mathrm{P(e)}]$

This lexical entry for the Perfect seems to capture the difference between stative and the eventive predicates considered earlier better than Parsons's account. It only indicates that the event occurs prior to the utterance time, but this doesn't require that the event or state have culminated prior to the reference time. To be sure, if P is a telic event description like an accomplishment or an achievement, then the only way for the it to satisfy (103) is for the event to be completed by t. This is because, were it not completed, it wouldn't fall under description P by t. On the other hand, because statives are homogenous, there can be an event falling under P prior to t, and yet still be ongoing, because continuations of e can still themselves be P-events.[27] If we plug this version of the Perfect into Hacquard's

[26]Of course, in (102a), we can attribute this to the semantic contribution of *still*. However, in contexts where the Present Perfect is acceptable *without* the temporal adverbial, it's still the case that interpretation where the state is completed is by no means obligatory. It's not contradictory to respond to the question *Have you ever been to Chicago?* with *I've lived in Chicago. In fact I still do.*

[27]This would mean that foregoing remarks also hold for activities, insofar as they are also homogenous in the way statives are. I think this is appropriate. For activity predicates like *eat*, an utterance of *I have eaten* tends to be used to indicate that the eating event has culminated, but this seems more like an implicature than an entailment. *I have eaten. In fact, I still am.* is a bit odd, but not contradictory. But if this sounds contradictory to you, there's another route to the kind of analysis I'm looking for. Katz [2003] suggests that the Perfect actually stacks on top of other aspectual operators like the perfective or the imperfective. What we are looking for is an account of the Perfect according to which embedded stative predicates do not have to be completed by the utterance time, but eventive predicates do. If we recall the analysis about stative and eventive predicates in the present tense, the suggestion was that eventive predicates trigger default perfective aspect, and statives imperfective. If the Perfect stacks on top of

event-relative modal semantics, but with the amendments I've been suggesting, we then get the following.

(104) a. MIGHT [PRES [PERF [Ice cream be in the fridge]]]

b. $\lambda w \lambda e_2$. [ASSERT$'$(e_2, w) & $\forall w' \in$CON(e_2): $\forall w'' \in \text{BEST}_{g(e_2)}(\bigcap f(e_2))$: [$t = t_u$ & $\exists e_1$ [time(e_1)$\prec t$ & Ice-cream-be-in-fridge$'$(e_1)(w'') = 1]]]

As a consequence of the first part of my proposal, e_2 is actually a past event, corresponding roughly to the time at which Sophie looked in the fridge. Because the epistemic modal projects from e_2, the modal base will be calculated to determine the set of worlds characterizing the information state at e_2. Also, t_u wouldn't be the time of the actual utterance, but the time of e_2. But with our characterization of the Perfect as above, we have the result that the ice cream's being-in-the-freezer can continue through to t_u, even though an entire ice-cream-being-in-the-freezer precedes $t{=}t_u$.

This gives us the reading we are after, but there are two questions that still need addressing before we can put this matter to rest. The first question is this: according to the semantics just offered, where the illocutionary event is deictically identified with a past discourse event, the truth conditions for present tense and imperfective would be entirely appropriate. So, while the truth conditions for (104a) give us truth conditions to describe the hindsight cases, so do the truth conditions we'd get with the configuration in (105).

(105) MIGHT [PRES [IMPF [Ice cream be in the fridge]]]

---

other aspectual operators, then assuming eventives and statives trigger the same perfective/ imperfective operators even when the Perfect is present, we can get the desires result. This would effectively make the perfective operator responsible for the cumlination intuition. For what it's worth, I prefer the non-stacking option, because it strikes me that activity predicates under the perfect aren't completed. But if my impression is wrong, the stacking option gives us what we want.

However, that would mean a sentence like (106) could just as well be used to secure the kind of hindsight reading we are considering.

(106) The might be ice cream in the fridge.

We just don't see this – why not? Why does it seem that the Perfect is obligatory in these hindsight interpretations, when something like (106) should do, going only by the truth conditions?[28]

It's worth noting that, if the claim that the hindsight readings are instances of FID is sustainable, the question doesn't arise. We can agree with von Fintel and Gillies that *might have* is an epistemic modal and a past tense morpheme (though we don't say that the former scopes over the latter), but we can then say that the past tense morpheme is a requirement of FID, but doesn't semantically contribute to the interpretation of the sentence. This would allow us to have our cake and eat it to, because it would mean that the sentence is pronounced *There might have been ice cream in the freezer*, but semantically it would be equivalent to (106).

Of course, since I wound up taking a different tack with my analysis of hindsight modals, we have to deal with this question. I don't have a definitive answer for it, but I can offer a few speculations. I suspect that the answer is pragmatic. When the modal is keyed to the actual discourse event, to the event at which the utterance takes place, it's clearer to the conversational participants when this event takes place – in spoken discourse, one's conversational participant is present to perceive the event and locate it in time. But when a discourse event has to be evoked in discourse, the exact contours

[28]Note that this is a problem that the von Fintel and Gillies do not have, precisely because they interpret the *might have* as past tense + modal. For them, the *have* is obligatory in hindsight intepretations because it's the instance of past tense which, when it scopes over the modal, will shift back the temporal perspective of the modal to give us the hindsight interpretation.

and temporal location of the *proposed* discourse event is not as apparent. Would it be the moment Sophie opens the freezer? The minutes leading up to it? The time at which she *decided* to check the freezer for the ice cream? Use of the Perfect renders the *exact* resolution of proposed discourse event moot, because the issue is the time at which the ice cream was thought to be in the freezer. And in using the Perfect, what's being uttered is that there was a full event of the ice cream being in the freezer prior to this hypothetical speech event, so it matters little the precise time at which the event is thought to be. This explanation is a bit sketchy, but the gist is that the use of the Perfect is pragmatically imposed on the speaker in virtue of their keying the modal not to the actual speech event, but to the evoked, past discourse event.

The second question is the following. Why should the interpretation of hindsight readings of modals require of us to resolve the location (even if only approximate) of a proposed discourse event? Isn't this a complication we should try to avoid in our analysis? My response here is unequivocal – the von Fintel and Gillies view, according to which past tense scopes over the modal, requires of us much the same thing. Here's why. Let's consider two different theories of tense, a more "weak" theory and a "strong" theory. The weak theory is just the view that past tense quantifies over past times. On this type of theory, PASTϕ is true at t just in case ϕ was true at some t' prior to t.[29] The reason I'm calling it a "weak" theory of tense is that it's truth conditions are fairly weak – since it existentially quantifies over past times, it is fairly permissive in the truth conditions it

[29]The way I've characterized the "weak" theory, it matters little whether we characterize tense in terms of an intensional operator or as an object-language quantifier. We can imagine it being the latter, so it's implementation would be consistent with the way I've treated tense elsewhere in this dissertation.

engenders. *Alexander went to sleep* on this occasion is true just in there was at least one instance prior to the time of utterance at which Alexander went to sleep.

Of course, the very permissiveness of such theories raises concerns that inspired an alternative approach to tense. For example, we have the famous examples in Partee 1973 which suggest that tense does not quantify over past times so much as refer to specific times, much as a pronoun refers to specific individuals. The example above is a case in point; an utterance of *Alexander went to sleep* doesn't seem to be true just because Alexander went to sleep at some point in time before in his life. No, intuitively, an utterance of this is meant to be true just in case Alexander is asleep at of some particular salient time, like 20 minutes ago, when his room fell silent.[30]

Now, let's return to the original scenario, reproduced again as (107).

(107) CONTEXT: Sophie is looking for some ice cream and checks the freezer. There is none in there. Asked why she opened the freezer, she replies:

a. There might have been ice cream in the freezer.

b. PAST(*might*(ice cream in freezer))

With respect to this question (why propose an analysis that has to resolve a past discourse event), there's a bit of a dilemma for the view that past tense shifts back the the temporal coordinate relative to which the modal is evaluated. If this is the case, then

[30]Of course, there is room to maneuver in this. It's a common move in philosophy of language to say that quantifiers are unrestricted, but we understand them in a restricted sense. The move here is that the locutionary content, or what is said, is unrestricted, but the illocutionary content is understood in a more restricted sense. An analogous move could allow one to say a similar kind of thing in response to this observation: tense might operate as the weak view suggests as far the locutionary content of uttered past tense sentences suggests, but the illocutionary content is restricted to be about a particular time. To forestall this kind of response, Partee canvases a number of other data which suggests that tense behaves as though it introduces pronouns into the object language.

it will do so either either according to the weak view of tense or the strong view of tense. But the weak view is much too weak, for it would mean, roughly, that at some point in time prior to now, it was consistent with Sophie's evidence that there is ice cream in the fridge. If this has the sense that Sophie's evidence merely leaves it open whether there is ice cream in the fridge, or that she doesn't have evidence ruling out there being ice cream in the fridge, then in all likelihood, there such a condition is met. For example, if Sophie ever had ice cream in her fridge before, then the sentence is true. If she was ever unaware of the fact that she *didn't* have ice cream in her fridge, then the sentence is true. So the weak interpretation of tense yields much too weak truth conditions – they are much too easy to make true. And because they are so easy to make true, it makes Sophies's utterance entirely uninformative, and we lose the appropriateness of Sophie's response to the question of why she told us to look in the fridge. If we want to know why she asked us to look in the fridge then, the fact that at some point in time in the past it was consistent with her evidence that there was ice cream in the fridge does nothing to explain this request.

Obviously, I don't think that the advocates of this analysis of the hindsight case are committed to the weak view, or even have something like this in mind. In fact, it is evident that they don't. On this conception, the truth conditions are much too weak – an advocate of the view that tense shifts the temporal perspective of the epistemic modal would do well to think of this proposal in terms of the strong view of tense. Let's recall one of the examples from above. In uttering *Alexander went to sleep*, I am not so much claiming that there is some past time (any old time) at which it is true that Alexander goes to sleep, but that there is a *particular* salient time relative to which it is true. If

my wife and I are worried about Alexander romping around and staying up late, the salient time will be a time in the interval at which we finally no longer hear him stirring in his room. What does this mean for the the modal case? On the strong view of tense, interpreting past tense will often require us to resolve the value of temporal variable – something like a temporal pronoun. And, to do this, we would need to determine the most salient time to serve as the value of the tense variable. But that would require of us to know a bit about past events and their temporal locations so we could actually determine the candidate value of the tense variable.[31]

The point is this. Either the Tense>Epistemic Modal view adopts the weak view of tense, which yields implausible truth conditions for the epistemic modal, or it adopts the strong view, which requires one to keep track of the location of past events in order to determine the salient time to serve as the value of the tense variable. I take it as given that proponents of the view don't adopt the implausible weak view. But this means, dialectically, that they face the same pragmatic task as the one I am proposing in having to resolve the time at which the modal base is to be calculated.

4.3. the Eventivity Constraint in Modals Beyond *Must*

There's something of an elephant in the room that needs to still be addressed. So far, the bulk of the discussion has focused on a series of examples involving *must* and its behavior when combined with prejacents of various types. I've made the case that *must*

[31]In fact, there's a more plausible version of the weak theory of tense, which I've put off until now. Rather than simply quantifying over the entirety of past times, the more plausible weak theory would quantify over past times within a restricted interval. I've put off considering this until now, because this makes the pragmatic task very similar to the one that faces us on the strong theory of tense – determining the candidate restricted interval which is quantified over requires being able to locate it among past events which make the interval salient.

conforms to the EC. However, one wonders whether *must* is an outlier in this behavior, or whether this this behavior is more general for modals. A quick glance at the kinds of sentences that motivated the postulating of a constraint like the EC suggests that other modals probably don't so conform.

(108) a. John must go to the store. (epistemic: bad; deontic: OK)

b. John may/ might go to the store. (epistemic: OK; deontic: OK)

c. John should/ ought go to the store. (epistemic: ???; deontic: OK)

So, the EC describes a constraint on interpretations of the modal *must*. However, the explanation I gave for the EC in this dissertation is not particular to *must*. Part of this is because the framework I based my explanation on is a general one that applies to modal auxiliaries across the board. In light of this, whether or not other modals conform to the EC is an issue that needs addressing. Put simply, are other modals subject to the EC constraint? If they are not, is there some explanation for why they are not, given that the explanation I gave for the EC would seem to make the prediction that they should be? This last point is important, dialectically – it does seem that my explanation would predict that something like the EC holds for other modals auxiliaries like *ought/ should*, and *may/ might*. That being said, the examples in (108) provide a prima facie challenge to this explanation, insofar as it seems that *any* of them admit of epistemic readings. My task, then, is the following. I will make the case that initial appearances are deceiving. The EC is plausibly a constraint that these other modals conform to as well. This is a more modest task than *providing a semantics* for these other modals. The reason for my adopting this more modest task is that there are complications that a semantics of each of these modals must address. I won't here venture such a semantics, but will instead

make the case that despite the intuitions we have in about putative epistemic readings in (108) is consistent with the EC holding for modal auxiliaries across the board. If this is so, the conclusion we can draw is that the the EC issues a constraint on the kinds of semantics these other modals can have.

This section will then breakdown into several subsections. In the first subsection, I will review *ought* and *should* to survey the evidence of these modals' conforming to the EC, determine how one might defend the view that they do. In the second, I will do the same thing with *may* and *might*. In a third, I will take a look at modals other than the auxiliaries focused on thus far. So I will look at whether modal adverbs like *possibly* conform to the EC.

4.3.1. On *ought/should*

To a first approximation, it seems like *ought/ should*-sentences with eventive prejacents plausibly have an epistemic reading. But this matter is complicated here by the fact that intuitions about the flavor of *ought/ should* sentences can be very subtle. In fact, they are subtle not only with respect to sentences with eventive prejacents, but also with stative prejacents. This is unlike with deontic interpretations, which are rather prominent and the intuitions are quite strong. For example, it seems difficult to get an epistemic reading of (119a) above, but the my judgment, nor that of several interlocutors I queried was as strong as with *must*. How about with a stative prejacent?[32]

(109) John should be at the store.

[32]I'll usually revert to using *should* for my examples and discussion in this section for ease of exposition. But what I say in this section is meant to apply equally to epistemic *ought* as well.

At first blush, it seems that (109) has an epistemic reading of *should*, based on something as coarse-grained as our intuitions about the sentence without any consideration of the context. If (109) has an epistemic reading, but (119a) does not, this bodes well for the claim that the EC holds for *should*. Still, it quickly becomes clear that such coarse-grained intuitions are unreliable with supposed epistemic *should*. The reason for this is as follows. *Ought* and *should* are typically thought to be weak necessity modals in the sense that are not quite as strong as *must*.[33] Evidence for this comes from the fact that it seems that, holding fast the context, MUST ϕ entails SHOULD ϕ, which in turn entails MIGHT ϕ.

(110) MUST $\phi \models$ SHOULD $\phi \models$ MIGHT ϕ

Indeed, this seems to hold when the flavor of the modal is deontic. An example like the following is pointed out by von Fintel and Iatridou [2008].

(111) a. You must wash your hands.

b. You should wash your hands.

c. You may wash your hands.

(112) a. You should wash your hands. In fact, you must.

b. # You must wash your hands. In fact, you should.

In (112), the *should*>*must* discourse is acceptable, though the *must*>*should* discourse is odd. This is as expected, if the strength of the *must* sentence is stronger. Having uttered the *must* sentence first, the *should* sentence offers no additional information. And, of course, it's eminently plausible that an obligation to ϕ entails permission to ϕ, so the relation between the sentences in (111) reflects the entailment relations affirmed in (110).

[33]This has been pointed out by Horn 1989.

It's not clear whether the apparent weakness of *should* has much of a bearing on whether or not it is an EC modal. However, there is a considerable and growing body of evidence presented in the recent literature militating against the idea that all the apparent epistemic readings are in bona fide epistemic modals. To start with, the following following was noted by Copley [2006].

(113) The beer should be cold by now.

That this is plausibly epistemic is bolstered by noting that it could be uttered in the following kind of circumstance: we've put the beer in the cooler, covered it with ice and let it sit in the cooler for a while. This has always cooled the beer to an acceptable temperature to drink. (113) seems appropriate in this circumstance. It also appears that, whatever (113) means, we are using knowledge to infer the high likelihood that the beer is cold. Moreover, *The beer must be cold by now* is equally acceptable, and there's little question that this has an epistemic reading. However, things start to look very different when we make note of the following pattern.

(114) a. #The beer must be cold not, but it isn't.

b. The beer should be cold now, but it isn't.

c. #The beer may be cold now, but it isn't.

The respective *must* and *might* sentences are odd, but the *should* sentence is fine. There are several explanations for why (114a) and (114c) are bad. One such explanation, suggested by Copley, suggests that with an epistemic reading, the sentences are Moore Paradoxical. That is, if the second conjunct conveys that one knows that the beer isn't cold, then it can't be the case that it follows from one's knowledge that the beer is cold;

or even that it's left open by one's knowledge that the beer is cold. Since (114b) isn't bad in the same way, this suggests that *should* isn't actually epistemic.[34]

Here's another example that appears in Yalcin 2016 (example numbering changed to reflect the numbering in this work).

> 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:
>
> (115) I should be dead right now.
>
> (116) I ought to be dead right now.
>
> In a similar vein, Jones's sister says things like this:
>
> (117) It's incredible! That quake was massive. He should be/ought to be dead. We're so lucky he survived.

While, Yalcin maintains, both of (115) and (116) are plausible candidates for epistemic interpretations, it would be entirely inappropriate to utter *I might be dead right now*. But if the entailment facts as I've described them hold, and *should* is epistemic, then *I might be dead right now* should actually be an entailment of (115) and (116). Yet, while *I might be dead right now* clearly has an epistemic reading, that it is inappropriate in this situation suggests that (115) and (116) are not themselves epistemic.

[34]In fact, there are competing explanations for the fact the *might* and *must* sentences are bad. The sentences in (114) have the form MIGHT ϕ & not ϕ. According to the Moore Paradoxical explanation, the sentences are pragmatically infelicitous but not actually contradictory. On the basis of embedding behavior, Yalcin [2007] contends that these kinds of conjunctions are contradictory, and calls them *epistemic contradictions*. On this basis, he offers semantics of modals which makes (114a) and (114c) contradictory. Cf. also the analysis in Mandelkern 2019a, according to which the schema $\ulcorner$MIGHT ϕ & not $\phi\urcorner$ is inconsistent, but the resulting semantics is more amenable to classical contextualism. The difference in explanation matters little for our present purposes, because all parties to the discussion agree that the oddity of (114a) and (114c) has essentially to do with the modal being epistemic.

Yalcin and Copley both suppose that in the examples just just discussed, *should* is not, contra initial appearances, an epistemic modal. Yalcin calls such readings *pseudo-epistemic* readings, and proposes a non-epistemic semantics for them. Rather, Yalcin claims, they are modals of *normality*, a reasonable paraphrase for which might be glossed as follows.

(118) α should/ought to F $\approx$ It would be normal, all relevant things considered, for α to F

Based on this suggestion, Yalcin characterizes the semantics for such normality modals as quantifying over expectation-laden information states. These are information states whose worlds are ordered by an expectation pattern to determine the set of normal worlds (normal with respect to that expectation pattern). In spite of the fact that the quantificational domain is partially determined by an information state – the same parameter that serves as the domain of quantification for epistemic modals, Yalcin emphatically contends that this is not an epistemic modal. Rather, the information state for normality modals is not a representation of a person (or group's) state of knowledge, but a state characterized by the circumstances a person is aware of. The normality ordering on those worlds represents the expectation a person has given those circumstances. Though this isn't an epistemic state, it does not float entirely free from knowledge, either, since the circumstances we are aware of supervenes on what we know.

In spite of this, this Yalcin doesn't think a Kratzer-style semantics with a circumstantial modal base and an ordering source (a stereotypical ordering source is perhaps the go-to candidate ordering source for this construction) is the best way to model this normality modals. Chief among his concerns is 1) the worry that other modals which

have circumstantial readings don't themselves have the kind of pseudo-epistemic readings that *ought* and *should* have, and 2) that the conception of the modal base needs to be finessed in order to account for the data, in particular, such that actual world needn't actually be part of the most normal worlds. I find (1) an interesting data point, but not a compelling reason to jettison the Kratzerian framework. It's already part of the framework that certain modals have lexical idiosyncrasies that rule out certain readings (think of *might* not being available for a deontic reading). So, some such explanation could be operative here, depending on how much lexical idiosyncrasy in one's theory one is willing to tolerate. On point (2), I concur, but again don't find this a reason to give up the framework. Copley's diagnosis is similar with respect to the epistemicity of these modals. She argues that these pseudo-epistemic readings actually involve a metaphysical modal base a la Condoravdi 2002, and a kind of *highest entropy* ordering source, where the latter is understood to rank the worlds according to worlds higher where no other events intervene, all else equal.

I will not myself hazard an account of the kind of *should* and *ought* here. But Yalcin and Copley's diagnosis is salutary for my account of the EC. The upshot of this discussion is that those instances of *should* which appear to be epistemic are not, in fact, and they are better modeled as being interpreted with respect to a set of worlds characterizing a set of circumstances. In event-relative semantics for modals adopted here, this would make pseudo-epistemic *should* a root modal.[35] Obviously, if there are no genuinely *should* sentences, it follows that there are epistemic *should* sentences with bare eventive prejacents. Granted, this fact have more to do with idiosyncrasies of *ought* and *should*, as opposed

[35]This means that giving an event-relative semantics for *should* requires that one take seriously Yalcin's worries about the nature of circumstantial modal bases.

to the EC, and so it becomes difficult to say with certainty that the EC is operative in ruling out epistemic readings of *should* with bare eventive prejacents. However, it could simply be that the lack of epistemic readings with should and ought is overdetermined. Whatever the case may be, the evidence here doesn't militate *against* the EC.

This is not the end of the story on *ought* and *should*, however. First of all, a number of theorists (Cariani 2013, Finlay 2014) have maintained that *should* and *ought* do have genuine epistemic interpretations, so we should at least seriously consider the possibility in our discussion of the EC. Second of all, the examples canvased above all had stative prejacents, and it seemed like the pseudo-epistemic readings of *ought* and *should* had hints of a counterfactual reading distinguishing it from the pure epistemic reading.[36] Yet, *should* with eventive prejacents doesn't seem to have this kind of counterfactuality, and so doesn't as clearly distinguish itself from, say, epistemic *might* in quite the same way as *should* with a stative prejacent.

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

b. #The beer might be cold by now, but it isn't.

Recall Copley's observation that (119a) isn't acceptable, but (119b) is. Though, as we saw, Copley gives a Moore-paradoxical explanation of the unacceptability of (119a), Yalcin treats (119a) as expressing a contradiction. Indeed, for Yalcin, the possibility of such contradictions is diagnostic of epistemic readings of a modal. If this is so, it's telling what happens when we replace the stative prejacent with an eventive one. Consider the comparison in (120).

[36]For example, in Yalcin's scenario with the utterance of (115), an utterance of *I should have died* seems equally as acceptable and appropriate as *I should be dead*.

(120) a. John should go to the store, but he won't.

b. #John might go to the store, but he won't.

First, it's clear that both (120a) and (120b) are future-oriented, unlike (119a) and (119b). To be sure, (120a) also has a strong – perhaps even *preferred* – deontic reading. And on the deontic reading, (120a) is fine.[37] (120b), on the other hand, is categorically pretty bad.

One thing to note is that *might* has no deontic reading, and it's on the deontic reading that (120a) is acceptable. When we train on the *epistemic* reading of (120a), it sounds pretty bad as well. Since it has the form ⌜MODAL Φ and NOT Φ⌝, it is reasonable to suspect that it is a kind of epistemic contradiction like (120b), and that it owes its badness to a similar cause. Since it's the contradiction-inducing conjunct (*but he won't*) that seems to renders the epistemic reading of (120a) unacceptable, this forces us to seriously contend with the possibility that *John should go to the store* does have an epistemic reading.

Pseudo-epistemic readings of *should* and *ought* aside, this makes some trouble for my proposal that *ought* and *should* both conform to the EC. My proposal, hinging on putative general facts about modals and their location in clausal architecture, suggests that there is no reading of *John should go to the store* that is epistemic, for the same reason there is no (non-futurate, non-habitual) reading of *John must go to the store* that is epistemic – because ⌜MODAL [John goes to the store]⌝ is ruled out because *John goes to the store* is unacceptable.

[37]This is one of the features that distinguishes deontic *should* from deontic *must*; a *must*-sentence conjoined with an assertion that the prejacent won't happen is usually taken to be unacceptable. Cf. Ninan 2005; von Fintel and Iatridou 2008.

It's clear, though, that on the semantics we adopted, we need some kind of intervening operator between the modal and the Verb Phrase, at the very least to explain the temporal and aspectual facts about the interpretation of the sentence. For example, (120a) is future oriented, though without some mechanism for generating the future orientation, the semantics I presented in section 3.7 would get the temporal interpretation of the sentence wrong.

Here we can consider the proposal discussed in section 4.1.2 by Kratzer [2010] and Matthewson [2012], that there is a prospective operator responsible for securing the future orientation of the modal. When I considered this proposal as the source of future orientation for *root* modals, I found that the proposal ran into some trouble.[38] But the trouble was with the EIP (event identification problem), and had to do with the fact that the argument to the modal base parameter was the same as the argument to the prospective operator. As I argued, these two couldn't be the same. But for *epistemic* modals, there is no such problem, because the modal base parameter takes an argument completely outside the vP to begin with.

However, simply appealing to this prospective operator is an unsatisfying move, which becomes apparent when you take a closer look at how to leverage this into an explanation of epistemic *should* that is friendly to my explanation of the EC. Recall that my explanation of the EC was as follows: *must* + eventive prejacents are out because eventive prejacents in present tense trigger perfective aspect, and present perfectives are unacceptable. By contrast, epistemic *should* + eventive prejacents is OK because in these environments, eventive prejacents in the present tense would trigger *prospective* aspect.

[38]Cf. section 4.1.2.

This is unsatisfying because it leaves unexplained why the prospective aspect solution is unavailable to the the complements of epistemic *must*. Why is it that, in order to get an epistemic reading, *must* complements with eventive predicates need to have either habitual or futurate operators but *should* can trigger prospective aspect?

So it's worth looking elsewhere for an explanation for epistemic *should* with eventive prejacents – one that can also explain why such sentences have a future orientation. But thinking through this case can prove instructive, because the issue that prompted us to consider the inclusion of prospective aspect was the need for something in the semantics to secure the sentence's future orientation. We've just considered the possibility that prospective aspect has this function, and rejected this possibility because it made the complements of epistemic *must* and *should* too different without providing an explanation for this difference. But as noted earlier, epistemic *must* can have a future orientation when a futurate operator is present. And, I've argued in section 2.4, this is consistent with the EC since the resulting expression is stative.

The proposal I wish to advance is that epistemic *should* with eventive prejacents, like those in (120a), actually have a futurate operator intervening between the vP and the modal, just like the *must*-sentences discussed earlier in the dissertation. In other words, the logical form of epistemic should sentences with eventive prejacents is actually:

(121) SHOULD [FUT Φ]

This proposal makes the logical form of eventive, epistemic *should*-sentences parallel that of eventive epistemic *must*-sentences, and moreover, and makes both consonant with the EC. This may be surprising, since utterance of an epistemic *must*-sentences with futurate prejacents are only acceptable in fairly exceptional contextual circumstances, and the

observation we started with was that certain *should*-sentences seemed to have a more permissive distribution. Grist for this mill included the observation that *must* can receive an epistemic interpretation *much* more easily when the eventive predicates embedded under the modal more easily admit of scheduled readings.

(122) a. The Yankees must play the Red Sox (tomorrow).

b. The plane must leave in the morning.

But it is harder to get an epistemic reading of *must* when the predicate for the complement is not as easily interpreted as describing a scheduled event, or one planned in advance.

(123) John must go to the store (tomorrow).

In (123), it is quite hard to get an epistemic reading of *must*, but it becomes much easier for this sentence to get an epistemic reading when the context makes clear that what is at issue are John's plans for tomorrow. Recall the example discussed in section 2.4 (cf. example (44)) that such a context can in fact support an epistemic reading of *must*.

> CONTEXT: You and I are are discussing an upcoming party John is hosting the day after tomorrow. We know that John intends on buying a lot of supplies for the party. We see that his cupboards are still bare, so we attempt to determine when John will go to the store for supplies. You note that John has to work late today and wonder aloud what his plans are for tomorrow. I retort:
>
> (124) John must go to the store (tomorrow).

Such a context makes it seem like it follows from our evidence that *John goes to the store tomorrow*, a futurate sentence, is true. But with epistemic *should*, we don't see the

same kind of acceptability profile – non-planned predicates seem to allow for an epistemic reading of *should* much more easily than with *must*.[39]

This can be explained by borrowing an old idea from Stalnaker [1975] and being sensitive to the pragmatics of epistemic modals. First, let me attempt a rough gloss on the contextual contraints under which epistemic *must* can embed a futurate. I suspect that the reason for the exceptional contextual circumstances required for epistemic *must* with futurate prejacents is due to the fact that it needs to be common ground among the interlocutors that the event described by the predicate is in some sense *settled*.[40] An event being planned or scheduled, and presumed not to be intervened with or prevented from happening, is sufficient for it to be treated as settled by the interlocutors. As mentioned earlier (cf. footnote 25), epistemic *must* has some interesting pragmatic features quite apart from our discussion of the EC. It is now relevant that an utterance of an epistemic *must* sentence seems to be "weaker" than an utterance of unmodalized analog of the sentence. There have been several proposals for this apparent weakness.[41] I will briefly present two: von Fintel and Gillies [2010] suggest that felicitous utterance of an epistemic *must* sentence requires evidence for the prejacent be indirect. This is why, if you look outside the window and see rain, it would be odd to utter: *It must be raining*. However,

[39]It might be that circumstantial and epistemic readings of *should* are often confused – enough that circumstantial *should*-sentences will sometimes have a pseudo-epistemic reading that Yalcin discussed. And *disambiguating* circumstantial from epistemic *should* requires making the kinds of contextual circumstances clear that were made apparent in the discussion of (124). (Cf. Condoravdi [2002, p. 80]'s useful discussion on a related point.) But the fact remains that *must* absolutely resists an epistemic interpretation unless these circumstances are clear to the conversational participants. The fact that there is even a potential for ambiguity in the case of *should* suggests that this constraint isn't nearly as stringent as with *must*.

[40]Cf. Kaufmann 2002 for the notion I have in mind.

[41]There are many. A point of contention in the debate surrounding this phenomenon is whether this weakness is to be accounted for semantically or pragmatically. It is the pragmatic solutions that I find most useful for my present purposes.

if you see people with wet rain gear entering the lobby, such an utterance is entirely felicitous. In the first scenario, you have direct evidence for the truth of the prejacent, so the modal utterance is out. In the second, your evidence is indirect, so the modal sentence is OK.

A related proposal is advanced by Mandelkern [2019b].[42] Mandelkern argues that a hallmark of utterances with epistemic *must* is that they are felicitous only if there is a shared argument for the proposition the *must* statement embeds. This is supposed to improve on von Fintel and Gillies's example showcasing the "indirect evidence" for the following reason. Suppose that we are entertaining skeptical scenarios, and I try to convince you that your direct perception of the rain is due to some conspiratorial deception. You retort by appealing to the clarity of your percept and the unlikelihood of such a deception. In this case, it is significantly less inappropriate for you to utter *So it must be raining*, even if your evidence remains direct (and, presumably, was initially direct as well). But this is so because, even though you have direct evidence of the rain (if perceptual evidence be considered direct), you can still utter the *must* sentence because, on account of the skeptical scenario being discussed, the conversational participants are aware of a shared argument (involving the perception of rain and the claim that it is veridical) which entails that it is raining.[43]

I will not defend either of the two views about the weakness of *must* directly, but I simply want to note how the both of them square nicely with the proposal I am making.

[42]Cf. also Mandelkern 2018.

[43]This is an improvement over the von Fintel and Gillies account because the alternative requires fretting about what exactly constitutes *indirect* evidence such that even perceptual evidence can be indirect in the right circumstances. Mandelkern's proposal allows us to avoid having to stake a commitment about this.

Notice that in the context spelled out in the lead-up to (124), a number of important facts are stipulated to be apparent to the interlocutors ("you" and "me") – that John is having a party for which he intends to buy a lot of supplies, that he has yet to do so, and that his only available time for doing so is tomorrow. Now, obviously we don't have *direct* evidence of John's going to the store tomorrow; it being a statement about something that hasn't happened yet, there's no way we could have this. But the circumstances spelled out give us sufficient *indirect* evidence that this will happen, as per the von Fintel and Gillies account.

Mandelkern's account provides us even more material we can leverage for an explanation of the distribution of epistemic *must*. Recall that on Mandelkern's account, felicitous use of *must* sentence requires that an argument for the prejacent be available as common ground to the interlocutors. Since we are here primarily concerned with epistemic *must* sentences with futurate prejacents, this condition has to square with the felicity conditions for futurates, which are such that the event described has to be able to be treated as settled by the conversational participants. Combining these two conditions yields a constraint on utterances of epistemic *must* sentences with futurate prejacents that an argument has to be common ground to the interlocutors to the effect that the event described by the prejacent is able to be treated as settled. This is precisely what we see in contextual circumstances to (124). These circumstances make the relevant premises apparent to the interlocutors, which in turn provide an argument for the settledness of event of John's going to the store tomorrow, by, for example, providing evidence that he has the requisite intention to do so. The particular distribution of the epistemic *must* is

such that without all of this being manifest to the conversational participants, the *must* statement would not be felicitous.

So much for the pragmatics of epistemic *must* with eventive (and so, futurate) prejacents. What about epistemic *should* and *ought*? What accounts for the more permissive distribution of epistemic *should* and *ought*? That is, there is an epistemic interpretation of *should* in (125) meaning something (roughly!) like "John is probably going to the store tomorrow".

(125) John should go to the store tomorrow.

It seems that we can get an epistemic reading of (125) without the particular constellation of circumstances that need to be manifest to the interlocutors, as would be needed for (124) to have an epistemic interpretation.[44] The way to explain this is to draw from an interesting proposal made by Stalnaker [1975, p. 276]:

> I take it that the subjunctive mood in English and some other languages is a conventional device for indicating that presuppositions are being suspended, which means in the case of subjunctive conditional statements, that the selection function is one that may reach outside of the context set.

Though his concern is primarily to explain some phenomena in conditional semantics by way of a difference between subjunctive and indicative conditionals, his proposal isn't exclusively about conditionals, but about the function of the subjunctive mood. This connects with our present discussion about *should* and *ought* as follows.

Should and *ought* have a noteworthy lexical property – they both have past tense morphology on them which seems to be moribund in that it doesn't function as a regular

[44]Cf. the discussion in fn. 39.

tense.[45] That is, *should* was originally the past tense form of *shall*, and *ought* was the past tense form of to *owe*; but while these expressions retained the past tense morphology, they have shed the past tense meaning. They are quite normally used with a present tense meaning. Iatridou [2000] and von Fintel and Iatridou [2008] have pointed out that, cross-linguistically, past tense morphology (especially that which doesn't function as a semantic past tense – what Iatridou has called "fake past") tends to denote subjunctive or counterfactual meaning. The idea is that *should* and *ought* bear a morphological marker (past tense morphology that doesn't semantically function as a past tense) associated with subjunctivity. And we treat this subjunctivity along the lines suggested by Stalnaker, as a conventional way to indicate that some presuppositions are being suspended.

We can use this observation to describe what is going on with epistemic *should*. Epistemic *must* places strong felicity constraints on the use of the *must*-sentence: *must* requires that an argument entailing the prejacent be available to the interlocutors. When the prejacent is a futurate, this means that the argument must directly entail something about the settledness of the prejacent. However, according to the proposal I am floating, *should* suspends this requirement – that is, an argument does not need to be available to the interlocutors for the felicitous utterance of an epistemic *should* sentence. So, epistemic *should* with futurate prejacents make an assertion as to the settledness of the prejacent (so, they make a claim about a future event occurring), but do not presuppose the settledness of the event, as is the case with epistemic *must*.

[45]Ramchand [2018] makes this observation as well, and it has important consequences for her proposed modal semantics, though I put the observation to slightly different use.

Since epistemic *should* does not require quite the collection of information to be common ground that *must* does, we expect the modal *should* to more easily receive an *epistemic* reading than *must*. And, as a consequence, epistemic *should* has a more permissive distribution than epistemic *must*. What I take this discussion to show is that we do not need a different treatment for *must* and *should* in our modal semantics when it comes to explaining the EC. We can think of both *must* and *should*/*ought* as conforming to the EC with the following two caveats: 1) some of the apparent epistemic instances of *should* or *ought* might only be apparent and are better analyzed as the kind of pseudo-epistemic readings described by Yalcin and Copley. 2) even though (genuine, not pseudo-) epistemic *should* conforms to the EC and so requires of non-habitual eventive prejacents to be futurates, the assertability profile of epistemic *should* makes the futurate structure less easy to detect than with epistemic *must*, which is more transparent.

4.3.2. On *might* and *may*

To make the case for attested epistemic readings of *might*/ *may*, we only need to remind ourselves of an example from last section, repeated here as (126). For it is evident that *might*/ *may* admit of epistemic contradictions, even with eventive prejacents.

(126) #John might go to the store, but he won't.

So, let's accept without further argument that there are epistemic readings of *might* with eventive prejacents. How does this square with the EC? To square this, my strategy is the same as above – to claim that, contra initial appearances, epistemic *might* sentences with eventive prejacents do indeed conform to the EC, and they do so because the prejacents are actually futurates. Again, this fact might not be as apparent at first blush

when compared with epistemic *must*-sentences with eventive prejacents, but as I pointed out in the previous section, that's because epistemic *must* statements have particular assertion conditions that make their assertability profile more limited than that of epistemic *should*/*ought* and epistemic *might*/*may* sentences. So, the strategy is to determine whether some fact about epistemic *might* or *may* can explain the relaxed assertibility conditions of the futurates I suspect to be embedded under the modal.

Now, a key distinction between *must* on the one hand and *should*/ *ought* on the other was the fact that the latter had subjunctive morphology, which we are supposing to allow one to "reach out of the context set" in asserting epistemic *should* statements with futurate prejacents, and relax the otherwise stringent assertibility conditions for futurates. For epistemic *must* with futurate prejacents, by contrast, an argument for the truth of the prejacent must be common ground for the epistemic statement to be assertible. This was what accounted for the limited distribution of epistemic *must* statement with futurate prejacents.

Might similarly has this subjunctive morphology, so we could avail ourselves of the same kind of story for this modal (*might* being the erstwhile past tense form of *may*). However, the auxiliary *may* emerges here as an outlier. *May* does not have subjunctive morphology, and so, by assumption, it doesn't signal the same kind of presupposition-suspending meaning as *might*, or as *should*/ *ought*. But nonetheless, *may* and *might* seem almost synonymous in their epistemic readings.[46] And that would mean that the eventive

[46]There are some lexical differences, of course. For example, *may* can admit of a deontic reading, whereas *might* has no such reading. And because it has this reading, a *may* sentence with an eventive prejacent may tend towards a root interpretation (so, *John may go to the store* may have a more prominent root interpretation), where no such tendency exists for the *might* sentence. But considering both in their epistemic interpretations, *may* and *might* sentences seem synonymous.

prejacents in epistemic *may* sentences are actually futurate prejacents as well. But they have a similarly permissive distribution as the corresponding *might* sentences, in spite of the lack of subjunctive morphology. So why are epistemic *may* sentences with futurate prejacents similarly permissive?

We can find a path to an explanation for this by considering the pragmatics of epistemic possibility statements in general. It's been argued, for example, that the utterance of epistemic possibility statements serve to communicate subjective uncertainty about the prejacent. If the purpose of a typical assertion is to update the common ground (as in Stalnaker 1978), utterances of epistemic possibility statements differ from typical assertions in that they don't seem to update the common ground so much as communicate the a certain proposition is *consistent with* the common ground, or that the proposition *shouldn't be excluded from the consideration* of live possibilities.[47] Some authors have taken this observation about the pragmatic role of epistemic possibility statements as not easily compatible with the kind of contextualism I've been assuming in the dissertation. The worry is precisely that contextualism erroneously treats epistemic possibility claims, an utterance of ⌜MIGHT Φ⌝, for example, an assertion of uncertainty about Φ, not simply as a speech act meant to counsel the conversational participants not to rule out Φ.

Mandelkern [2020] advocates a contextualist-friendly framework for modeling the pragmatics of epistemic possibility statements that simultaneously takes seriously the observation that such utterances serve to negotiate the live possibilities, and does not simply express that a proposition is consistent with a body of information. According to his

[47]Cf. Swanson 2006, 2009. The way Portner [2009] describes Swanson's contention, utterance of statements of epistemic possibility are no longer assertions, but another kind of speech act entirely, "perhapserting", for example.

prospective contextualism, an epistemic possibility statement like ⌜MIGHT Φ⌝ means that Φ will be compatible with the common grand after the assertion has been made and either accepted or rejected by all interlocutors. So in a nutshell, ⌜MIGHT Φ⌝ means that ⌜MIGHT Φ⌝ is compatible with the *prospective* common ground.

Such a framework can explain the distribution of epistemic *may* statements with futurate prejacents.[48] With respect to the typical felicity conditions of futurate statements, the condition imposed by prospective contextualism is pretty weak. To felicitously assert a futurate, the event obtaining has to be treated by the interlocutors as settled. When this is embedded under an epistemic possibility modal, this transforms the constraint somewhat – because now the obtaining of the event doesn't need to be treated as settled by the common ground, its being settled simply needs to be consistent with the common ground after the assertion has been made. So, an assertion of *John might go to the store* amounts to a proposal to make the common ground consistent with it being settled that John's going to the store obtains at some point in the future.

This is just a kernel of an explanation. For me to be able to offer it as a fully fleshed-out theory, it would need to be more precise about what it means for a future event to be settled, and it would need to explain the difference between futurates with an animate agent and so-called "natural futurates" – futurates that aren't thought to be planned or scheduled by some animate agent.[49] I leave this task to future work. However, for now I will draw a tentative but promising conclusion; even though *must* most prominently and most evidently conformed to the EC, an eminently plausible story could be told according

[48]It can explain the distribution of epistemic *might* sentences as well. But since *might* has the presupposition-suspending features associated with subjunctive morphology, there is more than one reason epistemic *might* sentences have a more permissive distribution.

[49]For example, *The sun rises tomorrow at 6:35AM.*

to which other modal auxiliaries do as well. If this story survives further scrutiny, then the EC will not describe the idiosyncractic behavior of one particular modal, but will describe a constraint on modals in general.

A final observation about the consequences of the arguments of this chapter are also worth making. If the arguments of the foregoing sections are on the right track, then it's something of a mistake to think of a modal sentence's future temporal orientation as a meaning component contributed by the modal itself. Rather, the modal by itself seems to have rather little to contribute to its temporal interpretation. As we saw earlier, its temporal perspective depends crucially on other elements in the clause the modal interacts with, as does its temporal orientation.

4.3.3. On lexical modals

In previous subsections in this chapter, I aimed to make the case that modals besides *must* conformed to the EC. This discussion focused entirely on modal auxiliaries. But modal expressions exceed just the auxiliaries. In this final subsection, I will ask whether the EC could plausibly be said to describe a constraint on *lexical* modals as well. An affirmative answer to this question would help further bolster the case that the EC is actually a constraint that pertains to the expression of modality generally, and is not simply an idiosyncratic feature of the auxiliaries.

I won't survey all lexical modals, but a good place to start is with some of the lexical modals that command attention in the philosophical literature. For example, adverbs such as *possibly* and *probably* are oft discussed by philosophers interested in modals – the latter being closely enough associated with epistemic modality that they are sometimes

used to suggest that epistemic modals ought to be given a probabilistic interpretation.[50] So, I will focus a brief discussion on these.

Let's look at the test used on (31) from chapter 2 to see the results of embedding prejacents from various aspectual classes under the modal *must*. The complication here is that while *probably* has an epistemic meaning, it has no clear root meaning (certainly no deontic interpretation; and no obvious teleological interpretation, for example). In light of this lack of root readings for *probably*, we should expect to find anomaly or unacceptability judgments where the prejacents are eventive.

(127) **States**:
   a. John probably knows the answer.
   b. John probably believes the report.
   c. John probably desires a vacation.
   d. John probably understands French.
   e. John probably is happy. / John is probably happy.

(128) **Activities**:
   a. John probably runs.
   b. John probably swims.
   c. John probably pushes a cart.
   d. John probably drives a car.

(129) **Achievements**:
   a. John probably recognizes the answer.
   b. John probably spots his car.
   c. John probably reaches the summit (by this afternoon).
   d. John probably dies.

(130) **Accomplishments**:
   a. John probably paints a picture.
   b. John probably makes a chair.
   c. John probably delivers a sermon.
   d. John probably draws a circle.

[50]For example, cf. Yalcin 2007, 2010; Cariani 2016.

e. John probably recovers from an illness.
f. John probably builds a house.

fact, this is precisely what we find, with only slight variations on the results from when the test was applied to *must*. The sentences will stative prejacents are fine, as expected. Some of the sentences with activity prejacents are OK, but we quickly find that these are OK when the prejacent is interpreted habitually, just as we saw with *must* as well. Otherwise, the probably-sentences with eventive prejacents are all bad. Interestingly, (129a) and (129b) sound acceptable, but the predicates seem to have a state interpretation on this acceptable reading.[51] So far, this all comports well with my defense of the EC. For, if the EC described a general constraint on the expression of modality, this is exactly what we'd expect to find.

We can observe a similar pattern with the modal adverb *possibly*.

(131) **States**:
a. John possibly knows the answer.
b. John possibly believes the report.
c. John possibly desires a vacation.
d. John possibly understands French.
e. John possibly is happy.

(132) **Activities**:
a. John possibly runs.
b. John possibly swims.
c. John possibly pushes a cart.
d. John possibly drives a car.

(133) **Achievements**:
a. John possibly recognizes the answer.
b. John possibly spots his car.
c. John possibly reaches the summit (by this afternoon).

[51]You can test this. If you placed these sentences into a discourse where they had this acceptable reading, they wouldn't advance the narrative time. Cf. Kamp and Reyle 2013 for this test.

d. John possibly dies.

(134) **Accomplishments**:

a. John possibly paints a picture.

b. John possibly makes a chair.

c. John possibly delivers a sermon.

d. John possibly draws a circle.

e. John possibly recovers from an illness.

f. John possibly builds a house.

Possibly seems to show the same interpretive pattern as probably.[52] Again, this is salutary for my claim about the EC.

There is, however, one complication worth addressing. In my discussion of *might* and *may*, I argued that there was actually an embedded futurate in the sentences, and that the pragmatics of epistemic possibility claims were such that it relaxed the felicity conditions on the embedded futurate. What we see in the case of *probably* or *possibly* is simple anomaly (or perhaps ungrammaticality). But these *possibly* and *probably* sentences are epistemic possibility statements as well.[53] So, shouldn't this also apply to these modal adverbs expressing epistemic possibility? In fact, their distribution suggests that they do not have an embedded futurate. Is this a problem for my claim?

[52]To my ear, some of the sentences sound a bit degraded, though still marginally acceptable. According to my idiolect, I'd prefer the periphrastic modal expression *It is possible that...* to the frame adverbial *possibly*. But changing the examples to *it is possible that....* overtly embeds full sentences, and makes an even stronger case for the EC holding of this expression.

[53]We might not want to assimilate *probably* to epistemic possibility since it's making a stronger claim than a typical epistemic possibility claim. On the other hand, it is making a considerably *weaker* claim than a typical epistemic necessity claim. This could very well suggest that epistemic modals are gradable. In any event, the point that is relevant to my discussion is that since it's a weaker claim than epistemic necessity, it seems like the pragmatics of *probably*-claims should similarly be such that they relax the felicity conditions on any futurates that they embed.

I don't think so, because the sentences are easily made better with *will*. Let's look at the examples with the accomplishment prejacents, which were uniformly bad without *will*.

(135) **Accomplishments**:[54]

a. John will probably paint a picture.
b. John will probably make a chair.
c. John will probably deliver a sermon.
d. John will probably draw a circle.
e. John will probably recover from an illness.
f. John will probably build a house.

So why is it the case that, when an epistemic possibility modal like *probably* is used, as opposed to an auxiliary like *might* or *must*, that *will* is preferred to the futurate when the complement is eventive? I think the answer to this is a simple one. *Will* is an auxiliary, and in English (some idiolects aside), multiple auxiliaries tend to be ungrammatical.[55] So, when *might* or *must* is used, it's ungrammatical for *will* to appear, and the futurate, not being an auxiliary, is the only grammatical option. And when adverbs like *possibly* and *probably* are used, this allows future orientation of the modal to be expressed will the auxiliary *will*, since it is not precluded by the presence of another auxiliary.[56] To

[54]Interestingly, surface order aside, the interpretation of the sentences in (135) all have the following scope: PROBABLY>WILL.

[55]Apparent exceptions usually involve semi-modals, which don't behave exactly like auxiliaries. So, *John might have to go to the store later* is OK, but this is because *have to* isn't a genuine auxiliary. Since *will* is a genuine auxiliary, it can't be stacked with a modal auxiliary.

[56]Famously, in the English spoken in the Southern United States, multiple modal constructions are attested, in particular configurations like MIGHT>COULD or MIGHT>SHOULD. My prediction, which I'm not currently position to test, is that in idiolects that allow multiple modal constructions, you would be more likely to find a configuration like MIGHT>WILL. I can offer a bit of preliminary evidence corroborating this suspicion, the MultiMo database of Multiple Modals at the University of South Carolina includes several entries with attested instances of MIGHT>WILL; for example *I think I'm going to go home. I might will feel a little better if I lay down.*, *Do you think all these potatoes will fit into this pan? They might will.*, and *This is the worst I've ever done. I honestly think I might will fail.*. Cf. http://artsandsciences.sc.edu/multimo/table Last accessed August 11, 2020.

summarize, the case of lexical modals seems to support the claim that the EC holds of modals generally.

CHAPTER 5

Ought and Agency

5.1. Introduction

In this final chapter, I detail how the machinery developed in the previous chapters can be put to use to provide a novel solution to a puzzle that has long exorcised moral philosophers and meta-ethicists. The word *ought* can be used to express a number of different senses (or "flavors," as semanticists tend to say) – epistemic, bouletic, teleological, deontic, and perhaps others besides. A thorny question surrounding the meaning of *ought* that has long been commented on by moral philosophers concerns a felt distinction between deontic uses of *ought*. One the one hand, there's the *ought* that seems to evaluate a state of affairs, and on the other, the *ought* that describes a requirement or obligation to perform an action.[1,2] The names given to this phenomenon have been diverse[3] and philosophers sometimes train their attention on different features associated with the phenomenon. But the action-enjoining and state of affairs-evaluating distinction

[1]Such a distinction traces back at least to Sidgwick 1874 and Prichard 1912, and is discussed in the early literature on deontic logic (e.g., in von Wright 1951) but for more recent discussion cf. Feldman 1986, Grice 2001, Horty 2001, Broome 2013.

[2]For the view that the felt distinction is merely apparent, cf. Chisholm 1964, Williams 1981, Ch 9.

[3]Here's a sample of the names given to a distinction in the conceptual neighborhood of the one drawn above: agentive vs. non-agentive (Chrisman 2015), deliberative vs. evaluative (Williams 1981; Schroeder 2011), relative vs. non-relative (Grice 2001) ought-to-do vs ought-to-be (Feldman 1986). This terminology doesn't capture exactly the same conceptual distinction. Broome [2013]'s distinction between owned and un-owned *oughts* is also in the conceptual neighborhood, but for Broome, control over one's action is not an essential ingredient of this *ought*, so the link to agency is not as direct as in other philosophers' conceptualization. There may very well be more than one distinction in this conceptual neighborhood. Cf. Humberstone 1991 for discussion.

is pretty consistent across characterizations of the phenomenon, as the quotations below make clear. (The emphasis is mine.)

> Harman [1973]: "In one use *ought* represents a predicate of the **possible state of affairs**... In another use, *ought* represents a relation between **an agent and a possible course of action**."
>
> Geach [1982]: "In the symbolism of von Wright's original article 'Deontic Logic' the operators 'O' and 'P' for obligation and permission are attached, not to propositional letters, but to letters which stand in for general terms, and answer to **kinds of actions**. [...] [O]bligation essentially relates to an agent, it is somebody's obligation; if instead we try to think of the ought-to-be-ness... of a situation involving the agent, then our thinking is going to be confused..."
>
> Schroeder [2011]: "['O]ught' often expresses **a relation between agents and actions** – the relation that obtains between an agent and an action when that action is what that agent ought to do. [...] 'ought' also has an evaluative sense, on which it means, roughly, that were things ideal, some proposition would be the case."
>
> Chrisman [2015, p. 125]: "There is one sense of 'ought' that is conceptually tied up with someone's agency and is related to practical deliberations and responsibility; it is the sense of 'ought' that implies the 'can' having to do with the kind of **control we have over our own actions**. Then there is another sense of 'ought' having a different semantic function; it **evaluates (possible) states of affairs** along various dimensions."

Examples tend to draw out the contrast vividly and evoke the felt distinction without much set-up. In line with the rough characterizations set out above, the sentences in (136) seem to concern an agent's performing of some action, and in (137) the evaluation of some state of affairs.[4]

(136) Agentive:

[4]The sentences are labeled according to the philosopher who supplied the example.

a. You ought to keep that promise. (Harman)
b. John ought to beat up Tom. (Geach)
c. Jay ought to give up smoking (Chrisman)
d. Alison ought to get a sun hat. (Broome)

(137) Non-agentive:

a. Tom ought to be beaten up by John. (Geach)
b. Luckless Larry ought to win the lottery (Schroeder/ Chrisman)
c. There ought not be childhood death and disease (Chrisman)
d. Milton, you ought to be living at this hour. (Wedgwood/ Chrisman, channeling Wordsworth)
e. The meeting ought to start at noon. (Schroeder)
f. Alex ought to get a severe punishment (Broome)

Informants tend to recognize a difference in these examples, and philosophers easily generate them whether or not they think there is any deep cause of these interpretive differences.[5]

Part of the debate surrounding this contrast has turned on how properly to account for it. Those philosophers who deny that there is a distinction, don't feel the need to provide a robust account of the distinction. But among philosophers who think there is some difference between the (136)-sentences and the (137)-sentences – agentive and non-agentive *ought*-sentences, as I will call them – there is little agreement. Some claim that there is a difference in the logical form underlying the sentences, and leverage this difference in logical form into an account that explains the perceived contrast.[6] The

[5]Williams [1981], for example, denies that there is a distinction in these sentences, though he has no trouble producing examples showcasing the different interpretations. Though, as I understand it (from Broome 2012), Williams came to reject the conclusion drawn in this paper in a subsequent, but unpublished, lecture entitled "*Ought*, *must*, and the needs of morality".

[6]In calling them "agentive", I purposely avoided the term "agential" *ought*-sentences, even though this appears in the literature. (Cf. Finlay and Snedegar 2014) This term is typically used to refer to the surface form of the (136)-sentences, differentiating sentences of the form 'S ought to ϕ' where from ϕ is some action, from sentences not of this form, like (137c). Though surface form will be important in what follows, what I want to denote are the deliberative interpretations of *ought*-sentences, agential

difference-in-logical-form approach comes in two flavors – depending on whether it locates the difference in an ambiguity associated with the complement on which *ought* operates, or in the very meaning of *ought* itself. The complement-ambiguity strategy is most often associated with the Agency-in-the-Prejacent hypothesis ("AIP") advanced by proponents of stit-logics like Horty and Belnap 1995, Horty 2001, or Belnap et al. 2001.

Other philosophers have embraced the *ought*-ambiguity strategy. In the quote above, Geach [1982] proposes that agentive *ought* is a distinctive sense of *ought* which does not operate on propositions, but on actions. Schroeder [2011] gives a more modern, linguistically nuanced version of this strategy, where the distinction between agentive and non-agentive *ought* is said to correlate with some syntactic facts. Specifically, agentive *ought* is said to behave like a control verb, "controlling" the arguments of the subordinate verb, and non-agentive *ought* is said to behave like a raising verb, have a control sense of *ought* and a raising sense of *ought*, where the surface argument of *ought* is semantically only an argument of the embedded verb.

But if *ought* is a kind of modal, then the standard account of modals in formal semantics militates against the *ought*-ambiguity strategy. It's widely thought that modals contribute a *uniform* meaning to the sentences containing them, and that differences in meaning the modal utterances exhibit is rather due to a difference in the values given to the parameters that form part of the modal's uniform "kernel" of meaning.[7] In order to

---

or otherwise, so I'm employing the term 'agentive' for this. In doing so, I adopt the terminology used by Chrisman [2015] (eg., cf. p. 124), and intend to mean the same thing Schroeder [2011] means by 'deliberative *ought*'.

[7]Cf. the work of Angela Kratzer – in particular, the papers in Kratzer [2012].

satisfy the apparent desideratum of keeping *ought* uniform, opponents of the difference-in-logical-form strategy propose accounts that keep the logical form of the relevant *ought*-sentence constant across agentive and non-agentive uses, and account for the difference in meaning in some other way.[8]

This gives is the following taxonomy of positions on agentive *ought*.

(138) a. **Denial**: There is in fact no difference between agentive and non-agentive *ought*.[9] (Williams, Chisholm [1964])

b. ***Ought*-ambiguity**: The agentive/ non-agentive difference is due to an ambiguity in the word *ought*. (Schroeder, Geach, Harman)

c. **Complement-ambiguity**: The agentive/ non-agentive difference is due to an ambiguity in the the complement of *ought*. (Horty, Belnap)

d. **Thorough Uniformity**: The agentive/ non-agentive difference is due differences that do not owe to the logical form of the sentence – so, neither to an ambiguity in *ought* nor to an ambiguity in its complement. (Chrisman, Wedgwood)

Two questions emerge as lynchpins in this debate. On the one hand, there is the question of logical form – whether or not there is a difference in the logical form of agentive and non-agentive *ought*-sentences. **Denial** and **Thorough Uniformity** fall on one side of this question, the remaining views on this other. On the other hand, there is the question of whether or not *ought* itself is ambiguous. ***Ought*-ambiguity** stands alone in claiming that it is.

---

[8]For example, Wedgwood [2006] posits an agent parameter in the index. If this parameter is given a value, the resulting *ought*-sentence is interpreted relative to this agent and results in an agentive interpretation. Chrisman [2015] invokes Castaneda [1975]'s distinction between propositions and practitions, arguing that agentive *ought* takes a practitional argument, whereas non-agentive *ought* takes a propositional argument. This may seem like a form of the complement-ambiguity strategy, but for Chrisman propositional and practitional arguments are of the same logical type, so from the perspective of the logical form of the sentence, agentive and non-agentive sentences apparently look identical.

[9]Thus construed, **Denial** is not just a position about the logical form of *ought*, but a deflationary position on agentive *ought* in general.

The primary aim of my paper is to give a novel account of the difference between agentive and non-agentive *ought*. My account will be a variant of **Complement-ambiguity**, but one that is not committed to AIP. This represents a departure from the literature because proponents of the other strategies have treated AIP as the only game in the **Compliment-Ambiguity** town. A few insights emerge as consequences of my account, which have been under-appreciated by the extant positions as canvassed above. First of all, my proposal falls squarely within the Kratzerian paradigm for the semantics of modals. So, my account will show that the Kratzerian position should not so quickly be associated with **Denial** as it has been, in particular by proponents of ***Ought*-ambiguity**. Second, it will show that **Complement-ambiguity** has considerably more resources at its disposal than has been appreciated by its opponents (and probably also by its proponents). And, finally, that once these resources are appreciated, positing a difference in logical form as an explanation for agentive *ought* won't thereby commit one to ***Ought*-ambiguity** if one already rejects AIP, as proponents of **Thorough Uniformity** seem to hold.

Here's the plan: §5.2 motivates the distinction between agentive and non-agentive *ought* in more detail. §5.3 sketches an informal version of the proposal. §5.4 describes some linguistic patterns an explanation should be sensitive to. §5.5 explains some problems with the erstwhile most prominent version of **Complement ambiguity**, AIP, but argues that some insights from this strategy are worth retaining. §5.6 lays out the proposal – the coercion theoretic account of agentive *ought* – and its motivations. §5.7 and §5.8 elaborate on some of the account's details, and §5.9 concludes by discussing the payoffs of the coercion theoretic account.

5.2. Agentivity motivated

Aside from the intuitive distinction appealed to in the introduction, why think there is any kind of special difference between agentive and non-agentive *oughts*? Here are additional reasons to think so.

Agency Sensitivity: Philosophers have pointed out that there's an intuitive difference between S's performing an action and S merely being a participant in an event where an action-like behavior was performed. The second kind of circumstance can obtain even if the event is one we'd unreflectively think of as a kind of action, but one's performance isn't a genuine exercise of one's agency. Consider the sentence in (139).

(139) Tom ought to kiss Bill.

As pointed out by Chrisman [2015, p. 115], there's an intuitive difference between (i) someone bringing it about that Tom kisses Bill by spiking his drink so that it foreseeably leads to his kissing Bill in his drunken haze, and (ii) Tom's just performing the action of kissing Bill "directly and actively". The difference can perhaps be illustrated even more explicitly by considering the role of deviant causal chains. We might say that Tom's proper exercise of his agency involves forming an intention and having this intention be the cause of his kissing Bill. Suppose that Tom resolves to kiss Bill, and that his having settled on this intention makes him so nervous that he stumbles forward. Rather than leaning in and deliberately planting his lips on Bill's, he lands mouth-first onto Bill. In this deviant causal chain, Tom winds up obtaining the intended result – his kissing Tom – but this outcome comes about through some "deviant" means rather than the paradigmatic way an agent's intention leads to action.

We don't need to delve deeper into action theory to see the challenge this poses to a standard, Kratzer style view of *ought*. If *ought* operates on propositions, and these are construed as sets of possible worlds, then the proposition *that Bill kisses Tom* is just the set of worlds where Bill kisses Tom. This set of worlds is the same regardless of whether Bill winds up kissing Tom in the "normal" way, as an expression of his agency, or in the "deviant" way, where his nerves get the better of him. It would seem that construing *ought* as a propositional operator like this is not sufficiently **Agency Sensitive**.

Argument Asymmetry: Now, suppose Tom and Bill are together attending a dance of some sort. Tom has been neglecting Bill all night, having failed to dance with him when he promised to do so. He really owes Bill an apology, and should finally dance with him while he still has the opportunity. Bill, by contrast, has suffered this indignation quietly – he has done nothing wrong. Now consider (140a) and (140b).

(140) a. Tom ought to dance with Bill.

b. Bill ought to dance with Tom.

Under the circumstances just described, it seems to many people that (140a) and (140b) are subtly different. (140a) is true, but (140b) doesn't seem to be. Tom owes Bill a dance; Bill doesn't owe Tom anything (except perhaps an earful). To say he's got anything like an obligation to dance with Tom under the circumstances seems perverse. By contrast, Tom owes it to Bill to dance with him.

The *dance-with* relation is symmetrical, and the state of affairs where Tom dances with Bill is co-extensive with the state of affairs where Bill dances with Tom. So if *ought* just evaluates states of affairs, the thought goes, there's no reason to treat these (140a) and (140b) as any different. But if *ought* can single out the agent of the described action

as the one to whom the obligation adheres, we get a purchase on the distinction between (140a) and (140b). And this **argument asymmetry**, as I call it, is naturally explained by the fact that as far as the *ought* sentence is concerned, one of the arguments seems to be singled out as being the bearer of the obligation and the agent of the obliged action. And again, if *ought* operates on propositions, and these are construed as sets of possible worlds, then the proposition *that Tom dances with Bill* is the same as the proposition *that Bill dances with Tom*, at least on standard ways of individuating propositions.[10]

5.3. An Informal Picture

As I pointed out in section 5.1, a tendency in the literature takes the Kratzerian paradigm as committed to **Denial**.[11] Though Kratzer hasn't weighed in on this particular issue, it is easy to see how a Kratzerian can provide the semantic underpinnings to **Denial**. The standard semantics has it instead that modals like *ought* uniformly operate on sentence-like expressions. And at first blush, it's hard to see how the agentive/ non-agentive distinction can be made tractable given such constraints. One of the motivations behind the proposal I will articulate is that the essential ingredients to Kratzer semantics need not be enlisted to serve the deflationary function of **Denial**. Once this is understood, the commitment to a uniform, non-ambiguous meaning of *ought* doesn't preclude the possibility of a difference in logical form for explaining agentive *ought*.

To give the flavor of what I have in mind for my proposal, consider what Recanati [2004] calls variadic functions. These are functions from relations to relations which

[10]This is also the thought that animates the discussion of passive transformations of sentences. In the circumstances described above, we tend to think (139) can be true without the passive transformation of (139) – *Bill ought to be kissed by Tom* – thereby being true.

[11]This tendency is not universal, though. Cf. Bronfman and Dowell 2018 and Dunaway 2017.

increase or decrease the adicity of the relation. So, adding a predicate modifier to a predicate expressing an n-ary relation results in a predicate expressing an $n + 1$-ary relation. Recanati calls upon veriadic functions to explain a certain kind of phenomenon, where the truth-functional content of a sentence includes components that aren't overtly linguistically specified.[12] For example, consider sentence (141).

(141) It's raining.

In (141), there's no location specified, yet hearers typically understand the sentence as concerning rain *at a specific location* – the location of the utterer, say. To get the sentence to reflect the truth conditions which include the the location, the veriadic function operator $\text{OP}_{location}$ will map the RAINING relation to the RAINING-AT relation by increasing the adicity of the RAINING relation by 1, thereby providing an argument position for a location (where l is the location of the raining event).

(142) $\text{OP}_{location}$ (Raining) = Raining-at(l),

A schematic first pass for what I have in mind can be put as follows. The difference between agentive and non-agentive *ought* sentences is that former have a kind of operator which allows for an additional argument position to be added to the clause which is related to the modal in a distinguished way. Consider (143).

(143) OUGHT [$\text{OP}_{\text{AGENT}=x}$ (ϕ)]

In (143), OUGHT can have a standard modal interpretation, largely along the lines proposed by Kratzer.[13] ϕ is clausal, so it is the kind of sentence-like expression required by

[12]Readers familiar with the literature about the semantics-pragmatics interface will recognize this as part of the debate surrounding unarticulated constituents. (Cf. Perry 1986)

[13]In fact, the proposal I will advance will require a slight revision of Kratzer's early formulation of her semantics (cf. section 5.6.2), but the main revision is one Kratzer herself endorses. Cf. Kratzer 2013.

Kratzer's semantics. And yet $\text{OP}_{\text{AGENT}=x}$ is a variadic function which takes the 0-place relation ϕ and turns it into a 1-place relation that can take an agent argument. Because the source of the agentive interpretation lies with a kind of ambiguity in the modal's complement, the proposal is a form of the **Complement-ambiguity** strategy.

If such an account were well-motivated, it would boost the linguistic plausibility of the complement-ambiguity strategy. However, it might seem like an uphill battle. Recanati [2004]'s veriadic functions are controversial, and enlisting them specifically to explain agentive *oughts* risks looking hopelessly *ad hoc*. Typically, they are thought to rely on a process of pragmatic enrichment, whereby the linguistic expressions associated with our utterances are pragmatically enriched with supplemental material before they are truth-conditionally evaluated. Proponents of pragmatic enrichment usually appeal to various sorts of evidence for the existence of these kinds of pragmatic processes. I haven't give any such account. So, to say that $\text{OP}_{\text{AGENT}=x}$ is a veriadic function in this sense is simply to say that, somehow, the black-box of pragmatic enrichment takes a normal *ought* sentence and gives us an agentive *ought* sentence.

Here is where my proposal diverges from the kind of veriadic functions envisioned by Recanati. As it turns out, I won't need to appeal to pragmatic enrichment to justify my proposal. I will argue that $\text{OP}_{\text{AGENT}=x}$ is actually *grammatically* motivated by several factors relevant to the interpretation of the sentence wherein it occurs.[14] The bulk of my arguments will therefore focus on two aims: first, making the case that the schema represented in (143) is linguistically well-motivated, and second, showing how the result gives us a good representation of agentive *ought*.

[14]The proposal therefore has more in common with the kind of contextualism proposed by Stanley 2000 and Stanley and Szabó 2000 than with enrichment accounts of Recanati.

5.4. Some Linguistic Generalizations

Before training specifically on the contribution *ought* makes to the meaning of the respective sentences, a careful observer may note some emerging patterns. The (agentive) (136)-sentences all share two grammatical and semantic features. First, they all contain verbs that describe an *action* some relevant agent (usually the subject of the sentence) should perform – they all contain action verbs. From a linguistic standpoint, what is noteworthy is that all the verbs describing these putative actions, are *eventive* verbs, which have a different temporal profile than stative verbs.[15] Secondly, all of the (136)-sentences are *future oriented* with respect to the time of evaluation.[16] This means that the event described by the complement of the modal is in the future with respect to the time of its evaluation.

By contrast, the non-agentive (137)-sentences are more varied on both counts. They tend *not* to describe actions. Sometimes they describe states of affairs, like (137c) and (137d), in which case the verbal predicates contained in the prejacents are stative. Other times they describe events, but not events the subject can properly be said to be the agent of. For example, (137a) is in the passive voice, and Tom is certainly not the agent of his being beaten up. While not grammatically passive, (137f) is passive in a certain sense, since the complement describes Alex being on the receiving end of some punishment doled out. They are also more varied in terms of their temporal orientation. They can be future oriented, as a number of them are (like (137b), (137e) or (137f)), but they can also be

[15]Some quick and dirty linguistic distinctions between eventive and stative verbs: Among other things, eventive verbs are unable to appear in the simple present tense without a habitual reading, but stative verbs are fine in the simple present. Eventive verbs describe events which *happen* or take time to occur; stative verbs describe states that simply hold. Eventive verbs can appear in the progressive, whereas statives tend to resist appearing in the progressive.

[16]The term "temporal orientation" originates with Condoravdi 2002. Cf. also Matthewson 2012.

present oriented, as (137c) and (137d) are. As non-agentive, evaluative *ought* sentences, they rank the state of Milton's being alive, or the state of childhood disease not obtaining, as being preferable to the alternatives *now*.[17]

When considering the putative meaning of agentive *ought*-sentences, neither of these observations will be too surprising. But the linguistic generalizations these observations suggest are noteworthy. The paradigmatic sentences that evoke the agentive interpretation of *ought* exhibit a grammatical pattern. We can summarize this pattern as follows.[18]

> Agentive *ought*-sentences (i) tend to have eventive complements, and (ii) are always future-oriented.
> Non-agentive *ought*-sentences (i) can embed stative *or* eventive complements, but we tend not to see complements that describe an intentional action. (ii) They are variously future or present oriented.

Something about the eventivity of *ought*'s verbal complements lends them to be easily identified with an agentive interpretation of *ought*. And this tendency is so strong, in fact, that we really don't need much stage-setting or context-describing to allow one to view it as an agentive *ought*-sentence. An account of agentive *ought* sentences should be sensitive to these patterns, so we'll return to it.

---

[17]There is a linguistic reason for this. In fact, eventivity and future orientation are related. Unless they are interpreted habitually, eventive predicates are obligatorily future oriented in a number of constructions – whether under modals, or in infinitival complements in general. (For example in attitude ascriptions. Compare: *Tom wants to be at the party* (present orientation) vs. *Tom wants to go to the party* (future orientation).) Stative predicates can be either future- or present-oriented. Insofar as the predicates in the (136)-sentences describe actions, they are eventive predicates. Given the obligatory future-orientation of eventive verbs, it is predictable that they would be future-oriented.

[18]Finlay and Snedegar [2014] call these paradigmatic agentive *ought*-sentences "agential". They mean to characterize the kind of grammatical form agentive *ought* sentences tend to have. I seem to have the same class of sentences in mind as they do, but they don't specify the pattern shared by this class of sentences quite the same way I have in this section.

5.5. Agency-in-the-prejacent theory

In this section, I will discuss the most prominent type of Complement-Ambiguity strategy, AIP, and why it falls short. In addition to summarizing some of the extant criticisms of AIP, I will add some of my own. I do this because I think that AIP contains a good insight – that the agentivity of the targeted reading of the *ought* sentences is a reflex of the way language encodes information about agents. It just exploits this insight in the wrong way. So my motivation in this section is to make explicit exactly how the proposal fails to provide a satisfactory explain of the distinction in order to save the legitimate insight.

AIP is motivated by a class of logics for action called stit logics.[19] Stit logics invoke a logical form that represents an agent as performing an action by specifying the agent, the proposition whose truth the agent is supposed to bring about, and connecting these by means of a stit operator which takes them both as arguments, so $\ulcorner \alpha \ \mathit{stit}\ \Phi \urcorner$ is read 'α sees to it that Φ'. In order to connect the logical form of such formulae with the natural language sentences they are purported to represent, stit logicians have proposed AIP. AIP exploits the linguistic insight that sentences have a way of encoding information about the agent of an action via the lexical semantics of the verbs contained in the sentence. It attempts to capture the notion that some *ought*-sentences of the form *S ought to* ϕ variably have agentive readings by appealing to an ambiguity in the prejacent, *S* ϕ*s*. The sense of agency owes to the fact that the sentence itself contains information about the agency of the participants in the event described by the Verb Phrase. Let's reconsider an example discussed earlier, repeated here as (144).

[19]Horty and Belnap [1995], Belnap et al. [2001], Horty [2001].

(144) Tom ought to kiss Bill.

According to AIP, the reason this sentence has a distinctly agentive sense is that the clause embedded under *ought* contains information pertaining to the agency of the subject, Tom. In verbal semantics, verbs are said to be related to their arguments via thematic roles; in virtue of saturating different argument positions, the arguments encode information pertaining to the object's role in the event or the action described by the verb.

A verb like *kiss* has two argument positions – an agent argument and a theme argument, where the former is the agent performing the action described by the verb and the former is object being acted upon in the event described by the verb.[20] Drawing on this fact allows the advocate of AIP to suggest that there's nothing different in the logical syntax of agentive *ought* as compared to the non-agentive *ought*; it's just that they have have their distinctive agentive sense in virtue of semantic information contained in the prejacent itself. (144) can then be analyzed as (145), which can then map onto the stit formula in (146).[21]

(145) Ought($\text{Tom}_{\text{AGENT}}$ kiss Bill)

(146) Ought [Tom *stit* (Tom kiss Bill)]

Critics of AIP typically point out that this strategy overgenerates. Schroeder [2011] gives the following example.

(147) a. Luckless Larry ought to win the lottery.

[20] *Theme* is sometimes used interchangeably with *patient*, though some theorists differentiate the two on the basis of whether or not the thing acted upon changes its state as a result of the event it undergoes. If we opt for this finer-grained distinction of *theme* versus *patient*, in the sentence *Tom cut the cake*, *the cake* would be the patient, and not the theme, because it's changed its state as a result of the cutting event.

[21] The virtues of stit logics don't rest on the viability of AIP, but AIP does provide a bridge principle for connecting stit logics to natural language.

b. It ought to be the case that Luckless Larry wins the lottery.

According to Schroeder, the natural reading of (147a) is non-agentive – it was one of the paradigm examples of a non-agentive *ought*-sentence in (137). Yet, Schroeder claims, there is a remote, but still accessible reading where (147a) is indeed agentive – think of a context where Larry is able to fix the lottery and ensure that he wins. AIP would explain the availability of the agentive reading of (147a) by appealing to an ambiguity in the prejacent. On the agentive reading, the predicate takes an argument which is theta-marked as AGENT, whereas on the non-agentive reading, it takes an argument not theta-marked as AGENT but instead, say, THEME. So far so good. Now, since (147a) and (147b) both presumably have the same prejacent, namely that *Luckless Larry wins the lottery*, (147b) should have an agentive reading also, but Schroeder denies that such a reading exists. This is troublesome for AIP because there is no apparent difference in the prejacent between (147a) and (147b) – why should there be an ambiguity in the embedded clause in (147a), but not in (147b)? AIP should then predict that there is an agentive reading for (147b) as well, but apparently there is not.

There is some dispute in the literature about whether Schroeder is right that (147b) can never have an agentive interpretation.[22] To my mind, the bigger problem for AIP is actually (147a) – AIP's answer for why (147a) can have an agentive interpretation is linguistically implausible. If (147a) does have an agentive reading, even if it is atypical, it's fairly mysterious on AIP why this is. AIP gains plausibility by hitching its wagon to a view about the semantics of verbs; that their distinctively agentive sense is due to the

[22]Chrisman [2015] and Bronfman and Dowell [2018] deny that (147b) is incapable of receiving an agentive reason. The interpretive issues here are rather delicate, and I discuss them later, so I will defer further discussion of this question until then.

underlying lexical semantics of the verbs and the kinds of arguments they take. It isn't simply a way to stipulate that some prejacents have an agent-y flavor sometimes and other times they do not. AIP's way of explaining the presence or absence of an agentive reading of (147a) will only be as plausible as the claim that the embedded verb can variously have or lack an agent argument. If AIP's explanation for how we get a deliberative reading of *ought* is to be linguistically plausible, it needs to maintain a fairly close connection between the AGENT-theta role and the sense of agency exhibited by the *ought* claim.[23]

This is what makes the agentive reading of (147a) a mysterious affair. Verbal polysemy doesn't work the way AIP needs it to in order to make good on the claim that the prejacent for (147a) is ambiguous with respect to its argument structure. *Win* belongs to a class of verbs – achievements, in the so-called Vendler [1957]/ Dowty [1979] taxonomy for lexical aspect – that typically do not have agent arguments. Rather, these verbs have EXPERIENCER-, PATIENT-, or THEME-roles.[24] If *win* doesn't have an agent argument, whence the agentive reading?

The most obvious response is the one floated by AIP. When the *ought*-sentence has an agentive reading, the embedded verb selects for an agent argument and when it does not have an agentive reading, it does not. So the answer buys modal uniformity at the cost of proposing a form of verbal polysemy. This may seem reasonable – verbal polysemy isn't particularly uncommon. However, the kind of polysemy needed to suit AIP is unlikely in two respects. First, as we will see in the next section, agent arguments are typically

[23]Perhaps there are circumstances where the AGENT-theta role can be de-emphasized, and a sentence with an agent argument in the prejacent has a non-agentive reading. If this is a possibility, we can imagine a non-deliberative reading of (144).

[24]A look at other achievement verbs will bring this point into further relief: *fall*, *die*, *recognize*, *find*. In the events these verbs describe there is certainly a salient person who undergoes the event. But that subject is not typically an agent with respect to that event.

sentential subjects. And sentential subjects are *external arguments*. So *Luckless Larry* is an external argument for *win* whether it has the natural THEME interpretation or the exceptional AGENT interpretation AIP proposes. External arguments are distinguished from *internal* arguments precisely because they tend *not* to influence a shift in meaning of the verb.[25] AIP's polysemy claim amounts to proposing that verbal meaning can be affected by the external argument, but the linguistic evidence seems to point in precisely the opposite direction.

Secondly, the kind of meaning for *win* AIP posits only seems to exist in the scope of the appropriate modal. Let's say that *win** is the word just like *win* except that it has an AGENT argument instead of a THEME argument. AIP effectively posits either a distinct meaning expressed by a new lexical item *win** that we simply don't detect in unembedded positions, but only under *ought*. This is not a good sign for the theory; if *win** only occurs in very specific embedding conditions, the hypothesis to beat would be that the phenomenon we are trying to explain owes to the embedding environment, not to lexical polysemy, which we should expect to observe unembedded.

In summary, AIP attempts to marry the agentivity of *ought* to a linguistic theory about how agentivity is encoded in language. This seems like a promising and well-motivated strategy. However, in explaining agentive *ought* by appealing to an implausible kind of verbal polysemy, AIP runs afoul of some of the data that underwrites this theory.

---

[25]For example, *win the race*, *win at life*, and *win his heart* are plausibly thought to involve different but related senses of *win*. It is precisely the internal arguments (*the race*, *at life*, and *his heart* respectively) that trigger the different senses at issue here. The external argument tends *not* to affect the meaning of the verb in this way. Cf. Bresnan 1982, Grimshaw 1990 and especially Marantz 1984 for this observation.

5.6. A Coercion-theoretic Account of Agentive *Ought*

In section 5.3, I compared my positive proposal to Recanti's veriadic functions. The sketch of the proposal was that agentive *ought* was can be thought of as due to an operation akin to a veriadic function operator as in (143), repeated below as (148).

(148) OUGHT [$\text{OP}_{\text{AGENT}=x}$ (ϕ)]

But I qualified this comparison by saying that the "operator" which triggers an agentive-*ought* is largely a grammatical matter rather than the result of pragmatic enrichment. I can now provide some additional details on what I meant by this. Rather than agentive-*ought* being triggered by a single semantic operator or function, it will actually be the result of the interaction of several independent grammatical phenomena. What allows for the agentive interpretation of *ought* is a kind of coercion operation caused by a mismatch between the semantic properties of the modal's complement clause, and the demands of the modal. Before presenting my proposal, I'll have to motivate this operation. So while it might seem that this section pivots into unrelated territory – how verbs compose with their agent arguments, and the syntactic position of deontic modals – this discussion will help explain the coercion-theoretic mechanism that gives rise to agentive-*ought*. It's a strength of the proposal, I take it, that the mechanism it posits isn't motivated by the desire to explain agentive *ought*, but by independent evidence in the semantics of modals. After suppling this motivation, I give explain the proposal in more explicit detail in 5.6.3.

5.6.1. Agent arguments

In a paper that spawned a rich and productive literature in semantics, Davidson [1967] proposes (i) that verbs are n-ary predicates of events, (ii) that prepositional modifiers and

many adverbial modifiers can be rendered as predicates of events as well, and (iii) that this modification is accomplished via conjunction of the modifier with the verbal predicate, according to the following schema: ⌜∃e[VERB(e,a,b) & MODIFIER(e)]⌝. The verb's arguments are represented as arguments of the verbal predicate VERB, which contains an extra argument position for the event variable, e.[26]

Subsequent "Neo-Davidsonian" work by Higginbotham [1985] and Parsons [1990][27] suggests that Davidson's proposal be combined with the work on thematic frames, such that the verb's arguments are themselves related to the event via a predicate for the thematic role which they occupy. Davidson's proposal and the Neo-Davidsonian proposal on the meanings of verbs are schematized by (149a) and (149b) respectively.[28]

(149) a. Davidson: $\lambda y.\lambda x.\lambda e[\text{VERB}(e,x,y)]$

b. Neo-Davidsonians: $\lambda y.\lambda x.\lambda e[\text{VERB}(e) \;\&\; \textsc{agent}(e,x) \;\&\; \textsc{theme}(e,y)]$

As I mentioned in my criticism of AIP, linguists have long pointed out that there is a systematic asymmetry between the argument realized by the *subject* of the sentence (often called the "external" argument) and those closer to the verb (the "internal" arguments). Kratzer [1996], drawing on an argument by Marantz [1984], points out that there are

[26]This approach allows for successive modification of an event variable, thereby preserving the kinds of entailments between sentences we intuitively associate with such modifications. Recalling a classic group of sentences (from Parsons [1990]), (A1) entails (A2) and (A3) but not vice versa; each of (A2) and (A3) entail (A4) but not vice versa, and (A1) entails (A4).

(A1) Brutus stabbed Caesar with a knife in the forum.
(A2) Brutus stabbed Caesar with a knife.
(A3) Brutus stabbed Caesar in the forum.
(A4) Brutus stabbed Caesar.

[27]cf. also Castañeda [1967]'s and Lemmon [1967]'s proposals.

[28]A caveat: ultimately these formulae – which I present as abstracted over – will be existentially bound. But I skirt over how exactly this happens. Such an explanation is not essential to our purposes.

many cases where an internal argument triggers a particular interpretation of the verb.[29] By contrast, there are vanishingly few cases where an external argument has this kind of effect. The examples in (150) and (151) showcase how differences in an internal argument can trigger a different meaning of the verb.

(150) a. throw a baseball
b. throw support behind a candidate
c. throw a boxing match (ie., take a dive)
d. throw a party
e. throw a fit

(151) a. kill a cockroach
b. kill a conversation
c. kill an evening watching TV
d. kill a bottle (i.e., empty it)
e. kill an audience (i.e., wow them)

Note that the meaning of *throw* and *kill* varies depending on the nature of the internal argument. This is not the case for *external* arguments, which are said *not* to affect the meaning of the verb.

On this basis on these and similar kinds of data, Kratzer proposes to "sever" the external argument from the verb, suggesting that the external argument slot is not properly part of the lexical meaning of the verb. This construes verbal meaning as in (152), where the variable y is an internal argument of the verb.

(152) Kratzer: λy.λe[VERB(e,y)]

The distinction between (152) and (149) is not merely notational, because this proposal has substantial consequences for the semantics of verbal meaning. For example, a verbal

[29]Cf. also Bresnan 1982 and Grimshaw 1990 for this observation, though they ultimately explain it in a different way than Marantz and Kratzer.

predicate must have an external argument in order to be a grammatical sentence. So, it means the external argument must enter into the semantic representation through a special grammatical process. According to Kratzer's proposal, this happens in the projection called the Voice Phrase ("VoiceP"). To see how this works, let's take the example *Brutus stabbed Caesar*. Using the schema in (152) as a template, and abstracting away from tense, the present proposal would mean that Voice takes λe.STAB(e, c) and λe.AGENT(e,b) and conjoins them as λe.[AGENT(e,b) & STAB(e, c)].

This gloss overlooks a bit of complexity, which is worth pausing over because it's important for our purposes. The first point is that the external argument enters the semantic representation through a special operation. Simply conjoining λe.STAB(e, c) and λe.AGENT(e,b) would not ensure that the event variable, which the respective predicates are the predicates of, denotes the same eventuality. A special semantic operation Kratzer calls *event identification* takes the two predicates of events and conjoins them – but in so doing, the operation ensures that the two event variables are identified in the computation.[30]

The second point is that the lexical verb and its internal arguments can constrain the kind of external argument it receives. The internal argument determines the shades

[30]In more technical terms, the need for a principle like event identification can be explained in terms of a type mismatch. Voice will first take the verb, then compose it with $\lambda x \lambda$e.AGENT(e,x), at which point it can get *Brutus* can fill the argument slot for the agent. So, in terms of the type theory, we have an expression of type <e,<s,t>> composing with an expression of type <s,t> and yielding an expression of type <e,<s,t>>:

(1) λxλe[AGENT(e,x)]$_{<e,<s,t>>}$, λe[STAB(e,c)]$_{<s,t>}$ $\rightarrow$ λxλe[AGENT(e,x) & STAB(e,c)]$_{<e,<s,t>>}$

If the only composition principle were function application, these expressions wouldn't compose because of a type-mismatch. But the framework Kratzer is supposing here takes function application to be the mode of composition of lexical arguments, and this process involves composition of non-lexical, functional items, which can have a composition principle other than mere function application.

of meaning the verb can have, such that *kill* can variably have the kind of meaning it exhibits in *kill a conversation* versus *kill a cockroach*, based on the kind of internal argument it has. But once the verbal predicate has this meaning, it can limit the kind of external argument that can compose with it. For example, *kill a conversation* needn't have an external argument denoting an agent. A misplaced, awkward joke can well kill a conversation. As an external argument *the awkward joke* wouldn't be an AGENT argument, but is instead an INSTRUMENT argument. By contrast, a person is well suited to kill a cockroach (with the heel of a shoe, perhaps), so the predicate expression *kill a cockroach* can easily receive an AGENT argument. This is an important point, because some verbs (like stative verbs, and as mentioned in the last section, achievements) do not have an AGENT argument to begin with. So, the state denoted by the predicate *own a dog* will not have an AGENT argument, but an EXPERIENCER argument.[31]

The points to emphasize from this discussion are as follows. First, the external argument composes with the verb via a special process called event identification, which will conjunctively add an agent argument to the logical form. Second, the kind of external argument that is added will be constrained by properties of the event described by the lexical verb and its internal arguments.

[31]There is some variation on the literature on precisely what thematic roles there are and how they are delineated. In her discussion of this point, Kratzer describes states as having a "holder" external argument. This variability doesn't affect the overall point, which is that the resulting predicate meaning can constrain the kind of argument that can compose with the verb through event identification.

5.6.2. Modals

Now, some discussion of modals.[32] If we treat *ought* as a modal along the lines suggested here, then we need to understand the semantics of modals. On Kratzer's canonical semantics, *ought* is treated much like the $\Box$ of modal logic, as a quantifier over sets of possible worlds.[33] One of Kratzer's innovations was to say that modals are interpreted relative to two kinds of *conversational backgrounds* – the modal base f and the ordering source g, which jointly restrict the domain of quantification. We can schematize the proposal for an *ought*-sentences as follows.

(153) OUGHT $(\mathrm{D})^{f(w),g(w)}$ ϕ

D is the domain, which is determined through context's selecting a modal base f, an accessibility relation on the world of evaluation w. Context also selects an ordering source g. From the worlds delivered by f, g then selects the highest ranked worlds among those. D is the set of the best ranked (according to g) worlds selected by f. The truth conditions are as follows: (153) is true in w just in case the best g-ranked f-worlds in D are ϕ-worlds.

For deontic modals, the modal base picks out a set of propositions characterizing the circumstances in w relevant to evaluating the modal, and the ordering source picks out a set of propositions which determine the relevant ideal or law (also in w). Evaluating (154), for example, would work as follows.

[32]To make clear how I intend on applying the machinery from the previous chapters to the account of agentive interpretations of modals, I will rehearse some of the relevant bits of earlier chapters, albeit in abbreviated form. This section draws on work from sections 1.5 and 3.6.

[33]Caveat: I ignore entirely here the fact that *ought* is often said to be "weaker" than other necessity modals, like *must*. This observation was made by Horn [1972]. To the extent that this weakness owes to a semantic property of *ought*, it is sometimes said that the domain of quantification is more restricted than for strong necessity modals like *must*. For more discussion, cf. Ninan 2005, von Fintel and Iatridou 2008. This is an interesting and perplexing difference between *ought*/*should* and *must*, but does not have a bearing on the present discussion, so I put it aside.

(154) Jay ought to give up smoking.

The modal base would pick out a set of propositions characterizing the relevant circumstances; *that Jay is a smoker*, *that smoking is expensive*, *that smoking causes cancer*, etc. Since propositions are themselves sets of worlds, intersecting each of these sets collects the worlds characterizing the conjunction of these propositions in a single set; the set of worlds consistent with the relevant circumstances. The ordering source picks out a set of propositions which characterize the relevant priorities; *that Jay do what he can to remain healthy*, *that Jay doesn't spend money frivolously*, etc, and it then ranks the worlds in the domain based on which of the priorities they are able to meet. Given these priorities, and holding fast these circumstances, the giving-up-smoking-worlds are ranked higher by the ordering source than the continuing-to-smoke worlds. (154) is then true just in case every such giving-up-smoking world is ranked higher than a continuing-to-smoke world.

On the face of it, we have a straight-forward application of Kratzer's semantics to the case of *ought*. But the semantic account just described abstracted away from linguistic phenomena like tense and aspect, and there are thorny questions about the interaction of modals with these other clausal elements. Chief among these is the apparent syntactic fact that epistemic modals typically sit high in the syntactic representation of the sentence, scoping above tense and aspect, and root modals like deontics tend to scope low, below aspect.[34] Accommodating this different syntactic behavior is challenging if one wants to maintain a *uniform* lexical entry for modals according to which they contribute the same linguistic meaning to the sentences containing them.[35] In terms of the schema

[34]Cf. Cinque [1999], Drubig [2001], Jackendoff 1972, Stowell [2004], etc.

[35]And recall, one of the major motivations of opponents of *ought*-ambiguity is to respect the contention, common in linguistics, that modals *are* uniform.

(153) above, epistemic modals differ from root modals like deontics due to a difference in D – the sets of worlds they quantify over are have different properties. Since the f parameter is responsible for the accessibility relation – g just being a ranking device – it is clear that epistemics and roots will differ by the choice of modal base. Root modals have a circumstantial modal base, which is characterized, as we saw above, by a set of circumstances obtaining at a certain time. Epistemic modals have an epistemic modal base, which is characterized by an information state – the state of what is known by an agent at a certain time, for example.

The work of Valentine Hacquard showed how modals can have a uniform semantics even in light of the height difference.[36] This is done by minimally re-configuring the conversational backgrounds. Notice in (153) how both f and g were functions from worlds to sets of worlds. Hacquard proposes to keep Kratzer's overall framework, but now f and g are functions from *events* to sets of worlds. This further integrates modal semantics with Davidsonian event-semantics, and maintains the uniform framework for modals advocated by Kratzer.

Since root modals sit "low" in the clause, closer to the verb, the modal base can take the VP's event argument as an argument instead of the world of evaluation.[37] The truth conditions for (154) remain more or less the same, except for the subtly different

[36]Cf. especially her 2006 dissertation and subsequent work, notably Hacquard 2010.

[37]Epistemic modals, sit "high" enough in the clause that the VP event is closed and is no longer "available". The modals instead take the utterance or illocutionary event as an argument. This allows Hacquard to give a principled explanation for why the height difference correlates with the difference in modal bases. Root modals are keyed to VP events and when f takes such an event as an argument, it yields a set of circumstances. Epistemic modals are keyed to events that have content – like illocutionary events or attitude events – and when f takes such an event as an argument it yields an information state. So, it's the kind of argument that determines the value of function; it's not that the modals themselves require different types of parameters.

way D is calculated, thanks to the re-configured modal base parameter reflected in (155), which replaces (153). For specificity's sake, (156) presents a side-by-side comparison of how Kratzer's and Hacquard's semantics determines the set of worlds that will ultimately determine D.

(155) OUGHT $(\mathrm{D})^{f(e),g(e)}$ ϕ

(156) a. $\bigcap f_{circ}(w) = \{w' |\ w'$ is compatible with the c-relevant circumstances in $w\}$

b. $\bigcap f_{circ}(e) = \{w' |\ w'$ is compatible with the c-relevant circumstances of $e\}$

But the schema in (155) is yet incomplete – according to Hacquard's analysis root modals scope below tense and aspect. To see what's at stake here, let's fill out (155) with a toy semantics for both of these. Aspect is standardly thought to turn largely on the perfective/ imperfective contrast, and concerns the way the grammar presents or packages the temporal viewpoint of the event described by the sentence.[38] Let's use the semantics for imperfective suggested by Kratzer [1996] according to which imperfective is characterized as follows: $\lambda t.\lambda e.[t \subseteq \tau(e)]$.[39] For present tense, we'll say that the truth conditions are defined if $t = t_u$ – that is, if t is identified with the time of utterance. So we update (155) with (157).

(157) Defined if $t = t_u$

$\exists$e $[t \subseteq \tau(e)$ & OUGHT $(\mathrm{D})^{f(e),g(e)}$ ϕ(e)$]$

[38]To draw on a common metaphor for explaining aspect, perfective aspect presents the event "from the outside," as a completed event, whereas imperfective aspect presents it "from the inside," as an ongoing and incompleted event. Cf. Comrie [1976] for an introduction to these notions and an explanation of the spatial metaphor usually invoked to explain them.

[39]τ, the temporal trace function, is a function from events to their run-times, or the interval throughout which they occur. Kratzer's proposal here encodes the "from the inside" metaphor for imperfective aspect by saying that the run-time of the event is included with a time t, which is the input to tense. If in present tense, it amounts to saying that the event is included in the time of utterance.

Two problems come into relief here.[40] Let's call them the Temporal Orientation Problem and the Event Identification Problem respectively. First, the Temporal Orientation Problem, which is this. Since e is both the variable of the event description *and* the input to aspect, the schema (157) would give truth conditions that make ϕ event – the event describing the action the subject is supposed to undertake – contemporaneous with the speech event. In a word, it would predict that the interpretation of (157) is *present* oriented instead of *future* orientated. Now for the Event Identification Problem. As indicated in (156b), the modal base uses e to calculate the worlds in D. But (156b) requires that the worlds given by the modal base be consistent with the circumstances of e. For *many* instances *ought* sentences, this just will not be the case. In these cases, the circumstances according to which *ought* is evaluated are what call for ϕ, and they can be incompatible with ϕ, which is precisely why we need to change them.

Our stock example in (154) can help illustrate both of these problems. Intuitively, (154) is future oriented. But the schema in (157) would have it that the utterance time contains the quitting-smoking event. So our initial schema does not capture the future orientation of (154) – worse still, it would represent it as present oriented. This is the Orientation Problem. Next, recall the circumstances that characterized the modal base in our original discussion of the example. The circumstances included *that Jay is a smoker*, *that smoking is expensive*, *that smoking causes cancer*, and other propositions like this. (157) says that the worlds in this set are those consistent with e, and in this case, e is the event of Jay giving up smoking. But the worlds at which the circumstances are such that Jay gives up smoking are *not* worlds where those same circumstances are that Jay

[40]The problems, and my solution for avoiding them, were introduced in section 4.1. I explain all of these briefly, and then explain how they figure in an account of agentive modals.

smokes.[41] So, we'd have to *give up* the idea that the modal is calculated relative to a domain characterized by the fact that Jay smokes. But this is unfortunate – it's *because* Jay smokes that he ought to quit.[42] Note that (154) isn't unique in being affected by the Orientation Problem and the Event Identification Problem – these considerations apply to *ought* sentences with eventive complements more generally. These problems seem *not* to affect *ought* sentences with stative complements.

What we need is two-fold: a way to shift the orientation forward, and a way for modal base to project from an event variable in or around the same clausal position, but without being so intimately tied to the event of the complement. In fact, we don't need to dramatically revise the framework to fix this problem. For present-oriented stative complements, (157) is still an adequate schema for the logical form of these sentences. So, all that's needed is a mechanism to introduce an additional event variable in between the modal and the most deeply embedded event description for the root modal sentences with *eventive* prejacents. And since the additional structure is required for the interpretation of the expression, we can appeal to a coercion mechanism that simply inserts the demanded structure into the logical form. Let's call the operator that inserts this structure 'Ω'. This new event variable will be related to the event description of the designated action in a way that reflects the temporal precedence of the former. Since the additional event

[41]At least, not unless you can ensure that the quitting-smoking event must *come after* the smoking in the relevant worlds. But the schema gives us no way to build this constraint in.

[42]There are fairly easy fixes for the Orientation Problem. In fact, in her dissertation, Hacquard gives a logical form for these problematic sentences (root modal sentences with eventive complements) that gets around the Orientation Problem. Cf. also Kratzer 2010 and Matthewson 2012. But the Event Identification Problem is a bigger problem, as long as deontic modals are keyed to the VP event, as they are in Hacquard's system.

variable is inserted just above the existing, the CAUSE predicate is a good candidate for effectuating the temporal relationship we have in mind.[43]

(158) Defined if t = t_u

$\exists$e [$t \subseteq \tau(e)$ OUGHT (D)$^{f(e),g(e)}$: $\exists$e$'$(ϕ(e$'$) & CAUSE(e,e$'$))]

Now we have workable schema for representing deontic modals in a framework that respects the scopal properties of different modals, and the coercion operator Ω allows us to evade the Orientation and Event Identification Problems. Moreover, the schema goes some distance to integrating the semantics of modals and the Davidsonian event semantics which figured in discussions in earlier sections of the paper. Importantly, the introduction of the additional event variable was motivated purely on grounds related to the semantics of modals – the temporal interpretation of modals, and the mechanism by which D is determined.[44]

5.6.3. Putting The Pieces Together

We can now raise the issue of agentive-*oughts* anew. Having laid out the issues in the last two subsections, the substance of my proposal might starting being obvious. But before I make it explicit, let me return the observations mooted earlier about agentive-*oughts* in general, originally made in section 5.4.

> Agentive *ought*-sentences tend to have eventive complements, as a matter of course. Relatedly, agentive *ought*-sentences are always future-oriented.

[43] In fact, there is independent evidence for the existence of an additional event variable in the relevant sentences. Some of the diagnostics for locating an underlying event variable corroborate its presence. Cf. Homer 2011 for discussion.

[44] It would be useful to differentiate the event variable introduced by Ω (the variable e in (158)) from the one introduced by the lexical verb in the complement of the modal (e$'$). When I need to distinguish them, I'll refer to e as the "additional" event variable.

> Non-agentive *ought*-sentences are variously future or present oriented. The embedded predicate can be stative or eventive, but we tend not to see *ought*-sentences with eventive predicates that describe an intentional action as of this non-agentive type.

What this first observation tells us is that *all* the sentences in (136) conform to the schema in (158) – they have the additional event variable. By contrast, *some* of the *ought*-sentences in (137), specifically, the present-oriented (137c) and (137d) are not so coerced – for these sentences, the schema in (157) provides an adequate representation of their logical form. There are other *ought*-sentences amongst those in (137), which are future-oriented and for which the schema in (158) is an appropriate representation of their logical form. In sum, the agentive *ought*-sentences are a proper subset of the sentences which include the coercion operator Ω.

We can turn to the discussion of external arguments to see how agentive *ought*-sentences differ from non-agentive but still future-oriented *ought*-sentences. According to the "severed" conception of external arguments, an external argument composes with an event description through a special composition rule called event identification. My proposal is that agentive *ought*-sentences are those sentences where the event identification composition rule is applied to a sentence to which the coercion operator Ω has already been applied. The resulting schema for an agentive *ought*-sentence will be as in (159).

(159) Defined if $\mathrm{t} = t_u$

$\exists \mathrm{e}\ [t \subseteq \tau(\mathrm{e})\ \&\ \mathrm{OUGHT}\ (\mathrm{D})^{f(e),g(e)}\colon\ \exists \mathrm{e}'(\phi(\mathrm{e}')\ \&\ \textsc{cause}(\mathrm{e},\mathrm{e}')\ \&\ \textsc{agent}(\mathrm{e},x))]$

Here is a quasi-English gloss on this schema: "There is an event in which the utterance time is contained, such that the worlds which are the best g-ranked worlds consistent with the circumstances of e are worlds where x causes a ϕ-event e′." Applying this schema

to (154) and letting x be Jay, we get: "there is an event in which the utterance time is contained, such that the worlds which are the best g-ranked worlds consistent with the circumstances of e are worlds where Jay causes an event e$'$ of him quitting smoking." So, given the circumstances – that he smokes, that smoking causes cancer, etc. – realizing one of the most highly ranked worlds (ranked highly according to the standards that he do what he can to stay healthy, that he avoid spending money frivolously, etc.) is matter of him causing a particular kind of action – namely one where he quits smoking.

The way the coercion theoretic account of agentive *ought* proposes something like a veriadic function operator can now be made more clearly. The coercion theoretic account of agentive *ought* proposes that what is responsible for agentive interpretations of *ought* is really the joint contribution of the coercion operator Ω and a subsequent application of event identification. The resulting expression is like the veriadic function operator in the following sense; it takes a 0-place relation (the fully saturated verbal predicate) and returns a 1-place relation, which can take an agent argument.

5.7. Flexibility Offered By the Coercion Theoretic Account

Those are the basics of the coercion-theoretic proposal for agentive *ought*. I think this view deserves serious consideration and is worthy of further development. There are a few details I've left sketchy – some because these are the details that are subject to further investigation and research, but others because the view offers some flexibility on how exactly to fill in these details. There are two parts of the account that offer some flexibility, which I will highlight in the sections that follow.

5.7.1. The Additional Agent Argument

I haven't said much about the additional AGENT argument that is introduced for agentive *ought*, leaving it open how exactly this argument is filled. Clearly, it is unpronounced, so it will need to be some kind of expression that is not phonologically realized, like a variable. This raises the question how the referent of such a variable will be resolved. Perhaps the simplest implementation was that the variable be resolved much like a pronoun. In principle, this would allow a certain permissiveness about who could be the agent of *ought*.

Some philosophers might not like this permissiveness. Luckily, the view offers some degree of flexibility on this, which is useful in light of the considerable disagreement in the literature about *who* can be the agent for an agentive *ought* sentence. Schroeder [2011, pp. 30 – 33], for example, argues that it can *only* be the subject of the complement. If this is the case, there may very well be grammatical reasons for this that are consistent with my proposal. For example, it could be that the argument moves from the original position as the agent of the embedded clause to the higher position as the agent of the new projection,[45] leaving a trace in the original position. This would mean that the agent of the *ought* and the sentential subject are always the same, thus constraining the interpretation of agentive *ought* as Schroeder would want.

Schroeder makes a further claim, which is that it is *incoherent* if the subject of agentive *ought* were someone or something other than the subject of the embedded clause.[46] If it is really *incoherent* to have someone other than the embedded subject as the subject

[45] from the Specifier of $v\mathrm{P}_1$ to the specifier of $v\mathrm{P}_2$, for example

[46] Cf. p. 31 "Similarly, 'Jon ought that Jon gets rich' is ungrammatical, but I don't think it is incoherent [...]. It means, pretty obviously, that Jon ought to get rich. What I don't understand is what it would be for it to be the case that Jon ought that Mary gets rich, unless it is supposed to be that Jon ought to make it the case that Mary gets rich, or to ensure that she gets rich..."

of the agentive *ought*, then we have to worry little about over-generating interpretations of the unpronounced pronoun with someone other than this agent – speakers will not resolve this pronoun with a value that renders the sentence incoherent. The only coherent interpretation will be the one that duplicates the lower agent argument.

However, I'm suspicious of Schroeder's claim here, and agree with Bronfman and Dowell [2018] and Chrisman [2015] that the agent of *ought* can be someone other than the sentential subject. Let's presume that the referent of the variable is determined in a similar fashion to pronoun resolution for deictic pronouns. Consider the following scenario:

> Tom and Bill, now happily married, visit a cardiologist for some health issues Tom has been having. The cardiologist is trying to get Tom to make some lifestyle changes for the sake of his health, but knows that Tom is stubborn and set in his ways. He knows that, if anyone, Bill can make sure that Tom actually makes some of these adjustments to his lifestyle. After Tom leaves, the cardiologist turns to Bill and says to him.
>
> (160) Tom ought to exercise daily.

Given the circumstances and the previous discourse, it seems reasonable to think that this is a genuine case of agentive *ought* where the agent is Bill instead of Tom. Of course, I'm trading in intuitions here instead of arguments, but the important point is that if this *is* an acceptable interpretation of the agentive *ought*, the coercion theoretic account could allow for this. One would merely need to maintain that the agent of *ought* is resolved to be the most salient person the discourse participants could coordinate on. There might very well be well-defined constraints on how such a pronoun could be resolved which are worth investigating further. To a first approximation, it would seem that the agent of *ought* can't have the status of "discourse new". As the discussion shows, the framework provides some flexibility on this issue, depending on one's antecedent theoretical inclination.

5.7.2. On "CAUSE"

I haven't said much about what CAUSE is, though the informal glosses I offered invoke our ordinary folk concept of causation. The reason for using CAUSE is that this predicate already features prominently in relations among events in decompositional analysis of verbal constructions – in using CAUSE, I'm not positing a novel kind of relation. This makes agentive *ought* look like a kind of causative construction. However, there is reason to suppose that the concept of CAUSE being invoked for agentive *oughts* isn't the same as the CAUSE in lexical causatives, since those tend to employ a kind of causal concept linguists and psychologists have called direct causation.[47] While the concept of direct causation might be appropriate for some instances of agentive *ought*, there are cases of agentive *ought* that where direct causation does not seem to be appropriate. If *Tom ought to exercise daily* from above is an instance of agentive *ought*, the relation between Bill's interventions and Tom's exercise is probably not accurately characterized by direct causation. However, just because the CAUSE we need is probably not direct causation, this is not a problem in and of itself for two reasons. First of all, even though CAUSE is related to our ordinary notion of causation, it's often thought that it doesn't overlap with this concept exactly. So it would be acceptable if the kind of relation invoked in the logical form of agentive *ought* runs afoul of some of our folk intuitions about causation. Secondly, it's not clear that all instances of CAUSE are the same – there might be multiple CAUSEs.

This shows another area of flexibility with the proposal. The account proceeds in two steps. First, Ω introduces an additional event variable into the logical form, which

[47]Cf. Wolff 2003 for an example of this work.

is related to the "old" event variable. Second, event identification composition principle adds an AGENT argument to the new event variable. The two event variables must be related in a way that represents the temporal priority of the former. If it turns out that CAUSE is an inappropriate relation linking the two events, you could posit a novel relation for this purpose (call it 'R_c').[48] But characterizing this relation more precisely requires additional work. Still, a novel relation like R_c is probably not needed.[49]

5.8. Clarifying the Role of Ω

One issue that's worth making clear is that the Ω operator applies to eventive complements of *ought* generally, not just those we'd interpret as agentive. For example, on its non-agentive reading, (161) has an application of the Ω operator.

(161) Luckless Larry ought to win the lottery.

(161) would be coerced into the following logical form.

(162) Defined if t = t_u

$\exists$e [$t \subseteq \tau(e)$ OUGHT (D)$^{f(e),g(e)}$: $\exists$e′(Larry-win-the-lottery′(e′) & CAUSE(e,e′))]

Why does (161) not then have an agentive reading? The answer is that although the Ω operator applies to the complement prior to composing with the modal, the resulting expression does not get an AGENT argument through event identification. Understanding why not is crucial, since the application of event identification is not haphazard. Let's

[48]The subscripted c indicates that the relation is likely to have a causal component to it, at least in the case of agentive *oughts*.

[49]There is work on linguistic constructions, like futurates and *have*-causatives, that employ the notion of dispositional causation for CAUSE, which can perhaps do the work we need. Cf. Copley 2018. Convergence between Copley's proposals and this one would not be surprising, as her work was an inspiration behind the coercion theoretic account of agentive *ought*.

give a gloss of (162) and walk through it. This says, roughly, "There is an event in which the utterance time is contained, such that the worlds which are the best g-ranked worlds consistent with the circumstances of e are worlds where e causes an event e$'$ of Larry winning the lottery." e is a current circumstance of Larry's environment, say the circumstance according to which the Powerball lottery is being played next week. The modal base is a function from e to the set of relevant circumstances – namely, that Larry is in financial distress, that he's a a deserving guy, that the lottery being played is fair. In order for e to cause an event of Larry's winning of the lottery (e$'$), a number of things need to happen. Larry needs to buy a ticket, decide on the appropriate numbers, etc. The causal chain that will lead to his winning is one where Larry's circumstances e causes certain lottery numbers to be drawn from the draw machine. But there is no plausible agent whose machinations bring about this outcome. This is just to say, there's no agent that brings it about that the Larry wins the lotto. It's just that a causal chain emerges from the circumstances and results in Larry winning the lottery.

Compare with a different context, where (161) does have an agentive interpretation – the admittedly remote one we spoke in section 5.5. We can get this interpretation if we suppose that, as part of Larry's circumstances, he has the knowledge and means by which he can rig the lottery, and we are attempting to offer advice on how Larry can deal with his debts. In this circumstance, the appropriate logical form would be (163).

(163) Defined if t = t_u

$\exists$e [$t \subseteq \tau(e)$ & OUGHT (D)$^{f(e),g(e)}$: $\exists$e$'$(ϕ(e$'$) & CAUSE(e,e$'$) & AGENT(e,x))]

(163) reflects a circumstance where there is a causal chain linking Larry's circumstances with the outcome of his winning the lottery – one that involves the actions Larry would undertake in rigging the lottery.

Recall that the underlying event can constrain the kind of external argument an expression might get through event identification. As mentioned earlier, the state of *owning a dog* doesn't have an agent, so event identification cannot compose with an argument theta-marked as AGENT. In the non-agentive interpretation of (161), the nature of the circumstances don't allow for an agent bringing it about that Larry wins the lottery. So interpreted, event identification would not allow composition with an AGENT argument. But in the agentive interpretation of (161), the circumstances *do* allow for someone to bring about Larry's winning – namely Larry.[50]

The fact that Ω can apply to a complement without yielding an agentive interpretation raises the possibility that all of the (1)-sentences – our paradigmatic agentive *ought* sentences – should at least *allow* non-agentive interpretations. This is because a sentence can have the extra structure given by Ω, without having an additional agent argument, the discussion of Luckless Larry showed. This is the correct result, I think. Sometimes such sentences are used to express the evaluative claim that it would be better for a particular event to occur rather than offering advice to or making a deliberative judgment.[51]

[50]Note that this doesn't mean we have to think that to interpret the sentence like this, we need a distinct interpretation of the verb *win*, as with the version of AIP discussed earlier.

[51]Some authors have noted that it's particularly unlikely for an it-cleft *ought* sentence (think: *It ought to be the case that Bill kisses Tom*) to have an agentive interpretation. At present, I don't have a robust explanation on offer for this, though I'm partial to the suggestions made in Finlay and Snedegar 2014. It's entirely plausible that the pragmatics of such constructions strongly disfavor agentive interpretations. What would be interesting to me is whether the explanation for this is purely pragmatic, or whether there are syntactic or semantic issues behind this tendency as well.

5.9. The Payoffs of the Coercion Theoretic Approach

I think the coercion theoretic account of agentive *ought* is worth taking seriously. Let's take stock of some of the explanatory payoffs of the coercion theoretic account of agentive *ought*. First, we can pay lip service to the intuition that there is something special about agentive *ought* sentences; they are tied to action in some fairly intimate way. The coercion theoretic account has this relationship encoded in the very logical form of the sentence in a pretty explicit way.

Secondly, adopting the coercion theoretic account allows us to explain the motivating data fairly directly. The explanation of Agency Sensitivity is related to the observation made above about the relationship between the obligation holder and the action being encoded in the logical form. Recall the example: if Tom intended to kiss Bill, and wound up doing so, but not through an exercise of his agency (because his forming the intention made him so nervous he tripped and planted his lips on Bill), this doesn't seem to capture what the agentive *ought* is counseling. In other words, if Tom is to be the agent that causes his kissing of Bill, he needs to stand as an agent to the event of him kissing Tom, and not as accidental participant. The scenario where nervous Tom accidentally kisses Bill would make the the non-agentive interpretation true, but it would fail to make the agentive interpretation true, as desired.

The explanation of Argument Asymmetry is likewise quite clear. The "agent" of the *ought* (if we may speak in this way) is the referent of the variable theta-marked as AGENT. As we saw in the previous section, there's some freedom in the proposal to think of this agent as being coreferential with the sentential subject or not. But consider (140a) (=*Tom ought to dance with Bill*). If the agentive interpretation of (140a) has Tom as the agent of

the *ought*, then this sentence will not be equivalent to (140b) (=*Bill ought to dance with Tom*). Even if we assume an occurrence of (140b) where *Tom* is the agent of the *ought*, the way Tom would cause himself to dance with Tom is very different than the way he might cause Bill to dance with him.

There are other theoretical payoffs of the coercion theoretic account. It provides an explanation of the grammatical patterns that characterize the paradigmatic agentive *ought* sentences, noted in the generalizations discussed in section 5.4. Agentive *ought* sentences are all future oriented because the structure that *makes* them future-oriented is a prerequisite for the agentive interpretation. This is the grammatical reason why we do not see agentive *oughts* which are present oriented. Relatedly, agentive *oughts* tend to have eventive complements because eventive complements are precisely those that require Ω coercion for their interpretation.[52]

Another reason to take the coercion theoretic proposal seriously; it shows that the taxonomy supposed by the literature misses an important possibility. The complement ambiguity strategy is not exhausted by AIP, as is often presupposed. The coercion theoretic account is not AIP (at least, not the way supposed by proponents of AIP), but still locates the source of agentive *ought* in an ambiguity in the complement. Also, we've developed the coercion theoretic account of agentive *ought* squarely within the Kratzerian paradigm of modal semantics, so it shows that adopting a Kratzer-style semantics doesn't

[52]This isn't to deny that you could have an agentive *ought* sentence with a stative complement, but even such cases support the coercion theoretic account. This is because (i) if the sentence truly is an agentive *ought*, it will be interpreted as future-oriented and (ii) it will have an interpretation that such-and-such a state should be brought about, as opposed to meaning that this state's obtaining is the best. *Ought* sentences with stative complements can have Ω coercion, but this is not obligatory, as it is with eventive complements. But the agentive interpretation of *ought* sentences with stative complements comes only from those sentences that show evidence of Ω coercion.

at all commit us to the denial position with respect to agentive *ought*, as others have claimed.

Bibliography

Dorit Abusch. Temporal and circumstantial dependence in counterfactual modals. In *Proceedings of the Sixteenth Amsterdam Colloquium*, pages 37–42. Citeseer, 2007.

Emmon Bach. On time, tense, and aspect: a essay in english metaphysics. In Peter Cole, editor, *Radical Pragmatics*. Academic Press, London, 1981.

Emmon Bach. The algebra of events. *Linguistics and Philosophy*, 9:5 – 16, 1986.

Nuel D. Belnap, Michael Perloff, and Ming Xu. *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford University Press on Demand, 2001.

Jurgen Bohnemeyer and Mary Swift. Event realization and default aspect. *Linguistics and Philosophy*, 27:263– 296, 2004.

Joan Bresnan. The passive in lexical theory. In Joan Bresnan, editor, *The mental representation of grammatical relations*, pages 3 – 86. Cambridge, MA: MIT Press, 1982.

Aaron Bronfman and J. L. Dowell. The language of "ought" and reasons. In Daniel Star, editor, *Oxford Handbook of Reasons and Normativity*. Oxford University Press on Demand, 2018.

John Broome. Williams on ought. In Ulrike Heuer and Gerald Lang, editors, *Luck, Value, and Commitment: Themes From the Ethics of Bernard Williams*. Oxford University Press, Usa, 2012.

John Broome. *Rationality Through Reasoning*. Wiley-Blackwell, 2013.

Fabrizio Cariani. Epistemic and deontic should. *Thought*, 2:73 – 84, 2013.

Fabrizio Cariani. Deontic modals and probability: One theory to rule them all? In Nate Charlow and Matthew Chrisman, editors, *Deontic Modality*. Oxford University Press, 2016.

Gregory N. Carlson. Generics, habituals and iteratives. In Alex Barber, editor, *Encyclopedia of Language and Linguistics*. Elsevier, 2005.

Rudolf Carnap. *Meaning and Necessity*. University of Chicago Press, 1947.

Hector-Neri Castaneda. *Thinking and Doing: The Philosophical Foundations of Institutions*. D. Reidel Pub. Co., 1975.

Lucas Champollion. The interaction of compositional semantics and event semantics. *Linguistics and Philosophy*, 38(1):31–66, 2015.

Roderick M. Chisholm. The ethics of requirement. *American Philosophical Quarterly*, 1 (2):147–153, 1964.

Noam Chomsky. *Aspects of the Theory of Syntax*. MIT Press, 1965.

Matthew Chrisman. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press, 2015.

Guglielmo Cinque. *Adverbs and Functional Heads: A Cross-Linguistic Perspective*. Oxford University Press, 1999.

Jonathan Cohen. Indexicality and the puzzle of the answering machine. *Journal of Philosophy*, 110(1):5–32, 2013. doi: 10.5840/jphil2013110143.

Jonathan Cohen and Eliot Michaelson. Indexicality and the answering machine paradox. *Philosophy Compass*, 8(6):580–592, 2013. doi: 10.1111/phc3.12039.

Bernard Comrie. *Aspect*. Cambridge University Press, 1976.

Cleo Condoravdi. Temporal interpretation of modals. In David Beaver, Stefan Kaufmann, Brady Clark, and Luis Casillas, editors, *The Construction of Meaning*, pages 59 – 88. CSLI Publications, 2002.

B. Jack Copeland. The genesis of possible worlds semantics. *Journal of Philosophical Logic*, 31(2):99–137, 2002. doi: 10.1023/A:1015273407895.

Bridget Copley. What should "should" mean? 2006.

Bridget Copley. Temporal orientation in conditionals. In Jacqueline Guéron and Jacqueline Lecarme, editors, *Time and Modality*, volume 75 of *Studies in Natural Language and Linguistic Theory*. Springer, 2008a.

Bridget Copley. The plan's the thing: deconstructing futurate meanings. *Linguistic Inquiry*, 39(2):261– 274, 2008b.

Bridget Copley. *The Semantics of the Future*. Routledge, 2009.

Bridget Copley. Causal chains for futurates. In Mikhail Kissine, Philippe de Brabanter, and Saghie Sharifzadeh, editors, *Future Times, Future Tenses*, pages 72–86. Oxford University Press, 2014.

Bridget Copley. Dispositional causation. *Glossa: a journal of general linguistics*, 3(1), 2018.

M. J. Cresswell. *Logics and Languages*. Methuen, 1973.

Osten Dahl. the marking of the episodic/ generic distinction in tense-aspect systems. In Manfred Krifka, Francis Jeffry Pelletier, Gregory N. Carlson, Alice ter Meulen, Gennaro Chierchia, and Godehard Link, editors, *the Generic Book*, page 412. University of Chicago Press, 1995.

Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *the Logic of Decision and Action*. University of Pittsburgh Press, 1967.

Ashwini Deo. Unifying the imperfective and the progressive: partitions as quantificational domains. *Linguistics and philosophy*, 32(5):475–521, 2009.

Edit Doron. Point of view as a factor of content. In *Semantics and Linguistic Theory*, volume 1, pages 51–64, 1991.

J.L. Dowell. A flexible contextualist account of epistemic modals. *Philosopher's Imprint*, 11(14), 2011.

David Dowty. *Word Meaning and Montague Grammar: the Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ*, volume 7 of *Studies in Linguistics and Philosophy*. Springer, 1979.

Hans Bernhard Drubig. On the syntactic form of epistemic modality. 2001.

Billy Dunaway. Review of the meaning of ': Beyond descriptivism and expressivism in metaethics. *Journal of Philosophy*, 114(3):155–159, 2017. doi: 10.5840/jphil2017114312.

Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Gerhard Preyer and Georg Peter, editors, *Contextualism in Philosophy*, pages 131–170. Oxford University Press, 2005.

Murvet Enç. Towards a referential analysis of temporal expressions. *Linguistics and Philosophy*, 9:405– 426, 1986.

Murvet Enç. Anchoring conditions for tense. *Linguistic Inquiry*, 18(4):633– 657, 1987.

Murvet Enç. Tense and modality. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*. Blackwell, 1996.

Fred Feldman. *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel Publishing Company, 1986.

Stephen Finlay. *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press, 2014.

Stephen Finlay and Justin Snedegar. One ought too many. *Philosophy and Phenomenological Research*, 89(1):102–124, 2014.

Kai von Fintel and Anthony S. Gillies. An opinionated guide to epistemic modality. In Tamar Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology: Volume 2*, pages 32–62. Oxford University Press, 2007.

Kai von Fintel and Anthony S. Gillies. Cia leaks. *Philosophical Review*, 117(1):77–98, 2008.

Kai von Fintel and Anthony S. Gillies. Must... stay... strong! *Natural Language Semantics*, 18(4):351 – 383, 2010.

Kai von Fintel and Irene Heim. Lecture notes on intensional semantics. 2002.

Kai von Fintel and Sabine Iatridou. How to say ought in foreign: the composition of weak necessity modals. In Jacqueline Guéron and Jacqueline Lecarme, editors, *Time and Modality*, Studies in Natural Language and Linguistic Theory, pages 115 – 141. Springer, 2008.

Robert Freidin. *Foundations of Generative Syntax*. MIT Press, 1992.

Peter T. Geach. Whatever happened to deontic logic? *Philosophia*, 11(1-2):1–12, 1982.

F. Giorgi and A. Pianesi. *Tense and Aspect. From Semantics to Morphosyntax*. Oxford University Press, 1997.

Michael Glanzberg. Context, content and relativism. *Philosophical Studies*, 136:1– 29, 2007.

Daniel Goodhue. Must φ is felicitous only if φ is not known. *Semantics and Pragmatics*, 10:14, 2017.

Paul Grice. *Aspects of Reason*. Clarendon Press, 2001.

Jane Grimshaw. *Argument Structure*. MIT Press, 1990.

Ian Hacking. Possibility. *The Philosophical Review*, pages 143–168, 1967.

Valentine Hacquard. *Aspects of Modality*. PhD thesis, MIT, 2006.

Valentine Hacquard. On the interaction of aspect and modal auxiliaries. *Linguistics and Philosophy*, 32:279 – 315, 2009.

Valentine Hacquard. On the event relativity of modal auxiliaries. *Natural Language Semantics*, 18:79 – 114, 2010.

Valentine Hacquard. Modality. In *Semantics: An International Handbook of Natural Language Meaning. HSK 33.2*, number 3, pages 1484–1515. Mouton De Gruyter, 2011.

Valentine Hacquard and Alexis Wellwood. Embedding epistemic modals in english: A corpus-based study. *Semantics and Pragmatics*, 5(4):1 – 29, 2012.

Peter Hallman. Instants and intervals in the event/ state distinction. MS, 2009.

Gilbert Harman. Review of the significance of sense: Meaning, modality, and morality. *Philosophical Review*, 82(2):235, 1973.

Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, UMass Amherst, 1982.

Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, 1998.

James Higginbotham. On semantics. *Linguistic Inquiry*, 16:547–593, 1985.

Charles F Hockett and Charles D Hockett. The origin of speech. *Scientific American*, 203 (3):88–97, 1960.

Vincent Homer. French modals and perfective: a case of aspectual coercion. In *Proceedings of the 28th West Coast Conference on Formal Linguistics*, pages 106–114, 2011.

Laurence Horn. *On the Semantic Properties of Logical Operators in English*. PhD thesis, UCLA, 1972.

Laurence Horn. *A Natural History of Negation*. University of Chicago Press, 1989.

John F Horty. *Agency and deontic logic*. Oxford University Press, 2001.

John F. Horty and Nuel Belnap. The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24(6):583–644, 1995.

I. L. Humberstone. Two kinds of agent-relativity. *Philosophical Quarterly*, 41(163):144–166, 1991.

Sabine Iatridou. The grammatical ingredients of counterfactuality. *Linguistic Inquiry*, 31 (2):231 – 270, 2000.

Ray Jackendoff. *Semantic Interpretation In Generative Grammar*. MIT Press, 1972.

Hans Kamp and Uwe Reyle. *From discourse to logic: Introduction to model theoretic semantics of natural language, formal logic and discourse representation theory*, volume 42. Springer, 2013.

David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes From Kaplan*, pages 481–563. Oxford University Press, 1977.

Lauri Karttunen. Possible and must. *Syntax and semantics*, 1, 1972.

Graham Katz. On the stativity of the english perfect. In Artemis Alexiou, Arnim von Stechow, and Rathert Monika, editors, *Perfect Explorations*. De Gruyter, 2003.

Stefan Kaufmann. the presumption of settledness. In *Proceedings of CLS*, volume 38, pages 313 – 328, 2002.

Stefan Kaufmann. Conditional truth and future reference. *Journal of Semantics*, 22(3): 231 – 280, 2005.

Stefan Kaufmann, Cleo Condoravdi, and Valentina Harizanov. Formal approaches to modality. In William Frawley, editor, *The Expression of Modality*, pages 71–106. Berlin–New York: Mouton de Gruyter, 2006.

Anthony Kenny. *Action, Emotion And Will*. Humanities Press, 1963.

Jeffrey C. King. Tense, modality, and semantic values. *Philosophical Perspectives*, 17(1): 195–246, 2003.

Jeffrey C. King. *The Nature and Structure of Content*. Oxford University Press, 2007.

Peter Klecha. Modality and embedded temporal operators. *Semantics and Pragmatics*, 9, 2016.

Wolfgang Klein. *Time in Language*. Routledge, 1994.

Angelika Kratzer. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(1):337 – 355, 1977.

Angelika Kratzer. The notional category of modality. In H.-J. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts. New Approaches in Word Semantics*, pages 38 – 74. De Gruyter, 1981.

Angelika Kratzer. Modality. In Arnim von Stechow and D. Wunderlich, editors, *Semantik: Ein Internationales Handbuch Zeitgenoessischer Forschung*, pages 639 – 650. De Gruyter, 1991.

Angelika Kratzer. Severing the external argument from its verb. In Laurie Zaring Johan Rooryck, editor, *Phrase structure and the lexicon*, pages 109–137. Springer, 1996.

Angelika Kratzer. More structural analogies between pronouns and tenses. In *Proceedings of Semantics and Linguistic Theory (SALT)*, volume 8, pages 92–110, 1998.

Angelika Kratzer. Situations in natural language semantics. *Stanford Encyclopedia of Philosophy*, 2008.

Angelika Kratzer. What "can" can mean. Lecture notes, University of Massachusetts, Amherst, 2010.

Angelika Kratzer. *Modals and Conditionals*. Oxford Studies in Theoretical Linguistics. Oxford University Press, 2012.

Angelika Kratzer. Modality for the 21st century. In *The language-cognition interface: Proceedings of the 19th International Congress of Linguists*, pages 179–199, 2013.

Manfred Krifka. Nominal reference, temporal constitution and quantification in event semantics. In Renate Bartsch, J. F. A. K. van Benthem, and P. van Emde Boas, editors, *Semantics and Contextual Expression*, pages 75–115. Foris Publications, 1989.

Manfred Krifka. Quantifying into question acts. *Natural Language Semantics*, 9(1):1–40, 2001.

Manfred Krifka, Francis Jeffry Pelletier, Gregory N. Carlson, Alice ter Meulen, Gennaro Chierchia, and Godehard Link. Genericity: An introduction. In *the Generic Book*. University of Chicago Press, 1995.

Saul A. Kripke. A completeness theorem in modal logic. *Journal of Symbolic Logic*, 24 (1):1–14, 1959. doi: 10.2307/2964568.

Saul A. Kripke. Semantical analysis of modal logic i. normal propositional calculi. *Zeitschrift fur mathematische Logik und Grundlagen der Mathematik*, 9(56):67–96, 1963. doi: 10.1002/malq.19630090502.

Nicky Kroll. Progressive teleology. *Philosophical Studies*, 172(11):2931–2954, 2015.

Fred Landman. The progressive. *Natural Language Semantics*, 1(1):1–32, 1992.

Marika Lekakou and Øystein Nilsen. What aspect can tell us about the future of must. In *Proceedings of the 8th international conference of Greek linguistics*, 2008.

E. J. Lemmon. Comment on d. davidson's "the logical form of action sentences". In Nicholas Rescher, editor, *The Logic of Decision and Action*. University of Pittsburgh Press, 1967.

Sarah-Jane Leslie. Generics: Cognition and acquisition. *Philosophical Review*, 117(1): 1–47, 2008.

Sarah-Jane Leslie. Generics oversimplified. *Noûs*, 47(3):28–54, 2013.

David Lewis. General semantics. *Synthese*, 22(1-2):18–67, 1970.

David Lewis. Index, context, and content. In Stig Kanger and Sven Öhman, editors, *Philosophy and Grammar*, pages 79–100. Reidel, 1980.

David K. Lewis. *On the Plurality of Worlds*. Blackwell Publishers, 1986.

Godehard Link. The logical analysis of plurals and mass terms: A lattice-theoretical approach. *Formal semantics: The essential readings*, pages 127–146, 2002.

John MacFarlane. Future contingents and relative truth. *Philosophical Quarterly*, 53 (212):321–336, 2003.

John MacFarlane. Nonindexical contextualism. *Synthese*, 166(2):231–250, 2009.

John MacFarlane. *Assessment Sensitivity: Relative Truth and its Applications*. OUP Oxford, 2014.

Matthew Mandelkern. A solution to karttunen's problem. In *Proceedings of Sinn und Bedeutung*, volume 21, pages 827–844, 2018.

Matthew Mandelkern. Bounded modality. *Philosophical Review*, 128(1):1–61, 2019a. doi: 10.1215/00318108-7213001.

Matthew Mandelkern. What 'must' adds. *Linguistics and Philosophy*, 42(3):225–266, 2019b.

Matthew Mandelkern. How to do things with modals. *Mind and Language*, 35(1):115–138, 2020. doi: 10.1111/mila.12231.

Alec Marantz. *On the nature of grammatical relations*. Linguistic Inquiry Monographs. MIT Press, 1984.

Ruth Barcan Marcus. *Modalities: Philosophical Essays*. Oxford University Press, 1993.

Lisa Matthewson. On the (non)-future orientation of modals. In *Proceedings of Sinn und Bedeutung*, volume 16, pages 431–446, 2012.

Eliot Michaelson. Shifty characters. *Philosophical Studies*, 167(3):519–540, 2014. doi: 10.1007/s11098-013-0109-7.

Richard Montague. Pragmatics and intensional logic. *Synthese*, 22(1-2):68–94, 1970.

Alexand Mourelatos. Events, processes and states. *Linguistics and Philosophy*, 2:415 – 434, 1978.

Stephen Neale. On a milestone of empiricism. In A. Orenstein and Petr Kotatko, editors, *Knowledge, Language and Logic: Questions for Quine*, pages 237–346. Kluwer Academic Print on Demand, 2000.

Dilip Ninan. Two puzzles about deontic necessity. In J Gajewkski, Valentine Hacquard, Bernard Nickel, and Seth Yalcin, editors, *New work on modality*, volume 51, pages 149–178. MIT Working Papers in Linguistics, 2005.

Dilip Ninan. Semantics and the objects of assertion. *Linguistics and Philosophy*, 33(5): 355–380, 2010.

Toshiyuki Ogihara. Tense and aspect in truth-conditional semantics. *Lingua*, 117(2): 392–418, 2007.

Terry Parsons. *Events in the Semantics of English: A Study in Subatomic Semantics*. Number 19 in Current Studies in Linguistics. MIT Press, 1990.

Barbara Partee. Some structural analogies between tenses and pronouns in english. *Journal of Philosophy*, 18(601 - 609), 1973.

John Perry. Thought without representation. *Proceedings of the Aristotelian Society*, 137: 137–152, 1986.

Paul M. Pietroski. On explaining that. *Journal of Philosophy*, 97(12):655, 2000.

Paul Portner. *Modality*. Oxford University Press, 2009.

S. Predelli. I am not here now. *Analysis*, 58(2):107–115, 1998. doi: 10.1093/analys/58.2.107.

Stefano Predelli. *Contexts: Meaning, Truth, and the Use of Language*. Clarendon Press, 2005.

H. A. Prichard. Does moral philosophy rest on a mistake? *Mind*, 21(81):21–37, 1912. doi: 10.1093/mind/XXI.81.21.

W. V. Quine. *From a Logical Point of View*. Harvard University Press, 1953.

Willard van Orman Quine. *Ontological Relativity and Other Essays*. Columbia University Press, 1969.

Gillian Ramchand. Stativity and present tense epistemics. *Proceedings of SALT*, 24(102 - 121), 2014.

Gillian Ramchand. *Situations and Syntactic Structures: Rethinking Auxiliaries and Order in English*. MIT Press, 2018.

François Recanati. *Literal Meaning*. Cambridge University Press, 2004.

François Recanati. *Perspectival Thought: A Plea for Moderate Relativism*. Oxford University Press, 2007.

Hans Reichenbach. *Elements of Symbolic Logic*. Dover Publications, 1980.

K. Romdenh-Romluc. 'i'. *Philosophical Studies*, 128(2):257–283, 2006. doi: 10.1007/s11098-004-7792-3.

Susan Rothstein. *Structuring Events*. Blackwell, 2004.

Gilbert Ryle. *The Concept of Mind*. Hutchinson and Co, 1949.

Philippe Schlenker. Context of thought and context of utterance: A note on free indirect discourse and the historical pr. *Mind and Language*, 19(3):279–304, 2004. doi: 10.1111/j.1468-0017.2004.00259.x.

Mark Schroeder. Ought, agents, and actions. *Philosophical Review*, 120(1):1–41, 2011.

Dana Scott. Advice on modal logic. In *Philosophical problems in logic*, pages 143–173. Springer, 1970.

Alan Sidelle. The answering machine paradox. *Canadian Journal of Philosophy*, 21(4): 525–539, 1991. doi: 10.1080/00455091.1991.10717260.

Henry Sidgwick. *The Methods of Ethics*. Thoemmes Press, 1874.

Alex Silk. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford University Press UK, 2016.

Carlota S. Smith. *the Parameter of Aspect*, volume 43 of *Studies in Linguistics and Philosophy*. Springer, 1991.

Carlota S Smith. Activities: states or events? *Linguistics and Philosophy*, 22(5):479–508, 1999.

Robert Stalnaker. Indicative conditionals. *Philosophia*, 5(3):269–286, 1975. doi: 10.1007/bf02379021.

Robert Stalnaker. Assertion. *Syntax and Semantics (New York Academic Press)*, 9: 315–332, 1978.

Robert Stalnaker. *Inquiry*. Cambridge University Press, 1984.

Robert C. Stalnaker. Possible worlds. *Noûs*, 10(1):65–75, 1976.

Jason Stanley. Context and logical form. *Linguistics and Philosophy*, 23(4):391–434, 2000.

Jason Stanley and Zoltan Gendler Szabó. On quantifier domain restriction. *Mind and Language*, 15(23):219–261, 2000.

Mark J Steedman. Verbs, time, and modality. *Cognitive Science*, 1(2):216–234, 1977.

Matthew Stone. Reference to possible worlds. *Technical Report 49, Rutgers University Center for Cognitive Science*, 1999.

Tim Stowell. Tense and modals. In Jacqueline Guéron and Jacqueline Lecarme, editors, *Syntax of Tense*. MIT Press, 2004.

Eric Swanson. *Interactions with Context*. PhD thesis, MIT, 2006.

Eric Swanson. Modality in language. *Philosophy Compass*, 3(6):1193 – 1207, 2008.

Eric Swanson. How not to theorize about the language of subjective uncertainty. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2009.

Barry Taylor. Tense and continuity. *Linguistics and Philosophy*, 1(2):199 – 220, 1977.

Paul Teller. Epistemic possibility. *Philosophia*, 2(4):303–320, 1972.

Hector-Neri Castañeda. Comment on d. davidson's "the logical form of action sentences". In Nicholas Rescher, editor, *The Logic of Decision and Action*. University of Pittsburgh Press, 1967.

Richmond H. Thomason. Indeterminist time and truth-value gaps. *Theoria*, 36(3):264–281, 1970.

Zeno Vendler. Verbs and times. *the Philosophical Review*, 66(2):143 – 160, 1957.

Hendrik Jacob Verkuyl. *On the compositional nature of the aspects*, volume 15. Springer Science & Business Media, 1972.

Emanuel Viebahn and Barbara Vetter. How many meanings for '? the case for modal polysemy. *Philosophers' Imprint*, 16(10), 2016.

G. H. von Wright. Deontic logic. *Mind*, 60(237):1–15, 1951.

Ralph Wedgwood. The meaning of 'ought'. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics: Volume 1*, pages 127–160. Clarendon Press, 2006.

Tom Werner. Future and non-future modal sentences. *Natural Language Semantics*, 14 (3):235–255, 2006. doi: 10.1007/s11050-006-9001-8.

Bernard Williams. *Moral Luck: Philosophical Papers 1973–1980*. Cambridge University Press, 1981.

Phillip Wolff. Direct causation in the linguistic coding and individuation of causal events. *Cognition*, 88(1):1–48, 2003.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983 – 1026, 2007.

Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–37, 2010.

Seth Yalcin. Modalities of normality. In Nate Charlow and Matthew Chrisman, editors, *Deontic Modality*. Oxford University Press, 2016.