

## ABSTRACT

Title of Dissertation: EPISTEMIC “MIGHT”:  
A NON-EPISTEMIC ANALYSIS

Quinn Patrick Harr, Doctor of Philosophy, 2019

Dissertation directed by: Associate Professor Alexander Williams  
Departments of Linguistics and Philosophy

A speaker of (1) implies that she is uncertain whether (2), making this use of *might* “epistemic.” On the received view, the implication is semantic, but in this dissertation I argue that this implication is no more semantic than is the implication that a speaker of (2) believes John to be contagious.

- (1) John might be contagious.
- (2) John is contagious.

This follows from a new observation: unlike claims with explicitly epistemic locutions, those made with “epistemic” uses of *might* can be explained only with reference to non-epistemic facts. I conclude that they express a relation, not to relevant information, but instead to relevant circumstances, and that uncertainty is implied only because of how informed speakers contribute to conversations. This conclusion dissolves old puzzles about disagreements and reported beliefs involving propositions expressed with *might*, puzzles that have been hard for the received view to accommodate. The cost of these advantages is to explain why the circumstantial modality expressed by *might* is not inherently oriented towards the future, as has been claimed for other circumstantial modalities. But this claim turns out to be false. The correct characterization of the temporal differences reveals that the modality expressed by *might* relates to propositions whereas other modalities relate to events. Neither sort is epistemic.

EPISTEMIC “MIGHT”: A NON-EPISTEMIC ANALYSIS

by

Quinn Patrick Harr

Dissertation submitted to the Faculty of the Graduate School of the  
University of Maryland, College Park, in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
2019

Advisory Committee:

Associate Professor Alexander Williams, Chair  
Professor Valentine Hacquard  
Professor John Horty  
Associate Professor Eric Pacuit  
Emeritus Professor Paul Pietroski

© Copyright by  
Quinn Patrick Harr  
2019

## **Dedication**

to Christine,  
and to Adelaide and Emmeline,  
in memory of Mark, husband and father

~

and to Grandma Savoie,  
on what would have been her 93<sup>rd</sup> birthday

## Acknowledgements

οὕτω δ' ἔοικε καὶ τοῖς φιλοσοφίας κοινωνήσασιν: οὐ γὰρ πρὸς χρήμαθ' ἡ ἀξία μετρεῖται, τιμὴ τ' ἰσόρροπος οὐκ ἂν γένοιτο, ἀλλ' ἴσως ἰκανόν, καθάπερ καὶ πρὸς θεοῦς καὶ πρὸς γονεῖς, τὸ ἐνδεχόμενον.

~ Aristotle, *Nicomachean Ethics* 1164b2-6

Aristotle's sentiment once seemed to me to be overwrought. No longer so.

To my advisor, Alexander Williams: for seeing the importance of a minor data point in a term paper on epistemic modals and encouraging me to develop it into a dissertation. For his guidance, and sometimes prodding, every step of the way. For his faith in the truth of so heterodox an idea, and his persistence in helping me to rise to the occasion of its defense, even when I was sorely lacking in both. Alexander has been my intellectual midwife, and I owe him a debt of gratitude that cannot be repaid. Much more should be said than can be said here. But let me add, as a final note of thanks, that, whenever I manage to write well, it is because Alexander's voice is in the back of my mind, instructing me how to do so.

To Valentine Hacquard: for encouragement, clarity, and direction in all of our meetings. For pointers to the literature and a pragmatic focus. And for the gift of a computer when I was trying to write this dissertation on a tablet. As with Alexander, so also with Valentine (and with the whole committee): much more should be said.

To Paul Pietroski: for the intellectual excitement of his courses and for the long exchanges in and after reading groups. To Jeff Harty: for his encouragement when I needed it most and for introducing me to default logic which, while it did not make it into the final draft of the dissertation, is always in the back of my mind. And to Eric Pacuit: for indefatigable kindness and good cheer, and for his blithe and contagious confidence that I would finish this dissertation in style.

To faculty not on my advisory committee. Most of all to Rachel Singpurwalla, but also very much to Georges Rey, for the friendship and mentorship that sustained me from early on in the program. To Jerry Levinson for his hospitality and our shared love of music. To Sam Kerstein and Sue Dwyer for their warmth and thoughtfulness. To Peter Carruthers and Dan Moller for advice. To John Brown for conversation. To Aidan Lyon for professional development. And to Norbert Hornstein, for incisive feedback, and Howard Lasnik, for helping me to fall in love with linguistics.

To faculty from my undergraduate institutions. Most of all to David Noe, who organized a retreat on Augustine's *de Magistro* that inspired me to continue my education and whose friendship I cherish to this day. To Kevin Culberson, Todd Bates, and Erik Root, all of them also present at that retreat, for opening my mind to new ways of thinking. To Kenneth Bratt, perhaps the most humane individual I have ever met and the best teacher I have ever had. To Jeff Winkle, for modeling how love for the profession can be combined with commitment to a community.

To Kevin Corcoran, for helping set me on the path to philosophy, and Gregory Mellema, Lee Hardy, and especially Rebecca Konyndyk DeYoung, for teaching me its history. To Georgios Kazantzidis, for helping me realize graduate school was within the realm of possibility, and also Elizabeth Baigent, Stan Rosenberg, David Crump, and Marlys Admiraal. I am fortunate to have all the teachers I had.

To my own students, in turn, especially my Aristotle and environmental ethics students: for their sincerity, intelligence, and passion, which provided inspiration when my own was lacking. Every time I teach, I become optimistic for the future.

To the support staff in the philosophy department at UMD. To Louise Gilman, truly the *sine qua non*, who made not only me but every other graduate student in the program feel loved and cared for from day one. To Antoinette Allen and John Boone, who fielded a thousand emails. And to Sheryl Joseph, my sister from another mister, whose raucous laughter, blended with mine, filled the whole department.

To my colleagues, many of them close personal friends. To Yu Izumi, Brendan Ritchie, Christopher Vogel, Lisa Leininger, Dimiter Kirilov, Lane DesAutels, Lucas Dunlap, and Deepak Mirchandani for teaching me the ropes. To Max Bialek, Ryan Ogilvie, Brock Rough, and Evan Westra for conversation. To Xuan and Max, for more than I can say. To Aleks, Dace, and Freya, who visited my family regularly in Michigan and quickly became part of it themselves. To Mike, the best officemate imaginable, and to his wife Emma, for their friendship and hospitality. To Julius.

To Alexey Alexei, Hong Chen, Kyley Ewing, Andrew Knoll, Jimmy Licon, Aida Roige, Aiden Woodcock, and Juanjuan Wu for friendship. And to the wonderful members of my cohort, and also dear friends, Kelsey, Cindy, and Christian.

To my colleagues from the linguistics department who took part in UMD's lively PHLING reading group and sometimes also in the venerable institution of Semantic Valueball: Dustin Chacón, Aaron Doliana, Allyson Ettinger, Jeffrey Green, Angela He, Mina Hirzel, Tyler Knowlton, Shevaun Lewis, Morgan Moyer, Max Papillon, Sigwan Thivierge, and Yu'an Yang. To Anouk Dieuleveut and Annemarie van Dooren, fellow travelers on the path of modality. And especially to Rachel Dudley and Alexis Wellwood, the very best colleagues imaginable, and also good friends.

To audiences at conferences where I have presented my dissertation research, and to referees at journals where I have submitted it. To Tom Grano, Tim Hunter, Dan Goodhue, and Pranav Anand for helpful meetings while they were at Maryland. To Christopher Gauker, for organizing a masterclass with Angelika Kratzer, and to Angelika Kratzer especially, for the conversations and correspondence that followed.

To others who made my sojourn in Maryland more enriching than I ever could have imagined. To Leslie Keiser, for providing a home away from home and being like a second mother. To Chelsea Stolt, for introducing me to Greek dancing and so much more. To the UMD Men's Chorus, for the creative outlet, camaraderie, and beauty it provided. To directors Greg Graff, Ianthe Marini, Steve Kim, and to friends Mae Almonte, Katy Hylton, Andrew Jones, and Eugene Won. To those who made it all possible: Dr. MacLary, Dr. Teddy, and Lauri Johnson. And to my friends Sean and Marissa, who hosted me in NYC and visited me in DC and are simply the best.

Most of all to my family, who drew me to Michigan summer after summer and sometimes for whole academic years at a time. My parents were my first teachers, and my mother my teacher for all of K-12. Most importantly, they have taught me how to live, and I aspire to be like them. My nine siblings and three nieces, in turn, have taught me how to love, and I cannot imagine the world without Stephanie, Kyle (and Jaime), Christine, Sherika, Jacob, Whitney, Drew, Austin, Darcie, or without Adelaide, Emmeline, and Jemma. My extended family also, my Aunt Kathy and Uncle Bruce especially, have been incredibly supportive of my research and have helped to make our small part of this globe the most welcoming, homey place ever.

To my Grandma Savoie, with all my love, for seeing me through to the defense. She would always ask if I was still working on my "paper". It is now finally done. Loved by all who knew her, we all miss her more than words can say.

Finally, to C.G.-S., for his helping me to realize that what I had long thought to be impossible may not in fact be so, perhaps the most important lesson of all.

# Table of Contents

<b>Dedication.....</b>	<b>ii</b>
<b>Acknowledgements.....</b>	<b>iii</b>
<b>Table of Contents.....</b>	<b>v</b>
<b>Chapter One: Introduction.....</b>	<b>1</b>
1.1    Epistemic uses of modals.....	1
1.2    Outline of the dissertation.....	7
<b>Part One: Fundamentals.....</b>	<b>12</b>
<b>Chapter Two: A Problem for Epistemic Analyses .....</b>	<b>13</b>
2.1    The basic problem.....	13
2.2    The full scope of the problem.....	16
2.2.1    Expanding the data set.....	16
2.2.2    Clarifying what is at stake .....	17
2.3    Potential force-modifier solutions .....	20
2.3.1    Traditional force-modifier accounts .....	20
2.3.2    Revised force-modifier accounts .....	22
2.4    Potential epistemic truth conditional solutions.....	25
2.4.1    Information states as an implicit part of content .....	26
2.4.2    Removing information states from content .....	34
2.5    Toward a non-epistemic solution.....	36
<b>Chapter Three: The Development of a Non-Epistemic Analysis .....</b>	<b>38</b>
3.1    Introduction.....	38
3.2    Contingent claims on a non-epistemic analysis.....	39
3.3    The framework for a non-epistemic analysis.....	41
3.3.1    The conceptual framework.....	42
3.3.2    The formal framework.....	46
3.3.3    The formal framework revisited.....	48
3.4    Do we need an epistemic overlay for the analysis?.....	50
3.4.1    Non-epistemic truth conditions and triviality.....	51
3.4.2    Non-epistemic truth conditions and speaker uncertainty .....	54
3.5    Conclusion .....	57
<b>Part Two: Applications .....</b>	<b>58</b>
<b>Chapter Four: Applying the Analysis to Modal Disagreement.....</b>	<b>59</b>
4.1    A puzzle about disagreement.....	59
4.2    Some attempted epistemic solutions to the puzzle .....	61
4.2.1    Denying the genuineness of disagreement .....	61

4.2.2	Redefining what is required for genuine disagreement.....	66
4.3	Other attempted epistemic solutions to the puzzle .....	70
4.3.1	Denying the problem of warrant.....	70
4.3.2	Redefining what is required for warrant.....	70
4.4	Dissolving the puzzle on a non-epistemic analysis .....	77
4.5	Truth-value judgments on a non-epistemic analysis .....	80
4.5.1	Truth-value judgments and the evaluation matrix.....	80
4.5.2	Binding and the evaluation matrix .....	83
4.6	An objection to a non-epistemic analysis .....	89
4.7	Conclusion .....	91
<b>Chapter Five: Applying the Analysis to Modal Belief .....</b>		<b>92</b>
5.1	A puzzle about belief reports.....	92
5.2	Belief reports without second-order commitments.....	93
5.3	Epistemic analyses with second-order commitments .....	94
5.3.1	Contextualists and relativists.....	94
5.3.2	Expressivists .....	96
5.4	Avoiding second-order commitments via anaphora .....	98
5.5	Challenges for an anaphoric analysis .....	100
5.5.1	Theoretical challenges .....	100
5.5.2	Empirical challenges .....	104
5.6	Belief reports on a non-epistemic analysis of <i>might</i> .....	109
5.6.1	Avoiding second-order commitments .....	109
5.6.2	Making sense of uncertainty implications.....	111
5.6.3	Comparison to epistemic analyses.....	113
5.7	An objection to a non-epistemic analysis .....	114
5.8	Conclusion .....	115
<b>Part Three: Objections.....</b>		<b>116</b>
<b>Chapter Six: The Meaning Difference Between <i>Might</i> and <i>Can</i> .....</b>		<b>117</b>
6.1	Two competing analyses.....	117
6.2	Kratzer's analysis.....	118
6.2.1	The continuity of Kratzer's analysis.....	119
6.2.2	The discontinuities in Kratzer's analysis.....	120
6.2.3	An implicit alternative to Kratzer's explicit analysis .....	124
6.3	On the intuitive motivation for Kratzer's analysis .....	126
6.3.1	Epistemic contradictions .....	126
6.3.2	Modal contradictions .....	127
6.4	On the theoretical motivation for Kratzer's analysis.....	130
6.4.1	The formal framework.....	130
6.4.2	The intensional assumption .....	133
6.4.3	Non-ambiguity without the intensional assumption.....	135

6.5	Revisiting the arguments for Kratzer’s non-ambiguity project .....	136
6.5.1	The explosion of meanings argument.....	137
6.5.2	The common kernel of meaning argument.....	138
6.5.3	The cross-linguistic argument .....	140
6.6	Conclusion .....	141
<b>Chapter Seven: The Temporal Interpretation of <i>Might</i> and <i>Can</i> .....</b>		<b>142</b>
7.1	Aspectual differences between <i>might</i> and <i>can</i> .....	142
7.2	Aspectual restrictions on <i>can</i> ’s complements.....	144
7.2.1	The standard analysis of the restrictions .....	144
7.2.2	The standard explanation(s) for the restrictions .....	144
7.3	Against the standard analysis of <i>can</i> ’s aspectual restrictions .....	148
7.3.1	Progressive and perfect complements .....	148
7.3.2	Bare complements .....	149
7.3.3	Circumstantial modal bases .....	152
7.4	Rethinking the temporal interpretation of <i>can</i> .....	152
7.4.1	Modal states.....	152
7.4.2	Representing modal states .....	154
7.4.3	Revisiting the representation of modal states .....	155
7.4.4	Aspectual operators and aspectual morphology .....	156
7.5	Revisiting the temporal interpretation of <i>might</i> .....	158
7.5.1	A puzzle about temporal orientation .....	158
7.5.2	Solving the puzzle about temporal orientation.....	160
7.5.3	Aspectual morphology revisited.....	162
7.5.4	A uniform semantics revisited.....	163
7.6	The counterfactual readings of <i>might</i> .....	164
7.6.1	An account in terms of aspect .....	164
7.6.2	An account in terms of tense .....	166
7.6.3	A second sort of counterfactual reading .....	169
7.6.4	Counterfactual readings and modal bases .....	170
7.7	Conclusion .....	171
7.7.1	Looking back .....	171
7.7.2	Looking forward.....	171
<b>Bibliography.....</b>		<b>174</b>



## Chapter One: Introduction

[I]t is precisely [the] confusion between *the inappropriateness of saying* certain things ... and the *possible truth* of such things which is a main source of subjectivist views about probability, certainty, morality, etc.

~ White (1975, 172)

This dissertation is about what are traditionally called *epistemic* uses of modals. These are uses of modals like *might* and *must* with which a speaker implies relative (un)certainty. With (1), for example, a speaker typically implies that she is relatively uncertain whether John is contagious, while with (2) she typically implies that she is relatively certain that he is.

- (1) John might be contagious.
- (2) John must be contagious.

It is standardly assumed that such (un)certainty is implied directly. Either the modals modify the force of the speaker's commitment to the proposition that John is contagious, or they express a relation between that proposition and the speaker's own information state. In this dissertation I am going to argue that this is not the case. I am going to argue that with so-called epistemic uses of *might* and *must* (un)certainty is implied only *indirectly*. Recognizing that this (un)certainty is implied indirectly is required, I argue, to make sense of how we explain the truth of the claims made with epistemic uses of *might* and *must*. It also makes sense of various properties of these uses which have both seemed anomalous and have been analyzed as such, often with far-reaching philosophical implications. The anomaly is only apparent, I argue, and depends on the assumption that a speaker who implies (un)certainty with the use of a modal must be doing so directly. But epistemic uses of *might* and *must* do not require far-reaching innovations to standard philosophical and semantic theory. All they require is for us to recognize that (un)certainty can sometimes be implied indirectly.

### 1.1 Epistemic uses of modals

With some uses of modals, a speaker implies that she is relatively (un)certain about a matter. But not all uses of modals are like this. While with (2) a speaker implies that she is relatively certain that John is contagious, with (3) her relative certainty about whether contagious people are quarantined is neither here nor there.

- (3) Contagious people must be quarantined.

Instead, the speaker simply expresses the requirement for contagious people to *be* quarantined, a so-called *deontic* use of a modal. My focus in this dissertation is on the uses of modals with which speakers do imply relative (un)certainty, uses that have traditionally been called *epistemic* (cf. Lyons 1977; Palmer 2001; Portner 2009).

There are two standard views about how speakers manage to imply relative (un)certainty with epistemic uses of modals. The first is that they do so because the modal directly modifies the force of their commitment to the proposition in question: that is, the proposition expressed by the modal's complement (cf. Halliday 1970; Palmer 2001; Schnieder 2010). On this view, the modal claims made with (1) and (2) express exactly the same proposition as the non-modal claim made with (4). They differ not with regard to their content but instead with regard to the force with which that content is put forth.

- (4) John is contagious.

With *must*, this content is put forth with strong, albeit less than fully assertoric, commitment (cf. Lyons 1977, 808–9; Westmoreland 1998). With *might*, it is put forth with very weak commitment. With the former sort of commitment, a speaker implies that she is highly, though not fully, certain that John is contagious. With the latter, she often implies that she is not at all certain whether he is.

The alternative view is that epistemic uses of modals make the content of the speaker's claim depend truth conditionally on an information state. They do so by expressing a relation between that state and the content of the modal's complement, commonly called the *prejacent* (cf. Kratzer 1981; DeRose 1991; Hacquard 2010; Bach 2011; von Stechow and Gillies 2011; MacFarlane 2014).<sup>1</sup> In its simplest version, this view would hold that that information state is the speaker's own, so that, relative to a given context, the truth conditions of the claims made with (1) and (2) would be, respectively, the same as those of the claims made with attitude reports more or less like (5) and (6).

- (5) I am uncertain whether John is contagious.  
(6) I am highly certain that John is contagious.

Of these two views, the truth-conditional one seems to fare better when we turn from modal sentences in main clauses to modal sentences in embedded ones, as in (7) and (8), where the modal sentence is embedded beneath a propositional attitude verb.

- (7) Mary thinks that John might be contagious.  
(8) Mary thinks that John must be contagious.

With the epistemic uses of *might* and *must* in (7) and (8) a speaker still implies relative (un)certainty. However, the (un)certainty is no longer her own but instead that of Mary, the reported attitude-holder. With (7) and (8), a speaker thus reports different attitudes from what she reports with (9).

---

<sup>1</sup> The relevant state had standardly been assumed to be one of knowledge, that is, a state comprising the propositions that were known, whence the name *epistemic*. Recently, it has been argued that this assumption is too strong: that propositions other than those that are known may also form part of the relevant information state (cf. Yalcin 2007; Tancredi 2007; Hacquard 2010; Kratzer 2012). However, the term *epistemic* has been retained, now used in a broader sense in the truth-conditional literature to refer to a modal's introducing truth-conditional dependency onto any sort of information state. I follow this practice throughout the dissertation.

- (9) Mary thinks that John is contagious.

She does *not* report the same attitude while simply introducing her own side comment on its content, as is possible for the epithet *dummy* in (10) and (11), and as we might expect to also be possible for *might* and *must* in (7)-(9) on a force-modifier view of those terms.

- (10) Mary thinks that that dummy John is contagious.  
 (11) Mary thinks that John is contagious.

Embedded occurrences of *might* and *must* have played a prominent role in debates about whether to adopt force-modifier or truth-conditional views of their epistemic uses (cf. Papafragou 2006; cf. also Hacquard and Wellwood 2012, and the references cited therein). While I cannot summarize that debate here, I think it is safe to say that it favors the truth-conditional view (cf. Swanson 2011, 257–60; Anand and Hacquard 2013, 7–14; MacFarlane 2014, 248–53). It is easier on a truth-conditional view to explain cases where epistemic *might* and *must* do not embed, which are relatively fewer anyway, than to explain, on a force-modifier view, those many cases where they do.

If a truth-conditional view seems to be required, the simplest epistemic version of it presented above is nonetheless subject to some very well-known difficulties. These have become central explananda in the development of alternative, increasingly more complicated truth-conditional views that have recently proliferated in the literature (cf., among others, Yalcin 2007; 2011; Anand and Hacquard 2009; 2013; Hacquard 2010; MacFarlane 2011; 2014).

One problem relates to *embedding phenomena*. While (7) and (8) are not used to report the same attitude as (9), they are also not used, respectively, to report the same attitudes as (12) and (13) (cf. Yalcin 2007; Hacquard 2010).

- (12) Mary thinks that she is uncertain whether John is contagious.  
 (13) Mary thinks that she is highly certain that John is contagious.

With (12) and (13), Mary's reported attitudes are about herself; with (7) and (8), they are about John. This difference is unexpected on a simple, truth-conditional view in which a speaker can express more or less the same contents with (1) and (2) as she does, respectively, with (5) and (6). In that case, we would expect the contents of the reported attitudes in (7) and (8) to be able to be the same as the contents of the reported attitudes in (12) and (13) too.

A second problem relates to *disagreement phenomena*. The speaker who responds to (1) or (2) with (14) or (15) seems to indicate disagreement with what has been said (cf. Stephenson 2007).

- (14) John can't be contagious.  
 (15) John doesn't have to be contagious.

However, if she responded to (5) or (6) in this same way, she would not seem to indicate disagreement, or at least not of the same sort. In the first case, she seems to deny what her interlocutor asserted; in the second, to accept it, but also to encourage her interlocutor to reconsider the attitude she has asserted herself as having.

These two problems do not depend on the choice of the attitude verb used to gloss the state of relative (un)certainly that a speaker implies with epistemic uses of *might* and *must*. Nor do they depend on the identification of the relevant information state as the speaker's either. Speakers who utter (7) or (8) do not seem to report attitudes about *any* particular information state. And hearers who utter (14) or (15) do not seem to deny that speakers who have uttered (1) or (2) are in any particular information state either. The problem is not *what* information state speakers imply with epistemic uses of *might* and *must* but instead *how* that state is implied.

The standard response to these and other problems has been to provide epistemic truth conditional theories that (i) are increasingly complicated; (ii) have far-reaching philosophical ramifications; (iii) make epistemic uses of modals anomalous; and (iv) do not provide full and principled empirical coverage of the relevant data anyway. Examples of such responses can be found in MacFarlane (2011; 2014) and Yalcin (2007; 2011).

MacFarlane (2014), for example, proposes to make sense of the disagreement phenomena by removing information states from content and then relativizing the truth of the claims made with epistemic uses of modals to contexts of assessment individuated by such states. The result is to make the content that is the *object* of disagreement to still be truth-conditionally dependent on information states without having the disagreement *itself* be about those states. The price is the introduction of a philosophically significant relativism about select modal claims. No one denies that a single modal sentence like (16) can be used in different contexts to make different claims with potentially different contents or truth values.

(16) John can jump 10 feet.

However, MacFarlane now proposes that, in an epistemic use of a modal, the truth value of a single *claim*, with a single content, can differ from one person and context of assessment to another. Nor does this relativism buy us everything we might want. For to make sense of certain responses to disagreement, MacFarlane is forced to concede that the claims made with epistemic uses of modals can sometimes have enriched contents that *include* information states, making the claims invariant in truth value relative to different information states and assessment contexts (2014, 258–60). This may not vitiate the proposal, but it does weaken its explanatory force.

Yalcin (2007), in turn, proposes to make sense of the embedding phenomena by having the information state already provided by the attitude verb be the state relative to which embedded epistemics express their epistemic modality. The result is to have an attitude report like (7), repeated here as (17), not report an attitude *about* any information state but instead a *sui generis* attitude toward the proposition that John is contagious. Yalcin (2011) proposes that this *sui generis* attitude is the state of mind reported by (18).

- (17) Mary thinks that John might be contagious.
- (18) Mary's belief state is compatible with John's being contagious.  
(And Mary is also sensitive to the question of whether John is contagious.)

The price of denying modal contents is a significant asymmetry between embedded epistemics and other modals. Embedded epistemics change the character, or strength, of the reported attitude: for example, from something like full belief to something more like mere suspicion. Other embedded modals, in contrast, change the content of the reported attitude. But the price we pay in positing such an asymmetry again does not buy us everything we might want. For to make sense of the fact that (19) and (20) are clearly *not* used to report the same state of mind, Yalcin is forced to concede that sometimes the relevant information state for embedded epistemics can be provided by context, rather than by the attitude verb, so that (19) is used to report the same attitude—with the same modal content—as (21) (2007, 1012–13).

- (19) Mary doesn't know whether John might be contagious.
- (20) Mary's knowledge state is *not* compatible with John's being contagious.  
(And Mary is also sensitive to the question of whether John is contagious.)
- (21) Mary doesn't know whether some contextually relevant information state is compatible with John's being contagious.

Again, this may not vitiate the proposal, but it does weaken its explanatory force.

In this dissertation, I argue that there is a simple, attractive alternative to these (and other) increasingly complicated epistemic truth conditions. It is to adopt the view that the claims made with epistemic uses of *might* and *must* have modal truth conditions that are *non-epistemic*, and that the (un)certainty that is implied with these uses is implied only indirectly.

Of course, I do not think that (un)certainty can only ever be implied indirectly, or that we can never make modal claims with epistemic truth conditions. In fact, I think we routinely do make such claims: for example, with modal adjectives like *plausible*. Rather, it is just that I think we never make such claims with epistemic uses of *might* and *must*. From a semantic standpoint, these modals, and the claims made with them, are systematically non-epistemic.

I argue for this conclusion about *might* and *must* on the basis of new data about how we explain the truth of modal claims. These data give us independent reason to abandon what I will call the *epistemic assumption*: the assumption that the claims made with epistemic uses of modals must have epistemic truth conditions. I show how rejecting this assumption dissolves the problems for epistemic truth conditional views in a principled way that makes epistemic uses of modals continuous with other uses and does not require significant innovations to standard philosophical and semantic theories either. And I show how we can make sense of (un)certainty being implied with the use of a modal without that modal's either altering the force of a speech act or making its content depend truth-conditionally on an information state. I preview these main points below, before outlining the structure of the dissertation as a whole.

*Independent Motivation:* Modal facts are often contingent, allowing us to explain what makes the content of a given modal claim true by appeal to relevant contingent features of the world. To explain what makes the content of a *deontic* claim true, for example, we can appeal to the contingent nature of the law, as in (22).

- (22) a. Wisconsinites can shoot pigeons during hunting season.
- b. You're right. But the only reason they can is that their legislature is still controlled by the gun lobby.

And to explain what makes the content of an *epistemic* claim true, we can appeal to the contingent nature of a speaker and her interlocutors' current information state, as in (23).

- (23) a. It's plausible that John is contagious.
- b. You're right. But perhaps the only reason it's plausible is that we still don't know about his test results.

However, we cannot make this same appeal with all apparently epistemic uses of modals, as (24) demonstrates. Not knowing John's diagnosis may explain why it is plausible that he is contagious, but it apparently *cannot* explain why he might be so.

- (24) a. John might be contagious.
- b. You're right. #But perhaps the only reason he might be is that we still don't know about his test results.

The best explanation of this fact, I argue, is that, unlike the content of the claim made with *plausible* in (23), the content of the claim made with *might* in (24) does not turn out to have epistemic truth conditions. Instead, it has circumstantial ones. Its truth depends on how John's being contagious stands with regard to circumstances, rather than on how his being so stands with regard to an information state (for the circumstantial-epistemic distinction, cf. Kratzer 1981; 1991). As such, it is to the contingent nature of the circumstances, rather than of any information state, that we must appeal to explain its truth, as in (25).

- (25) a. John might be contagious.
- b. You're right. But the only reason he might be is that he still works in that old hospital ward.

*Dissolution of Outstanding Problems:* If the truth of the claim made with ((25)a) *does* depend on how John's being contagious stands with regard to the circumstances, then it follows automatically that disagreement with that claim is not disagreement about an information state. Similarly, it follows automatically that to believe the content of that claim to be true is not to have a belief about an information state. At the same time, if believing this content to be true is believing that the circumstances bear a *particular* modal relation to the proposition that John is contagious, this will be distinct both from believing that they bear other modal relations to that proposition

and from believing that proposition itself. So we also readily make sense, as force-modifier views cannot, of why (7)-(9) report different attitudes on the part of the reported attitude-holder.

*Implied (Un)certainty*: Circumstances can affect the objective probability of a proposition. John's working in the worst hospital ward, for example, increases the objective probability that he is contagious. His having a powerful autoimmune system decreases it. And so on. Say for the sake of illustration that the claims we make with so-called epistemic uses of modals are about how circumstances affect the objective probability of a proposition.<sup>2</sup> A speaker who makes such a claim can often reasonably be assumed to be interested in the *actual* truth value of that proposition, and to have made the claim she did only because she was less than completely certain about it. With certain uses of modals, speakers could thus imply relative (un)certainty even without those modals modifying the force of a speech act or making its content truth-conditionally dependent on an information state. Relative (un)certainty would be implied in virtue of the fact (i) that these modals make the content of a speech act depend truth-conditionally on how circumstances affect the objective probability of a proposition; and (ii) that speakers can generally be assumed to have a purely extrinsic interest in this matter, one that depends solely on their unsatisfied interest in the actual truth value of that proposition.

## 1.2 Outline of the dissertation

This dissertation is composed of three distinct parts, of two chapters each. The first part deals with fundamentals: with the argument against epistemic analyses of *might* and *must* from the data about how we explain the truth of modal claims, and with the development of a framework for a non-epistemic analysis. The second part deals with applications of this analysis to the disagreement and embedding phenomena that have proven so difficult for epistemic analyses to accommodate. The third part deals with objections that arise within the standard Kratzerian framework for modal semantics: specifically, from its use of the epistemic/circumstantial distinction to differentiate *might* from *can* (cf. Kratzer 1981; 1991), and also to account for the different temporal properties of the two (cf. Thomas 2014; Rullmann and Matthewson 2018).

### *Part One: Fundamentals*

In *Chapter Two*, I present data about how we explain the truth of modal claims. I show that with so-called epistemic uses of modals there is a robust divide between those uses that license subsequent epistemic explanations and those that do not. This divide cross-cuts a number of important distinctions, including distinctions in modal force (possibility vs. necessity), in modal polarity (negated vs. non-negated), and in the grammatical category of a modal expression (auxiliary vs. adjective). This divide does *not* seem to persist when apparently epistemic uses of modals are restricted by epistemic modifiers, like *given what we know*. (These modifiers have been alleged to make explicit the otherwise implicit information state relevant for the epistemic use.) However, I show that such modifiers license epistemic explanations independently of whether the use they restrict is alleged to be epistemic. I also show that there is no

---

<sup>2</sup> This is not the view I ultimately adopt, but I will often use it for illustrative purposes.

similar divide for other uses of modals in the sorts of explanations licensed by their allegedly implicit and explicit counterparts.

Having presented these data, and having shown that they cannot be attributed to anything about the particular modal concepts allegedly expressed by *might* and *must*, I critically review existing theories. I argue that we cannot make sense of the data by assuming that the relevant uses of these modals function either (solely) to modify the force of a speech act or to make its content truth-conditionally dependent on an information state. Truth-conditional views incorrectly predict epistemic explanations to be available for *might*- and *must*-claims, even after the role of information states is fundamentally re-envisioned, as with the removal of such states from content to the evaluation matrix (cf. MacFarlane 2014). Force-modifier views, in contrast, correctly predict such explanations to be unavailable, but also incorrectly predict non-epistemic explanations to be unavailable too. Neither of these standard views for how speakers imply (un)certainly with the relevant epistemic uses of modals makes good sense of our data then.

In *Chapter Three*, I propose we adopt a non-epistemic analysis for those epistemic uses of modals that do not license epistemic explanations. I focus primarily on *might* and on the idea that it could express *circumstantial* possibility. Without committing to any particular analysis of circumstantial possibility (though I allude to one in Chapter Seven), I show how we can make ready sense of the data from the previous chapter whatever circumstantial analysis we adopt. For if circumstances leave open (or close off) certain possibilities, then different ones could have been left open (or closed off) if circumstances had been different. This allows us to appeal to contingencies of the circumstances to explain what makes claims about these sorts of possibilities true, though understandably *not* to contingencies of information states.

Our data provide good reason to think that the contents of possibility judgments are truth-conditionally independent of the states of those who make them. However, the analytical difficulty is to effect this separation. I show how we can do so, just as we have learned to do for actuality judgments, if we allow the contents of possibility judgments to be general, or independent of specific sets of circumstances. This is something I argue we should do for the claims made with uncontested circumstantial modals anyway, so nothing is required to make sense of so-called epistemic uses of modals, on my view, that is not also already required for other modals. I show how to formalize this view within an adapted version of a standard Kratzerian framework for modal semantics (cf. Kratzer 1981; 1991). And I also respond to objections intended to show that any non-epistemic analysis cannot make sense of the non-triviality of *might*-claims, or of a general requirement for speaker uncertainty when making them. In both cases, I argue, a non-epistemic analysis in fact fares better in accounting for the relevant data than epistemic analyses do, thereby independently reinforcing the argument against such analyses from Chapter Two.

### *Part Two: Applications*

In *Chapter Four*, I apply a non-epistemic analysis to the disagreement phenomena. The basic problem for epistemic analyses is that genuine and warranted disagreement about what might be the case seems to be possible, even between complete strangers. Yet it is not clear how this could be so if that disagreement were about an information

state. For complete strangers could not engage in warranted disagreement about each other's information states. And if their claims were exclusively about their own states, then the disagreement would not be genuine. I argue that responses to this problem intended to show that the disagreement is not genuine do not succeed (cf. von Fintel and Gillies 2008; Schaffer 2011). Something similar holds for the response that the disagreement is genuine, of a sort, but that it does not involve assertion and denial of the same content (cf. Huvenes 2015; Khoo 2015). Alternative responses focus on the warrant for epistemic disagreement, either denying that a special problem of warrant arises (cf. Yanovich 2014) or adopting custom-made theories of that warrant (cf. von Fintel and Gillies 2011; MacFarlane 2014). Neither response is compelling, I argue. I conclude there is no good epistemic solution to the problem of apparently genuine and warranted disagreement between strangers about what might be the case.

Yet such disagreement, I show, turns out to be analogous to disagreement about what can happen. For the latter, there is no problem in explaining how genuine and warranted disagreement between strangers is possible. This is because disagreement here is about the potentials that the *circumstances* leave open. And while speakers' different information may lead them to make different claims, that information does not enter into either the content of, or the intended truth conditions for, those claims. Different information thus does not threaten the possibility of genuine disagreement, and so also does not require custom-made accounts of warrant for that disagreement. As such, if we analyze *might*, like *can*, as expressing circumstantial modality, we can make straightforward sense of the disagreement phenomena.

In *Chapter Five*, I apply the non-epistemic analysis to the embedding phenomena, where the basic problem is more or less the same. The attitudes reported when modals like *might* are embedded do not seem to be *about* information states (cf. Yalcin 2007). This has led some researchers to propose that the relevant information state for an embedded epistemic is provided directly by the attitude verb, the embedded modal simply expressing the compatibility of its prejacent with the content of this state (Yalcin 2007; 2011; Hacquard 2010; Anand and Hacquard 2009; 2013). The proposal works best for belief reports, for which the prejacent often is compatible with the content of a reported belief-holder's belief state. However, I show that not all attitude verbs, nor even all epistemics, could interact in this specialized, idiosyncratic way. Further, even for belief verbs, there are exceptional cases in which compatibility is not required. And accounting for these cases undermines the explanatory force of the proposal. A new explanation of these attitudes reports is in order.

On a non-epistemic analysis of *might*, we would never have expected the relevant attitudes reported to be about information states in the first place. As such, we are free to maintain a standard view of the interaction of attitude verbs and embedded modals. For both epistemic and non-epistemic modals alike, and across all attitude verbs, we can take embedded modals to uniformly contribute to the content of the reported attitude, as opposed to changing its character. As for those exceptional belief reports, which undermine the explanatory force of epistemic analyses, these turn out to be a special case of a broader phenomenon involving *all* circumstantial modals. Generally, circumstantial modal beliefs are formed and reported on the basis of all one's beliefs about the circumstances. Yet for pragmatic reasons modal beliefs formed on the basis of a subset of the presumed circumstances may sometimes be reported too. This latter

sort of belief, while continuous with the former, does not require the compatibility of the belief state with the prejacent. Hence, both the standard cases and the exceptional ones can be accommodated without requiring anything special from the attitude verb itself, unlike on the epistemic analysis.

*Part Three: Objections*

In *Chapter Six*, I respond to an objection that arises from the standard, Kratzerian analysis of the meaning difference between (26) and (27) (cf. Kratzer 1981; 1991).

- (26) John might easily be a military man.
- (27) John can easily be a military man.

On this analysis, (26) and (27) differ in how they characterize the facts relative to which the prejacent is said to be a possibility. For *might* the facts are characterized epistemically, for *can* non-epistemically, a difference that would be eliminated on my non-epistemic account of *might*.

Yet I show that Kratzer's analysis is problematic independently of this account. Its intuitive motivation comes from so-called *epistemic contradictions* (cf. Yalcin 2007), which occur when a prejacent is denied in conjunction with a *might*-claim, as in (28), but not when it is denied in conjunction with a *can*-claim, as in (29).

- (28) Hydrangeas aren't growing here, #but they might be growing here.
- (29) Hydrangeas aren't growing here, but they can grow here.

And its theoretical motivation comes from the Kratzerian attempt to account for the phenomenon of *modal flexibility* (or the ability of a single modal to be used to express multiple "flavors" of modality: epistemic, deontic, etc.) without having to posit *modal ambiguity* (or multiple lexical entries for that modal) (cf. Kratzer 1977). However, I show that both motivations are ultimately ill-served by Kratzer's analysis. So-called epistemic contradictions arise not just with allegedly epistemic uses of *might* but also metaphysical ones, undermining the Kratzerian explanation of *might*'s difference with *can* in this regard. Further, for a Kratzerian non-ambiguity project to succeed, modal flavors must be differentiated by something other than the intensional characterization of facts and norms (cf. Nauze 2008).

I thus propose an alternative, developed more fully in *Chapter Seven*: while *might* is used to attribute relative modal properties to *propositions*, *can* is used to attribute them to *individuals*, *situations*, *locations*, etc. The crucial meaning difference is thus not in the *facts* that these modals express relations to, but instead in the *relations* they express to these facts (for inspiration for this account, cf. Palmer 2001).

In *Chapter Seven*, I respond to a second objection related to a difference between *might* and *can*, this time in apparent temporal properties. While *might* freely accepts complements in the progressive and perfect forms, as in (30), *can* does not do so, as (31) shows.

- (30) John might {speak/be speaking/have spoken} French.
- (31) John can {speak/#be speaking/#have spoken} French.

A fairly standard explanation for this difference relates it to the difference between epistemic and circumstantial modality (Werner 2006; Thomas 2014). Circumstantial possibility, expressed by *can*, is supposed to be inherently oriented toward the future. Since the progressive and the perfect orient a possibility toward the present and the past, respectively, they are thus incompatible with *can*. As for their compatibility with *might*, this is evidence that *might* expresses epistemic possibility, the only flavor of possibility able to be oriented toward either the present or the past.

I argue that this explanation for the contrast between (30) and (31) is inadequate. *Can* sometimes does accept complements in the progressive and perfect forms. And sometimes the clearly non-epistemic possibilities it expresses can only be analyzed as being oriented toward the present. Additionally, for the future-oriented possibilities it expresses, this orientation could not be attributed to anything about the circumstances to which those possibilities are relativized. Temporal orientation does not depend on the epistemic/circumstantial distinction.

A better distinction to help make sense of the contrast in (30) and (31) is the one proposed in Chapter Six. The progressive and the perfect help to denote propositions, by relating event descriptions to particular reference times in particular ways. And *might* is just used to attribute modal properties to propositions, which explains why it freely accepts perfect and progressive complements. Yet when it comes to the event descriptions used to specify the modal properties of individuals, etc., often nothing turns out to be gained by imposing the particular relations to particular reference times that that the progressive and the perfect would. This explanatory difference in the sorts of entities to which modal properties are attributed is something that I show how to formalize in this chapter.

In addition to its main argument, Chapter Seven also concludes the dissertation as a whole, indicating open questions and pointing out avenues of research it opens up.

## **Part One: Fundamentals**

## Chapter Two: A Problem for Epistemic Analyses

[M]odal statements of the sort we have considered so far are contingent, they are neither necessarily true nor necessarily false. That Jockl must have been the murderer (in view of what we know) is a fact of our world, but it is not a necessary truth. Had our knowledge been different, it might not have implied anymore that Jockl is the murderer.

~ Kratzer (1991, 641)

### 2.1 The basic problem

The truth of claims about our information should be contingent, as Kratzer suggests, dependent on the contingent nature of our information itself. Yet the truth of the claims made with some apparently epistemic uses of modals does not seem to be contingent in this way. The truth of the claims made with ((32)a) and ((33)a), for example, cannot be explained in terms of the contingent nature of the speaker and her interlocutor's current information state, whether appeal is made to the absence of a particular piece of information, as in ((32)b), or to its presence, as in ((33)b).

- (32) a. John might be contagious.  
b. You're right. #But {perhaps} the only reason he might be is that we still don't know about his test results.
- (33) a. John might speak French.  
b. You're right. #But the only reason he might is that we know about his semester in Paris.

Not all epistemic uses of modal predicates are like this. The claims made with ((34)a) and ((35)a), for example, support the epistemic explanations that the claims made with ((32)a) and ((33)a) did not.

- (34) a. It's plausible that John is contagious.  
b. You're right. But {perhaps} the only reason it's plausible is that we still don't know about his test results.
- (35) a. It's plausible that John speaks French.  
b. You're right. But the only reason it's plausible is that we know about his semester in Paris.

As we might expect, claims made with propositional attitude vocabulary also support the same sort of epistemic explanations.

- (36) a. We suspect that John is contagious.  
b. You're right. But {perhaps} the only reason we suspect as much is that we still don't know about his test results.

- (37) a. We suspect that John speaks French.
- b. You're right. But the only reason we suspect as much is that we know about his semester in Paris.

It is thus both unexpected and puzzling that claims made with some apparently epistemic uses of modals should not support epistemic explanations.

In this chapter, I will argue that we cannot resolve this puzzle given the traditional assumptions about apparently epistemic uses of modals. We cannot make sense of the fact that some apparently epistemic uses of modals do not license epistemic explanations, that is, if we assume that those uses must either modify the force of a speaker's commitment to her speech act or make the content of that speech act truth-conditionally dependent on an information state. In Section 2.2 of this chapter, I first expand and clarify the scope of the relevant data (though my focus afterwards will primarily be on *might* and *plausible* as paradigmatic of this data set). In Sections 2.3 and 2.4, I argue against attempted force-modifier and truth-conditional explanations of the data. And in Section 2.5, I conclude by showing how the data point to a non-epistemic analysis for those apparently epistemic uses of modals that do not license epistemic explanations. I develop the general outline of a non-epistemic analysis in the following chapter.

Before anything else, however, a simplifying note on the data. The examples used so far have made explicit that what is being explained is also being endorsed. However, other continuations less explicit in this regard do not differ in their relative acceptability as responses to the claims made with apparently epistemic uses of modals. (38) and (39) are marked in response to the *might*-claim made with ((33)a), for example, but not in response to the *plausible*-claim made with ((35)a). Similarly, ((40)a) is marked while ((40)a') is not.

- (38) Only because we know about his semester in Paris.
- (39) That's {only} because we know about his semester in Paris.
- (40) It's only because we know about his semester in Paris that ...
  - a. ... #John might speak French.
  - a'. ... it's plausible that John speaks French.

The same holds for continuations that do not aim to provide actual explanations but only to speculate on possible ones, as in (41), (42), and (43).

- (41) Could that be because we know about his semester in Paris?
- (42) Perhaps that's because we know about his semester in Paris.
- (43) That's probably because we know about his semester in Paris.

It also holds for continuations that explicitly do *not* endorse what is being explained, as in (44), but instead provide only hypothetical explanations.

- (44) I'm not so sure I agree.
  - a. #But *if* John might speak French, it's only because we know about his semester in Paris.

- a'. But *if* it's plausible that John speaks French, it's only because we know about his semester in Paris.

Finally, the relative acceptability of a continuation is independent of its use of anaphora, and so independent of any potential anaphoric targeting of the *prejacent*, or the content of the modal's complement.<sup>3</sup>

The uniformity of the contrast between these continuations suggests that what is at stake is what would constitute even a *potential* explanation for the content of the claims that could be made with modals like *might* and *plausible*. Note that the problem is not the former sort of claims do not support any sort of explanations at all. (45) and (46) clearly show the contrary.

- (45) a. John might be contagious.  
 b. You're right. But the only reason he might be is that he still works in that old hospital ward.
- (46) a. John might speak French.  
 b. You're right. But the only reason he might is that he spent his semester abroad in Paris.

Instead, the problem is that even though the *appropriateness* of epistemic uses of modals like *might* can depend on the contingencies of an information state, as (47) shows,<sup>4</sup> the *truth* of claims made with such uses does not seem to (be able to be explained in terms of such contingencies), as (48) shows.

- (47) John might be contagious. We can't rule out that the test results aren't positive.
- (48) John might be contagious. #But the only reason he might be is that we can't rule out that the test results aren't positive.

In this regard, apparently epistemic modal claims are analogous to non-modal claims, as (49) and (50) show, something that is unexpected on epistemic truth-conditional analyses.

- (49) John is contagious. We can establish that much.
- (50) John is contagious. #But the only reason he is is that we can establish that much.

---

<sup>3</sup> The availability of the *prejacent* as a target for anaphora is a recurring theme in the literature (cf. von Stechow and Gillies 2008; Portner 2009; Dowell 2011; Braun 2012; Huvenes 2015). It tends to be raised in contexts in which the aim is to explain away the availability of certain responses to *might*-claims that appear to provide counterexamples to one's own theory of epistemic uses of modals. Yet it is never shown that these same responses, including ascriptions of falsity, are *not* also available in response to claims made with *might*'s polar opposite, *cannot*. Insofar as they are, *prejacent*-targeting hypotheses are not to be taken seriously. In responding to *John might be contagious* by saying *That's false*, I could perhaps be construed as denying the *prejacent*. In responding to *John can't be contagious* by saying *That's false*, there is no way I can (Alexander Williams, p.c.).

<sup>4</sup> The example is due to an anonymous reviewer for *Semantics and Pragmatics*.

Given the uniformity of the contrast that illustrates this basic problem, I thus move freely between continuations that aim to provide actual, possible, or instead merely hypothetical explanations.

## 2.2 The full scope of the problem

### 2.2.1 Expanding the data set

The data from Section 2.1 may make it seem that the divide between apparently epistemic uses of modals is a grammatical one.<sup>5</sup> Only modal adjectives like *plausible* license epistemic explanations; modal auxiliaries like *might* do not. However, as (51) and (52) show, the divide is not grammatical. Some apparently epistemic uses of modal adjectives, just like some apparently epistemic uses of modal auxiliaries, also do not license epistemic explanations, even though they do license non-epistemic ones.

- (51) a. It's possible that John is contagious.  
b. #Is the reason it's possible that he hasn't told us about his test results?  
b'. Is the reason it's possible that he works in that old hospital ward?
- (52) a. It's possible that John speaks French.  
b. #Is the reason it's possible that he told us about his semester in Paris?  
b' Is the reason it's possible that he spent his semester abroad in Paris?

The divide is also not limited to modals that express possibility (or possibility-related) concepts. Necessity (and necessity-related) concepts show the same divide, as (53)-(54) show. (Imagine that John is known to be silent on a matter only when he has bad news.)

- (53) a. John must be contagious.  
b. #Is that because it's been two weeks and we still haven't heard from him about his test results?
- (54) a. It's obvious that John is contagious.  
b. Is that because it's been two weeks and we still haven't heard from him about his test results?<sup>6</sup>

This divide persists with negated possibility modals, as (55)-(57) show.

- (55) a. John can't be the *only* hospital worker to be contagious.  
b. #You're right. And the rumors of a general outbreak there are the reason he can't be.

---

<sup>5</sup> For an account of when grammatical differences could matter, cf. Hacquard (2013; 2016).

<sup>6</sup> Unfortunately, the dual of *possible*—*necessary*—is generally thought not to have (what have been analyzed as) epistemic readings (cf. Palmer 1986, 58). Thus, it is necessary to rely on possibility modals to show that the general divide does not seem to be a grammatical one.

- (56) a. It's not possible that John is the *only* hospital worker to be contagious.  
 b. #You're right. And the rumors of a general outbreak there are the reason it's not possible.
- (57) a. It's not plausible that John is the *only* hospital worker to be contagious.  
 b. You're right. And the rumors of a general outbreak there are the reason it's not plausible.

And the divide persists with negated necessity modals too, as (58)-(59) show.

- (58) a. Antonio doesn't have to be dead.  
 b. #Is the rumor about the boss's feeling merciful the reason he doesn't have to be dead?
- (59) a. It's not obvious that Antonio is dead.  
 b. Is the rumor about the boss's feeling merciful the reason it's not obvious that he's dead?

In all these cases, uses of modals that do not license epistemic explanations do license non-epistemic ones. Thus, (60) is fine as a continuation to ((53)a), (61) as a continuation to ((55)a), and (62) as a continuation to ((58)a).

- (60) Is that because he works in that old hospital ward?  
 (61) That's because of the general outbreak there.  
 (62) Is the boss's feeling merciful today the reason he doesn't have to be dead?<sup>7</sup>

Hence, the problem of epistemic uses of modals not licensing epistemic explanations is not limited in scope to non-negated modals, to possibility modals, or to modal auxiliaries, though I will focus almost exclusively on such modals going forward.

### 2.2.2 Clarifying what is at stake

A common first response to the data from Section 2.1 is to assume that it is something about the particular epistemic concept expressed by *might*—epistemic possibility—that makes attempted epistemic explanations bad. This assumption of the uniqueness of epistemic possibility seems to be implicit in MacFarlane (2011)'s response to the badness of examples like (63).

---

<sup>7</sup> These sorts of non-epistemic explanations are also available in response to those epistemic uses of modals that *do* license epistemic explanations. This is not an issue. Propositions are plausible, obvious, etc., to individuals because of the other propositions that they believe to be true. However, if a speaker believes a proposition to be true, and also believes that its being true stands in some sort of causal relation to her believing it to be so, then she can also attribute one proposition's being plausible, obvious, etc., to some other proposition's being true. Epistemic relations are parasitic on non-epistemic ones and can be explained in both sorts of terms. Non-epistemic relations are not. Thus, the issue is not whether an allegedly epistemic claim supports a non-epistemic explanation but instead whether it *fails* to support an epistemic one.

- (63) #It isn't possible that John is the murderer, but if no one had looked in this desk, it would have been.

On the basis of such examples, MacFarlane concluded that “counterfactual changes in what we know do not induce counterfactual changes in what is epistemically possible” (2011, 169). This was not an explanandum for MacFarlane, but rather a brute fact. And yet MacFarlane’s conclusion was premature, as (63)’s contrast in acceptability with (64) and (65) shows.

- (64) It isn't *epistemically* possible that John is the murderer, but if no one had looked in this desk, it would have been.  
(65) It isn't a live possibility {for us} that John is the murderer, but if no one had looked in this desk, it would have been.

Provided that we are explicit that it is epistemic possibility that we are talking about, then counterfactual changes in what we know *can* induce counterfactual changes in what is epistemically possible. It does not seem to me a very promising project, then, to assume that something about the particular epistemic concept allegedly expressed by *might* accounts for its failing to license epistemic explanations.<sup>8</sup> That said, let me respond to four potential conceptually-based reasons for why *might* might not license epistemic explanations.

First, it could be that an individual or group for whom something is an epistemic possibility is not in a position to reasonably speculate—much less explain—why this is so. But this is clearly false. Say that I know that conclusive tests have been run to determine whether John is contagious but have not yet been told their results. Then I know, and can also explain, why it is epistemically possible for me that he is so (cf. Teller 1972; DeRose 1991). (66) and (67) also demonstrate that explanations of epistemic possibilities, or of uncertainties, are perfectly conversationally appropriate.

- (66) It's an {epistemic/live} possibility for us that John is contagious. But {perhaps} that's only because we haven't asked him his diagnosis.  
(67) We don't know whether John is contagious. But {perhaps} that's only because we haven't asked him his diagnosis.

Second, actuality entails possibility, and yet the explanations we have considered could appear to implicitly deny that this is so. Say that we grant that the only reason it is possible that John is contagious is that we have not heard his test results. If we did hear them, and found out that he *was* contagious, surely there would not cease to be the possibility of his being so. So perhaps epistemic explanations are bad in response to *might*-claims because they appear to suggest otherwise. But this possibility is also clearly false. For when we are explaining why something is an epistemic possibility for us, we are explaining why it is a *mere* epistemic possibility. This is what accounts for the standalone acceptability of (66), and also for the acceptability of (68) as a continuation to it.

---

<sup>8</sup> Unless we are also to assume that *might* expresses a *sui generis* epistemic concept for which there is neither good paraphrase nor good analysis.

(68) If we had, we would know for sure one way or the other.

Third, many factors often contribute to something's being an epistemic possibility, and so it could be that epistemic explanations are bad for their suggestion that only a single factor does. Yet many factors contribute to something's being plausible too, and epistemic explanations are not bad in response to *plausible*-claims for that reason. Additionally, this response could not make sense of the acceptability of (66), nor could it make sense of the fact that replacing *only* with *mainly* or *partially* in any of the examples above leaves their relative acceptability unchanged.

Finally, it has sometimes been suggested that epistemic possibility can involve a subjective element (cf. Lyons 1977; Palmer 2001). A subjective/objective distinction for uses of *might* was first tentatively introduced by Lyons (1977) and has since subsequently been enlisted to explain *might*'s distribution and/or interpretation in conditionals (Papafragou 2006), with regard to quantified subjects (Tancredi 2007), and beneath so-called attitudes of acceptance like *believe* (a doxastic attitude of acceptance) and *claim* (a proffering one) (Anand and Hacquard 2009). The distinction has not always been very well understood: sometimes it seems almost to correlate with an epistemic/non-epistemic distinction, other times to crosscut it (cf. Lyons 1977; Papafragou 2006).<sup>9</sup> Yet however the distinction is to be understood, if the subjective does not admit of explanation, or at least not in conversation, then perhaps this distinction can be enlisted to explain why *might*-claims do not support epistemic explanations (cf. Papafragou 2006 for an argument that subjectivity places restrictions on acceptable discourse moves).

In response, I would note that while subjectivity may have a role to play in explaining our data, it cannot explain that data by itself. Say that the subjective correlates with the epistemic. Then no epistemic possibility claim, however expressed, should support an epistemic explanation. Yet some do, as (66) shows. Now there may be something about the different *ways* in which subjectivity is expressed in these different cases that explains the difference in admissible explanations (cf. Section 2.3). But then it is not any subjective element to epistemic possibility that explains our data but instead how that subjective element is expressed.

Say instead that the subjective/objective distinction crosscuts the epistemic/non-epistemic one. It would then seem that this distinction could be drawn not just for epistemic possibility but also for other epistemic concepts too, including plausibility. Yet while marking this distinction may constrain the range of explanations that *plausible*-claims support, as in (69) and (70), it does not eliminate them altogether.

- (69) a. It's {subjectively} plausible that John has lung cancer.  
b. Only because you haven't seen his x-rays.

---

<sup>9</sup> Lyons writes, in introducing the distinction, that it "is not a distinction that can be drawn sharply in the everyday use of language; and its epistemological justification is, to say the least, uncertain. It is also difficult to draw a sharp distinction between what we are calling objective epistemic modality and alethic modality" (1977, 797). Gagnon and Wellwood (2011) also seem to express skepticism about the distinction, writing about the "absence of an explicit theory of subjective versus objective modality" and claiming that "adequate tests for subjective versus objective modality are difficult to assess" (2011, 41).

- (70) a. It's {objectively} plausible that John has lung cancer.  
b. Only because his x-rays are indecipherable.

In particular, while unmarked *plausible*-sentences license explanations appealing either to what we might call subjective epistemic facts, as in (69), or to objective ones, as in (70), sentences marked for subjectivity *do* continue to license the former even though they do not license the latter (and vice versa for sentences marked for objectivity). So subjectivity does not seem to eliminate the possibility of acceptable epistemic explanations, not even in conversation. Unmarked *might*-sentences, in contrast, license neither explanations that appeal to subjective epistemic facts nor explanations that appeal to objective ones, as (71) shows.

- (71) a. John might have lung cancer.  
b. #Only because you haven't seen his x-rays.  
b'. #Only because his x-rays are indecipherable.

Hence, a subjective/objective distinction, insofar as it is well-founded, does not seem to explain our data, or at least not by itself.

To conclude: nothing about the concept of epistemic possibility *per se* explains the badness of the epistemic explanations considered so far. What is at stake is thus not the concept of epistemic possibility but instead how that concept is expressed in epistemic uses of *might*. I now turn to the two main types of proposals in the literature, both of which assume that it is expressed directly, and argue that neither can provide a solution to the problem of epistemic uses of *might* not licensing epistemic explanations.

## 2.3 Potential force-modifier solutions

### 2.3.1 Traditional force-modifier accounts

Of the two standard views of epistemic uses of modals, the force-modifier one seems to provide the more immediately promising solution to our problem. Proponents of force-modifier views, recall, propose that some epistemic uses of modals function solely to modify the force of a speech act (cf. Halliday 1970; Palmer 2001; Schnieder 2010). They do not introduce any truth-conditional dependency on an information state into the content of a speech act but instead indicate a non-assertoric commitment to that content, the prejacent, on the part of a speaker. Given such a view, we can readily make sense of why certain epistemic uses of modals do not license epistemic explanations, provided we assume those explanations target speech act content.

First, the prejacent considered so far have had non-epistemic truth conditions, yet the explanations have assumed epistemic ones, thereby committing a sort of category error. Second, even apart from a category error, the prejacent have explicitly *not* been asserted. They have not been (proposed to be) incorporated into the common ground but are rather still at issue (cf. Stalnaker 1978). As such, any non-hypothetical explanation would be out of place.

As for the assumption that the relevant explanations target the content of a speech act, this is entirely reasonable and explains why ((72)b) is odd in response to ((72)a).

- (72) a. John is contagious.  
 b. #Only because you haven't read today's newspaper.

The speaker's *believing* that John is contagious is not part of the content of her claim (cf. Frege 1956; Moore 1962), even though this attitude (or commitment) is implied by the assertoric force with which that content was put forth. This lends support to the force-modifier view that the attitude of (un)certainty implied with the epistemic use of a modal is not part of the content of a modal claim either.

Attractive as it initially is, however, a force-modifier analysis of our data faces an immediate problem. For some apparently non-hypothetical explanations do not seem out of place in response to *might*-claims, as (73) shows, though force-modifier analyses would seem initially to predict that they should.

- (73) a. John might be contagious.  
 b. Only because he still works in that old hospital ward.

To make sense of examples like (73), proponents of force-modifier views would have to hold that the explanation is implicitly hypothetical, explaining what would make the prejacent true, assuming that it is. However, in some cases the prejacent can also be contrasted with the target of an explicitly hypothetical explanation, as in (74), suggesting that the two, in fact, are not always the same.

- (74) I seriously doubt that John is contagious. Still, if he even might be, it's only because he works in that old hospital ward.

How could the proponent of a force-modifier view make sense of this contrast? She could not take the target of explanation to be some modal content of the antecedent over and above the prejacent, for she denies that there is any such content. Instead, she must presumably take it to be something like the first speaker's non-assertoric commitment to the prejacent, or the appropriateness of such a commitment. However, if something like speaker commitment is available as a target of the hypothetical explanation in (74), then the force-modifier account of our data comes undone. For speaker commitment could just as well be explained in epistemic terms as in non-epistemic ones, as (75) and (76) show.

- (75) You only say that because we still don't know about his test results.  
 (76) You only say that because he still works in that old hospital ward.

And yet as we already know, and as (77) once more demonstrates, when it comes to *might*-claims, epistemic and non-epistemic explanations are generally not on par.

- (77) I seriously doubt that John is contagious. #Still, if he even might be, it's only because we don't know about his test results.

Yet the proponent of a force-modifier view could no longer provide a reason why, having granted that a speaker's commitment, just as well as the prejacent, is available as a target of the hypothetical explanation.

Perhaps force-modifier theorists could deny that the target of explanation is contrasted with the prejacent in (74) and (77). In (78), for example—insofar as it is acceptable—it is plausible to interpret *perhaps* as making the hypothetical more remote, rather than as changing the target of the explanation from the prejacent to something else.

- (78) I seriously doubt that John is contagious. Still, if he even perhaps is, it's only because he works in that old hospital ward.

This is relevant because adverbs like *perhaps* are commonly viewed as force modifiers, even when it is contested whether auxiliaries like *might* are to be analyzed as such.<sup>10</sup> Yet this interpretation of *perhaps* in (78) does not seem to me to be any more available for *might* in (74) than it does for *chance* in (79).

- (79) I seriously doubt that John is contagious. Still, if there's even a chance he is, it's only because he works in that old hospital ward.

In both these cases, the modal seems to change the target of the explanation rather than (merely) reinforcing its hypothetical nature. Yet once we grant that the modal does this, we have to find some suitable target of explanation other than the prejacent. And here, as we have seen, the proponent of the force modifier view who denies a distinctive modal content to the conditional antecedents of (74) and (79) would seem to be in a bind.<sup>11</sup>

### 2.3.2 Revised force-modifier accounts

Apart from its inability to make sense of our data, I think there is a general problem for force-modifier accounts that should lead us to turn our attention elsewhere. This problem, alluded to in Chapter 1, is that force-modifiers view in their traditional form do not readily account for the effect of epistemic uses of modals in embedded contexts (cf. Papafragou 2006; MacFarlane 2011; Swanson 2011; Hacquard and Wellwood 2012; Anand and Hacquard 2013). One way to see this is to consider (80)-(83) (the examples and discussion are drawn from, yet also expand upon, Papafragou 2006).

---

<sup>10</sup> Bach (2011), for example, does not even consider a force-modifier analysis for *might* while taking it for granted that *perhaps* sometimes serves to modify force, not content (2011, 26, fn. 12).

<sup>11</sup> Related to the view of epistemic uses of modals as force-modifiers is the view of them as evidential markers. In the simplest form, modals as evidentials would mark the nature of the evidence for the speech act content without contributing to it (cf. Westmoreland 1998; Drubig 2001). But this view is subject to the same difficulties in any attempted explanation of our data that the force-modifier view is. In alternative forms, epistemics encode an evidential component of meaning over and above their modal one (cf. von Stechow and Gillies 2010; Matthewson 2015), or evidentials themselves have a modal semantics (cf. Matthewson et al. 2008). But these views are subject to the same difficulties that other truth-conditional views will be shown to be subject to in attempting to explain our data.

- (80) I'm going home since my son is coming to visit.  
 (81) I'm going home since my son might be coming to visit.  
 (82) I'm going home since I'm so excited that my son is coming to visit.  
 (83) I'm going home since, {yay! / yech!}, my son is coming to visit.

With (80)-(82) a speaker expresses different reasons for going home; with (83), however, she expresses not a different reason for going home but instead a different attitude *toward* her reason for going home. The latter is what we expect in these embedded contexts from terms that merely express the speaker's commitment or attitude toward the content of a speech act, without also contributing to that content itself. Yet this is not what we see from the embedded *might*. Instead, in (81), *might* contributes, somehow, to the content of the reason given for going home. On standard force-modifier views, it is not clear how to account for this contribution.

In recent years, force-modifier views have been revised to do just this. The results are compositional views that account for the effect of embedded-context occurrences of *might* while also maintaining the core commitment to matrix-context occurrences not contributing to the content of an assertoric speech act (Swanson 2006; 2011). Some of these accounts are even truth-conditional (Yalcin 2007; 2011). A number of authors have raised problems for these accounts, however, including their doing away with distinctive modal contents (cf. MacFarlane 2014, 277–79). I will not consider these problems here. Instead, I want to show that such revised accounts cannot make sense of our data, and that there is no good reason to adopt a force-modifier view anyway once we have had to make it compositional.

First, the potential application to our data. Yalcin (2007) has argued that a *might*-sentence could have a compositional effect as the antecedent of a conditional even without expressing a distinctive modal content. This holds out the hope of avoiding the problem for traditional force-modifier views from the previous section, for Yalcin will assign a different compositional semantics to (84) and (85).

- (84) If John is contagious, then there's been a general outbreak at the hospital.  
 (85) If John {even} might be contagious, then there's been a general outbreak at the hospital.

Given this difference, Yalcin could hope to explain why the target of explanation in (85) could be contrasted with the prejacent, without being either a modal content or the speaker's commitment to that prejacent.

The hope is short-lived. On Yalcin's view, conditionals are interpreted relative to information states, modeled as sets of worlds compatible with the content of that state (2007, 998–1000). A conditional is true (or to be accepted) at an information state if and only if the non-empty subset of worlds compatible with that state in which the antecedent is true is also a subset of the worlds in which the consequent is. Thus, (84) is true (or to be accepted) at an information state if and only if restricting that state to those worlds in which it is true that John is contagious leaves only worlds in which there has been a general outbreak at the hospital.

Now, on Yalcin's view, epistemic uses of modals serve as a *global* test on an information state, and the test that possibility modals perform is to ensure that the

prejacent is true in at least one of the worlds compatible with that state (2007, 1004). But if this condition is met, then the modal leaves that stage unchanged. Assume this condition *is* met: then (85) is true (or to be accepted) at an information state if and only if every world compatible with the content of that state is one in which there has been a general outbreak at the hospital. But surely this is too strong. I can accept (85) without thinking that there has been such an outbreak. I can do so, for example, if I think that the only way there could be a certain possibility (which I am still unsure about) is if some specific condition has obtained (which I am also still unsure about).

Abstracting away from conditionals momentarily, it is not clear how epistemic uses of modals could have the sort of informational effect they do on *any* force-modifier view, whether traditional or revised. The problem with conditionals just considered is just a specific instance of this more general problem. For example, you tell me that John might be contagious. I accept your claim and respond that there's been a general outbreak at the hospital then. But how, on a force-modifier view, could my accepting your claim justify this response? Surely your weakened commitment to the prejacent does not by *itself* allow me to infer that there has been a general outbreak. The same holds for the compatibility of your and my information state with John's being contagious. On learning that our information states *are* so compatible, how could I be justified in adding some proposition about the state of the world to those information states? Surely the process should be reversed, so that on adding some proposition to my information state I recognize that other propositions are now compatible with or entailed by that state.

Perhaps there are ways that revised force-modifier accounts could interpret conditionals to make sense of the informational effect of epistemic uses of modals, and to account for the contrast in (74) between the prejacent and the target of the hypothetical explanation. I will not attempt to consider them here. For over and above any problems with its application to our data, I do not see any reason to adopt a revised force-modifier view. That is, I do not see any reason why *might's* apparent contribution to the content of the reason given for going home in (81) could not also be its contribution to the content of an assertoric speech act. To hold that matrix-context occurrences of *might* are used assertorically is not to hold that they could not also be used to perform another speech act along expressivist lines (cf. Portner 2009, 172–77, and the references cited therein). Hence, if we have to develop a compositional account anyway, then we need some positive reason to think that matrix-context occurrences are never used assertorically. I do not know what this reason is.<sup>12</sup>

In fact, some proponents of force-modifier views even grant that what is part of modal force for one expression may be “objectified” and become, more or less, part of modal content for another (Halliday 1970, 336–37). This objectification process, for Halliday, would seem to be exemplified in (86)-(88).

---

<sup>12</sup> Lyons (1977), you may recall, distinguished between objective and subjective uses of *might*. In doing so, he also proposed that the former embed but the latter do not. Reasons to think that *might* did not embed would *ipso facto* be reasons to think that it was being used non-assertorically. However, the question here is whether we have reason to think that *might* is being used non-assertorically *despite* being embeddable.

- (86) John might be contagious.
- (87) It is possible that John is contagious.
- (88) There is a possibility that John is contagious.

Yet if (88) can be used to assert the existence of a possibility, then I am unsure why (86) should not also be able to be so used.<sup>13</sup>

Deep-seated philosophical discomfort with the existence of possibilities, which some authors seem to evince (cf. Yalcin 2011), provides no reason to deny that (86) can be used assertorically once we have granted that (88) can. And such discomfort provides no more reason to deny that (88) can be used assertorically than deep-seated discomfort about the existence of pain would provide reason to doubt that (89) can.

- (89) There is a pain in my left foot.

There *is* a clear assertoric/non-assertoric distinction to be drawn between (89) and (90), which even skeptics about the existence of pain have to acknowledge.

- (90) Ouch!

It is not clear that there is any such distinction to be drawn between (86) and (88). (86), unlike (90), has all the hallmarks, including embeddability, that suggest association with a distinctive content that could be the object of an assertion. Thus, even apart from its inability to explain our data, I do not think we should adopt the force-modifier view. We *should* take apparently epistemic uses of modals like *might* to contribute to the truth-conditional content of a speech act. We should just not take this contribution to be an information state, as I will now argue.

## 2.4 Potential epistemic truth conditional solutions

For proponents of epistemic truth-conditional views, there are three different ways in which the content of a speech act could depend truth-conditionally on an information state.<sup>14</sup> First, that state could be part of the proposition expressed in the speech act (cf. DeRose 1991; von Fintel and Gillies 2008; 2011; Hacquard 2010; Bach 2011; Dowell 2011; von Fintel and Heim 2011; Braun 2012; Yanovich 2014; a.o.). Second, that state could be used to help *determine* the proposition expressed, without itself being part of it (cf. Silk 2017). Third, that state could be part of an expanded evaluation matrix against which the proposition expressed was evaluated for truth or falsity (cf. Egan et al. 2005; Egan 2007; Stephenson 2007; MacFarlane 2011; 2014; Beddor and Egan 2018). The first view is by far the most commonly held, though the third has gained in popularity in recent years. Yet neither can make sense of why claims made with apparently epistemic uses of *might* do not support epistemic explanations. And while the second view can, it does so only by incorrectly predicting that *might*-claims

<sup>13</sup> Again, this is *not* to say that (86) could not also be used performatively in a way that (88) is not.

<sup>14</sup> In talking of the truth-conditional content of a *speech act*, I intend my claims to be compatible with views that maintain that (either some or all) sentences-in-contexts lack truth conditions, while still maintaining that sentences can be used in contexts to make *claims* that have truth conditions (cf., respectively, Bach 1994a; 2005; Pietroski 2005; 2018).

will not support any non-trivial explanations at all, epistemic *or* non-epistemic. I elaborate on these points below, starting with the view that takes information states to be part of content.

## 2.4.1 Information states as an implicit part of content

### 2.4.1.1 A pragmatic effect of implicitness?

If an information state is part of the content of the claim made with (91), it would have to be an *implicit* part of content, something that is not explicitly linguistically referenced but instead supplied by the extra-linguistic context (Kratzer 1977; 1981; 1991).

(91) John might be contagious.

On one important view in the literature, that implicit part of content can be made explicit through the use of restrictor phrases specifying whose information state is at stake (cf. Kratzer 1977; 1981; 1991; cf. also Bach 2011; von Stechow and Gillies 2011; Schaffer 2011). Some candidate phrases include those listed in (92) and (93) (cf. also Braun 2013).

(92) {Given/According to/In view of} what *x* knows

(93) {As far as/For all that} *x* knows

And, as it turns out, allegedly explicit counterparts of the claims that could be made with (91) do in fact support epistemic explanations, as (94) shows.

(94) a. Given what we know, John might be contagious.

b. That's only because we still don't know about his test results.

It might thus be thought that leaving an information state implicit serves some sort of pragmatic function that restricts the explanations that a claim supports.

Alleged implicitness, it should be noted from the start, does not in general have this effect, neither for modal claims nor for non-modal ones. Without their restrictor phrases, for example, the claims made with ((95)a) and ((96)a) could be either deontic or ability ones. And, without those restrictors, the deontic and ability claims support exactly the same sort of explanations that they do with them.

(95) a. {Given their laws,} Wisconsinites can shoot pigeons in hunting season.

b. That's only because their legislature is still controlled by the gun lobby.

(96) a. {Given their training,} Wisconsinites can shoot pigeons in hunting season.

b. That's only because they start learning marksmanship in grade school.

When it comes to non-modal claims, something similar holds. They also support explanations that target part of the content of the claim that has not been made

explicit. (97) and (98) show as much with examples of non-explicit meaning that Bach (1994a) calls *implicitures*.<sup>15</sup>

- (97) a. John started his own business and put himself through school.  
 b. That's only because his parents refused to pay his tuition and there was no other way he could afford it.
- (98) a. Ron won.  
 b. That's only because the Yahtzee board was rigged.

In (97), the explanation only makes sense if there is taken to be a temporal component to the speaker's claim, indicating the sequence in which the conjuncts became true. Yet this temporal component is not made explicit, nor does it need to be for the explanation to be acceptable. And in (98), the explanation only makes sense if the speaker's claim is taken to be about a specific contest. Yet again, this component does not need to be made explicit for the explanation to be acceptable. Hence, being explicit does not generally seem to matter for the explanations that a speaker's claim supports.

In fact, in some cases leaving a part of content implicit seems not to restrict but instead to expand the range of explanations than a claim supports. Take (99), which can be used to make a claim about a particular location. (Imagine a worried mother is speaking on the phone to her vacationing son.)

- (99) Hey, be careful when you go out. It's raining acid.

The location the mother's claim is about could have been specified either indexically, as in (100), or descriptively, as in (101).

- (100) Hey, be careful when you go out. It's raining acid there.  
 (101) Hey, be careful when you go out. It's raining acid where you're at.

With the location left implicit, the mother's claim can more readily support not only explanations that trade on an indexical specification of that location, as in (102), but also explanations that trade on a descriptive one, as in (103).

---

<sup>15</sup> These implicitures are supposed to be generated, respectively, by what Bach calls the non-literal use of a sentence (as opposed to any of its constituents) and by what he calls the semantic incompleteness of a sentence (or its failure to determine a complete proposition). Especially in the case of semantic incompleteness, there is little debate that the meanings that Bach calls implicitures can be the content of a speaker's claim—that is, of her direct illocutionary speech act (cf. Fillmore 1986; Partee 1989; Condoravdi and Gawron 1996; Williams 2012). There is debate, however, about whether (and, if not, why not) these meanings should also be taken to be the content of the sentence-in-context that the speaker has uttered (Sperber and Wilson 1986; Bach 1994b; 2001; 2005; Carston 1988; 2002; 2004; Stanley 2000; Recanati 2001; 2002; Cappelen and Lepore 2002; 2005; 2007; Borg 2007). Since my concern in this dissertation is primarily with speech act content, and since theorists agree that the content of a sentence-in-context constrains, even if it does not determine, this content, I will not engage with this debate here.

- (102) Ugh! That’s because the Czechs have never made the environment a priority.
- (103) Ugh! That’s because we’re vacationing in the Czech Republic and not in Spain.

Even if (103) sounds somewhat worse than (102) as a response to (99), nonetheless it seems to sound better as a response to (99) than as a response to (100). If this is correct, then not making a part of content explicit can in fact have the opposite effect that it is alleged to with apparently epistemic uses of modals. Hence, if implicitness is to restrict the sort of explanations that epistemic claims support, it must presumably be due to some pragmatic function that is served by leaving an *information state* implicit.

Several authors have proposed that leaving an information state implicit does, in fact, serve such a function. Bach (2011), for example, has claimed that when you mention a particular information state, or in his terms a *perspective*, “you make it the focus.” “Not mentioning a perspective,” he continues,

is a way of keeping the question of the possibility in focus and, moreover, of hedging the question as to whose/which perspective is at issue (2011, 57–58).

The sort of proposal that Bach makes is the centerpiece of von Fintel and Gillies (2011)’s *cloudy contextualism*. These authors propose that the discourse context, or the speech situation, often does not determine a unique context of evaluation. In the relevant case, for example, it does not determine a unique assignment of values to (covert) variables ranging over information states. This point about discourse contexts not metaphysically determining evaluation contexts has been made before, usually in conjunction with the point that a speaker’s (reasonable) communicative intentions help to *epistemically* determine the intended evaluation context (cf. Bach 2005). In this case, however, von Fintel and Gillies suggest that speakers may themselves not have any intended evaluation context. Instead, they exploit the indeterminacy of the discourse context, and the “cloud of admissible [evaluation] contexts” it leaves open, to “put into play” a set of propositions individuated by information states (2011, 118–19).

Putting a set of epistemic propositions into play is supposed to change conversational dynamics. While typically a speaker is in a unique position to assert one of these propositions (the one involving her own information state), her interlocutors are in a unique position to take up and respond to another (the one involving their own state).<sup>16</sup> And a cooperative interlocutor in fact *should* respond to a speaker’s utterance on the basis of the epistemic proposition that she herself is in a unique position to assert or deny. Accepting or rejecting the utterance on the basis of this proposition is more appropriate because more informative: it lets a speaker know whether the prejacent is also compatible with that interlocutor’s information (2011, 121–22). And putting a set of epistemic propositions into play, as opposed to asserting any single one, is just supposed to be a way of probing for such information and attempting to redress a relevant information asymmetry (2011, 123).

---

<sup>16</sup> For non-epistemic implicit content, this sort of mismatch could not generally be expected.

On these views in which the implicitness of an information state serves an important pragmatic function, the problem with an example like (104), and the reason for its contrast with an example like (105), would be that the attempted explanation engages the wrong epistemic proposition.

- (104) a. John might be contagious.  
 b. #That's only because you haven't asked him his diagnosis.
- (105) a. Given what I know, John might be contagious.  
 b. That's only because you haven't asked him his diagnosis.

The response purports to explain the truth of a proposition that the speaker chose not to assert, meanwhile maintaining silence on the proposition the speaker would really like to have confirmed or denied.

There are two general problems for these and other potential pragmatic views for why implicitness should constrain the explanations that a *might*-claim supports. One problem is that they do not account for contexts in which the implicitness of an information state could serve no pragmatic function, including in embedded contexts and also in discourse contexts in which there is no relevant information asymmetry between a speaker and her interlocutors.

Start with the second sort of context. With the agreement in (106), the relevant information asymmetry has just been ruled out.

- (106) a. John might be contagious.  
 b. You're right. #But {perhaps} that's because we haven't asked him his diagnosis.

And with the monologue in (107), an information asymmetry is impossible.

- (107) John might be contagious. #But {perhaps} that's because I haven't asked him his diagnosis.

In both these cases, an epistemic explanation is still bad. Yet in the latter case, failure to reference an information state could not function as a way to probe for information. And in the former case, the epistemic proposition being explained—the one involving the interlocutors' joint information state—would be one that had been added to the common ground as a result of a successful probe (cf. von Stechow and Gillies 2011, 122–23). In neither case, then, would the epistemic explanation be engaging the “wrong” proposition. Yet still these explanations are bad.

Turn now from discourse contexts to embedded ones, as with the final *might*-sentence in ((108)b).

- (108) a. Given what we know, John might be contagious.  
 b. You're right. Given what we know, John might be contagious. #But it's only because we haven't asked him his diagnosis that John might be contagious.

In ((108)b), the embedded *might*-sentence could not be used to assert any proposition, and so presumably could not be used to float any set of propositions either. Further, given the resolution of the contextually relevant information state up to that point, it should be clear that the particular epistemic proposition it is used to express is the group proposition asserted with ((108)a) and agreed to with ((108)b). And yet, despite all this, the explanation in ((108)b) is bad even though the explanation in (109) is perfectly fine.

- (109) But it's only because we haven't asked him his diagnosis that John might be contagious, given what we know.

This contrast suggests that restrictor phrases like *given what we know* do not serve to make explicit an otherwise implicit part of the content of a *bare might-sentence*, i.e., a sentence unrestricted by such phrases.

The second general problem for pragmatic views is that they cannot explain why the implicitness of an information state should not matter in the same way for other epistemic claims as it is alleged to for *might*-claims. No particular information state is referenced in ((110)a), for example, any more than in ((104)a).

- (110) a. It's plausible that John is contagious.  
b. That's only because you haven't asked him his diagnosis.

Yet what is plausible varies with the information state just as what is epistemically possible does. ((104)a) and ((110)a) should thus be equally able to be used to probe for information via putting a set of epistemic propositions into play. That they do not license the same responses suggests that it is thus not any probing function that an information state's implicitness might serve that matters for an epistemic explanation's acceptability. Acceptability would instead have to be correlated with some arbitrary lexical property of modals.

One property that might seem relevant is *flexibility*, or the ability of a modal to be used to express multiple flavors of modality (both epistemic and non-epistemic, for example). *Might*, in contrast to *plausible*, is often assumed to be flexible: to have not only epistemic readings, as allegedly in (111), but also non-epistemic ones, as in (112), which seems to express a metaphysical possibility about the open future (cf. Klecha 2016; Condoravdi 2002).

- (111) John might now be contagious.  
(112) John might be contagious in two months.

That said, the temporal modifier in (111) is alleged to force an epistemic reading of *might*, eliminating modal flexibility (Condoravdi 2002). Yet epistemic explanations are no more acceptable in response to (111) than in response to any of the previous *might* sentences. Hence, flexibility does not seem to matter in a way that could make sense of our data either.

To summarize: there does not seem to be any pragmatic function that leaving an information state implicit could serve that could make sense of the fact that bare *might*-sentences (or BMSs) do not license epistemic explanations, not even when embedded. This suggests that if such states *are* an implicit part of content, then they are not a part of content that can be made explicit by restrictor phrases like *given what we know*. For the difference in the explanations that are supported with and without these restrictors cannot be attributed to any pragmatic effect of leaving an information state implicit, and so must presumably instead be attributed to a semantic difference in the content of the claims made with and without such phrases.

#### 2.4.1.2 A semantic effect of implicitness?

I am not the first to suggest that restrictor phrases like *given what we know* do not function in the way that they are standardly assumed to. Braun (2013), for example, notes that such modifiers can also restrict non-modal sentences, as in (113), and avers that “it would be hard to justify the claim that these [phrases] perform different semantic functions [in the two cases]” (2013, 503).

(113) Given what we know, John speaks French.

On Braun’s view, such phrases do not supply a missing part of an incomplete proposition but instead uniformly attribute a property to the complete proposition expressed by the sentence they restrict (2013, 502). This is the property, it can easily be inferred, of that proposition’s being compatible with the relevant information state.

Braun’s position makes ready sense of why sentences like (113) license epistemic explanations like (115) even when their non-restricted counterparts like (114) do not.

(114) John speaks French.

(115) That’s only because we know about his semester in Paris.

It could readily make sense of the same contrast between restricted *might*-sentences and BMSs too. However, Braun still takes sentence like (116) to be able to be used to make epistemic claims, even though he obviously does not take sentences like (114) to be so used. Hence, something needs to be said about these particular epistemic claims that could make sense of their not supporting epistemic explanations.

(116) John might speak French.

Braun’s challenges are instructive in this regard. He proposes that *might* semantically expresses a property of propositions that he calls *weak possibility* (2012, 469–72), a generic property entailed by more specific properties such as epistemic, metaphysical, or nomological possibility (2012, 470). On Braun’s account, a speaker who utters a BMS thus says something that is almost certainly (and trivially) true, namely, that the prejacent is a weak possibility.<sup>17</sup> However, what she says is likely not what she means. Instead, the contents of her *locutionary* and *illocutionary* speech acts

<sup>17</sup> “Obviously, many propositions possess ... weak possibility. Perhaps all do” (Braun 2012, 471).

come apart, the speaker likely asserting that the prejacent is an epistemic possibility for herself (2012, 470; cf. Austin 1962). But this, it turns out, is similar to, or perhaps even the same as, what a speaker asserts in uttering a sentence like (117).

(117) Given what I know, John might speak French.

For the apparent semantic content of (117), for Braun, is that the prejacent's being weakly possible is an epistemic possibility for the speaker (cf. Braun 2012, 475; 2013, 502). This is not what a speaker is likely to mean in uttering (117), however.<sup>18</sup> Instead, a speaker is likely to have meant the far simpler proposition that the prejacent *itself* is an epistemic possibility for her. In this case, in uttering (116) and (117) she will have said different things but meant the same.

Once Braun grants that both bare *might*-sentences and restricted ones can alike be used to assert epistemic possibility propositions, how does he account for the fact that only the latter license epistemic explanations? Presumably he has to say that it is because only the latter are also used to say (or, in his terms, to *locute*) epistemic propositions. But this response carries no conviction. For the non-epistemic explanations that BMSs license must clearly be explanations of asserted content, not locuted content. (Presumably, there can be no interesting explanation of what makes a weak possibility proposition true.) So if BMSs can be used to assert epistemic contents, and if the explanations that they license seem to target asserted contents too, then it is unclear why some of the contents they can be used to assert—namely, epistemic ones—cannot be explained in terms appropriate to them.

The only potential solution to this problem that I can see is to define the information state that is supposed to be an implicit part of content expansively, as do several authors in the literature. Yanovich (2014), for example, has proposed that a speaker's claim in uttering a BMS is that the prejacent is compatible with the relevant and readily available information, where this may be more expansive than the information of any individual or group within the conversation (2014, 76–78). DeRose (1991) makes a similar proposal: the speaker's claim is that the prejacent is compatible with what is known by the relevant community, or could come to be known by them in the relevant ways, relevance in both cases depending on the context (1991, 594).

Assume that the information that our epistemic explanations have appealed to so far—John's test results, his diagnosis, etc.—is in all cases readily available. The problem with these explanations, then, would be their appeal to the status of that information as not yet *known* as being what grounds the truth of the speaker's claim. For whether the prejacent is compatible with the readily available information is clearly independent of whether that information is known. As a result, for Yanovich and DeRose, the attempted explanations of the speakers' claims we have considered so far would all be misplaced.

The problem for this suggestion is that epistemic explanations in terms of information not being available, as in (118), are just as bad as any of the explanations considered so far in terms of information not being known.

---

<sup>18</sup> And if it is, then she will have still asserted an epistemic possibility proposition all the same.

- (118) John might be contagious. #But that's only because *you* didn't ask his doctors his diagnosis the one chance you had.

It is not clear how Yanovich and DeRose could account for this fact. For with (118) a perfectly fine explanation is provided of why the prejacent might be compatible with the readily available information, as (119) shows.

- (119) We don't know whether John is contagious, and we can't know for certain anytime soon. But that's only because *you* didn't ask his doctors his diagnosis the one chance you had.

Additionally, if we appealed to expansive information states for BMSs, we should also do so when no particular information state has been explicitly referenced for terms like *epistemically possible* and *plausible* either. For what motivates DeRose's and Yanovich's accounts of *might*—truth value judgments about the claims made with BMSs—would motivate similar accounts for these terms. Consider a speaker's retrospective truth value judgment, as in (120), for example.

- (120) a. [S1] Have you seen John's rash? It's quite plausible he's contagious.  
 b. [S2] You can't have known this but, as John's physician, I can tell you categorically: there's no way his rash is symptomatic of a contagious disease.  
 c. [S1] Oh. I guess I was wrong then.

Accounting for the reasonableness of such retrospective judgments is one phenomenon motivating the appeal to expansive information states (cf. Yanovich 2014, 69–94). For such retrospective judgments often do not seem reasonable when a particular information state has been explicitly referenced: ((120)c), for example, would *not* seem reasonable if S1's first utterance had been (121) instead of ((120)a).

- (121) Have you seen John's rash? It's quite plausible to me that he's contagious.

That said, if we appealed to expansive information states for *epistemically possible* and *plausible*, then they should no more license epistemic explanations in terms of information not being known than do BMSs. Yet with such terms appeal to the status of information as not being known is just as fine as is the appeal in (122) to the status of information as not being available.

- (122) The only reason it's {epistemically possible/plausible} that John is contagious is that *you* didn't ask his doctors his diagnosis the one chance you had.

The data cross-cut predictions then. When information *can* be appealed to to explain the truth of a claim, the status of that information as merely not yet possessed, or instead as also not being readily available, does not matter. The same holds for when information *cannot* be appealed to. The availability of information does not

explain the acceptability of an epistemic explanation. Nor do I see any better way to make sense of *might*-claims not supporting epistemic explanations if we take information states to be an implicit part of the content of those claims. If we are to maintain a truth-conditional dependency on information states, while still explaining our data, it seems that information states must be removed from content.

## 2.4.2 Removing information states from content

### 2.4.2.1 From content to character

Some theorists have proposed that the truth-conditional role of an information state is merely to determine, without itself being part of, the content of the claim made with a BMS (Silk 2017). Reference to information states plays the role of a Kaplanian *character* on this account: that is, it provides a function from speech contexts to speech contents (Kaplan 1989). For modals like *might*, the content determined by their character is the set of propositions that some particular, contextually relevant information state comprises. And the content of the claim made with a BMS is that the prejacent is compatible with this set (Silk 2017, 1783).

The promise of this response is that the content-character distinction is itself a way of separating out the contingent determinants of truth conditions that can be appealed to in explanations of the truth of a claim (i.e., contingencies of content) from those that cannot (i.e., contingencies of character) (Kaplan 1989). The relevant contingencies are not those determining *what* content is expressed in a given context, that is, but instead those determining whether the content that *is* expressed is true. And if the contingent nature of an information state only determined what content was expressed in a given context, then we would not expect such contingencies to be able to be appealed to in an explanation of the truth of that claim.

The problem with this response is that the claims made with BMSs seem to have contingent contents, as we have seen from Section 2.1, while logical relations between sets of propositions are non-contingent (cf. Kratzer 1991). (123), for example, is a perfectly coherent and reasonable thing to say.

- (123) The only reason that John might be the killer is that he never learned to control his temper. If he had, there wouldn't even be this possibility.

Yet it would not be reasonable if the possibility the speaker claims to be contingent on John's temper were a non-contingent relation between sets of propositions. To secure the contingency of the contents of modal claims, we need the descriptions that contingently characterize sets of propositions—descriptions like *our information*, for example—to *not* play the role of a Kaplanian character. And yet if we secure contingency by taking descriptions like *our information* to be part of content, then it is not clear why we cannot explain what makes a claim true in terms of the contingencies of that information.

### 2.4.2.2 From content to the evaluation matrix

Perhaps we should remove information states, as well as the sets of propositions they comprise, from content altogether. *Relativists* hold that information states matter truth-conditionally, for example, but they regard them as contextually given aspects,

not of the proposition expressed, but instead of the evaluation matrix relative to which it is true or false (Egan et al. 2005; Egan 2007; Stephenson 2007; MacFarlane 2011; 2014). The same proposition may thus be true relative to one information state but false relative to another. The truth of a speaker's assertion of that proposition is likewise supposed to be relative, varying with the information states determined by the contexts in which that assertion is assessed.<sup>19</sup>

On this view, it is plausible that explaining what makes an assertion true would require explaining what makes it true relative to the *explainer's* context. This, after all, is the context that grounds the judgment of truth in the first place. We could thus make sense on this view of the badness of any examples in which we could plausibly assume that a speaker was explaining why the prejacent was an epistemic possibility only for an interlocutor. Problematically, however, we could not make sense of the examples in which a speaker was explaining why it was an epistemic possibility for herself (potentially as well as for others).

Perhaps the relativist should hold, then, that even the information state determined by the assessment context cannot always be appealed to. In many cases, an assessor should recognize, after all, that she could later occupy an assessment context at which an assertion true relative to her current context would be false (cf. MacFarlane 2014, 255, fn. 12). And perhaps you should not attempt to explain the truth of a claim whose truth you recognize is potentially variable in this way. With this stricture in place, the relativist could thus make sense of the badness of examples in which a speaker recognizes that she could gain information which could make the prejacent cease to be an epistemic possibility for her. However, it is not clear how she would make sense of the badness of an example like (124) in which two specialists have determined that the prejacent's truth value is unknowable.

- (124) a. John might be contagious.  
 b. #But {that's} only because we've just determined that there are no tests that anyone could run to establish whether he is.

Perhaps someone might respond that the prejacent's truth value is not strictly speaking unknowable, given the possibility of an omniscient being. But this possibility is irrelevant. It would make explanations of the truth of epistemic possibility claims viable only when (and in terms of the fact that) the prejacent itself was true. Yet even this absurd result is not consistent with the data. (125) is an admittedly odd sort of dialogue for anyone to have, and yet, even still, it is coherent in a way that (126) is not.

- (125) a. It's plausible that John is contagious.  
 b. Only because we know that he is.

---

<sup>19</sup> Often these will be the information states of single individuals, but relativists could also hold that they can sometimes be joint information states (cf. Egan et al. 2005, 153). It should also be noted that Stephenson departs from the other theorists listed in only taking the content of the assertion, and not the assertion itself, to be relative.

- (126) a. John might be contagious.  
b. #Only because we know that he is.

MacFarlane himself also points out the absurdity of appealing to the information states of omniscient beings when he writes, in a different context, that:

[I]t is not incoherent for a relativist to aspire to have more information, while recognizing her current information state as the information state that determines whether she can correctly say that something is “possible”.... [N]obody would say that, because information states can be better or worse, all epistemic modals should be interpreted relative to the best possible information state—full omniscience (2014, 148, fn. 6).

Just as a possibility claim should be evaluated relative to the current information state for a relativist, so, we might think, its truth should be explained in terms of that state too, and not in terms of the information state of an omniscient being.

Anyhow, even if relativists could somehow make sense of (124), they would still leave unaccounted for *might*'s contrast with *epistemically possible* and *plausible*, both of which behave on important diagnostics like presumed relativist terms, including being acceptable under subjective attitude verbs like *find*. Both also display the same perspectival restriction under *find* that presumed relativist terms like *tasty* do: namely, that to find something to be a certain way is to find it to be that way for *oneself*. These facts are shown by (127) and (128) (Alexander Williams, p.c.; cf. Lasersohn 2017; Kennedy and Willer 2016).

- (127) I find Joe's flan tasty {#for Sue}.  
(128) I find Joe's moon-landing theory {epistemically possible/plausible} {#for Sue}.

So the relativist no more than any other theorist is in a position to explain our initial contrast between *plausible* and *might* in their licensing of epistemic explanations.<sup>20</sup>

## 2.5 Toward a non-epistemic solution

Apparently epistemic uses of modals like *might* license explanations that seem to target a content other than the prejacent. This suggests, contra force-modifiers views, that such uses contribute to the truth-conditional content of a speaker's claim. However, such uses do not license *epistemic* explanations, something that truth-conditional views struggle to explain, even if they remove the relevant information state from content to character or even to the evaluation matrix.

---

<sup>20</sup> Stephenson (2007) does argue for a difference between predicates of personal taste (PPTs) and epistemic modals based on their behavior under doxastic attitude verbs like *think*. What Stephenson proposes is that PPTs can be used non-relativistically in some cases. As for *plausible* and *epistemically possible*, their behavior under *think* would group them with Stephenson's PPTs. This could seem to—but in fact does not—provide a basis for differentiating *might* from *plausible* and *epistemically possible*. For none of the cases considered so far would be non-relativistic by Stephenson's lights. Indeed, the fact that any epistemic modals pattern with PPTs complicates Stephenson's proposed distinction and suggests that it might rest on a misanalysis of *might* as an epistemic modal.

Proponents of epistemic truth-conditional views might have hoped to find some support for their position in the fact that (129), an allegedly explicit counterpart of (130), licensed epistemic explanations even though the latter did not.

- (129) Given what I know, John might be contagious.
- (130) John might be contagious.

Perhaps however, as suggested by our discussion of Braun, this fact is best explained in the same way in which we explain the fact that ((131)a), but not ((132)a), licenses an epistemic explanation.

- (131) a. Given what I know, Wisconsinites can shoot pigeons during hunting season.
- b. That's only because you haven't read their state laws.
- (132) a. Wisconsinites can shoot pigeons during hunting season.
- b. #That's only because you haven't read their state laws.

The contrast between (131) and (132) is explained by the epistemic modifier making explicit the *grounds* the speaker would have for making the unmodified, non-epistemic claim. In doing so, the modifier restricts the deontic quantification to worlds in which the known laws obtain. The modified claim thus has an epistemic component to its truth conditions, the unmodified one none. Hence, only the modified claim supports epistemic explanations. There is no epistemic part of truth conditions left implicit in ((132)a), and so no epistemic part of truth conditions to be made explicit by the modifier in ((131)a) either, and so ultimately also no puzzle about the contrast in the explanations supported either. The exact same could hold for *might*. The problem may not be the alleged implicitness of an epistemic contributor to truth conditions but instead the assumption of epistemic truth conditions itself.

## Chapter Three: The Development of a Non-Epistemic Analysis

The philosophically popular name of ‘epistemic possibility’ is a misnomer and the ideas based on it are mistaken.

~ White (1975, 86)

### 3.1 Introduction

In this chapter, I argue that we can make sense of our data about the (in)admissible explanations for the truth of modal claims if we assume, as White suggests, that not all modal claims that appear to be epistemic are in fact so. Instead, with some apparently epistemic uses of modals we make modal claims that are, in fact, non-epistemic. Take the claim made with (133).

(133) John might be contagious.

With (133), I propose, we make a claim about how John’s being contagious stands with regards to the circumstances. Specifically, we are claiming that circumstances leave open the possibility that he is contagious: that John’s being contagious is a *circumstantial* possibility.

In Section 3.2, I show that we can give a simple and straightforward account of the data from the previous chapter if I am correct about the content of our *might*-claims. For whether something is a circumstantial possibility depends on contingencies of the circumstances, not information states. The truth of claims about such possibilities should thus be explicable in these terms too. In Section 3.3, I develop the framework of a semantics and a pragmatics for circumstantial possibility claims. This framework is independent of any particular analysis of the circumstantial possibility relation, and instead focuses on the non-epistemic nature of that relation. In Section 3.4, I respond to objections intended to show that we cannot give truth conditions for the claim made with (133) in terms of any non-epistemic possibility relation, but instead must appeal to an epistemic relation overlying it. The question here is not whether there *is* a relevant non-epistemic relation—even proponents of epistemic truth conditions will need to grant as much.<sup>21</sup> Instead, the question is whether we need the *epistemic overlay* of this relation for analysis of the claim made with (133). I argue that we do not. I conclude in Section 3.5 by pointing forward to the following chapters in which the framework I develop is deployed to make sense of phenomena about modal discourse and beliefs.

Before anything else, however, an important terminological note. In the literature, the sort of modality expressed in a non-deontic use of (134) is often referred to as a circumstantial possibility (Kratzer 1981; 1991).

(134) John can win tomorrow’s race.

---

<sup>21</sup> If the content of an information state ever leaves open the possibility that John is contagious, this can only be because *what is represented* by that content does so.

Obviously, I do not think this sort of modality is the same sort expressed by *might*. However, the standard terminology to some extent prejudices the difference between the two sorts of modalities. To avoid prejudging the issue, I will thus reserve the term *possibility* for the sort of modality expressed by *might* and use the term *potential* to refer to the sort of modality expressed by non-deontic uses of *can*. This leaves open the analytical possibility that both *might* and *can* could express modal relations to (or involving) sets of circumstances, albeit different ones. Hence, it also leaves it as an empirical question whether *might* is better analyzed as expressing circumstantial or epistemic possibility. If the latter, as I argue, a question does arise as to the difference between the circumstantial modalities expressed by *might* and *can*. But this question does not require an immediate answer, and so I defer my response till Chapter 6.

### 3.2 Contingent claims on a non-epistemic analysis

On a non-epistemic interpretation of (133), the explanation of the data from the previous chapter is simple and straightforward. The contents of the claims made with *might* have non-epistemic truth conditions, but the explanations have assumed epistemic ones. The badness of these explanations is the result of a category error. And the same holds for epistemic explanations of the truth of the claims made with *must*.

The only role for information on a non-epistemic interpretation of *might* and *must* is as the grounds for making or accepting a modal claim. But information plays this same role for making or accepting *any* claim, modal or non-modal. To assume that the information state that is implied in making a modal assertion has to be part of the content of that assertion is to make exactly the sort of mistake that Frege and Moore have taught us not to make in the case of non-modal assertions (cf. Moore 1962).

Because information's only role for the claims made with *might* and *must* is as the grounds for those claims, and not as a truth-maker for them, it is only in the former role that contingencies of information can be appealed to in response to these claims. And this is precisely what we find. In contrast to the previous chapter's appeals to the contingent nature of an information state as making a *might*- or *must*-claim true, appeals to the contingent nature of an information state as the grounds for making or accepting such claims are perfectly fine, as in (135) and (136).

- (135) The only reason you say that John {*might/must*} be contagious is that he still hasn't told you about his test results. You wouldn't say that if he had.
- (136) The only reason I agree that John {*might/must*} be contagious is that I still haven't heard from him about his test results. I wouldn't agree if I had.

When we do want to explain what makes the content of a *might*- or *must*-claim true—and not just why someone would make or accept it—it is to non-epistemic facts that we must appeal. Such appeals are made in (137)-(140), and are perfectly fine, just as a non-epistemic analysis of *might* and *must* would predict.

- (137) The reason John *might* be contagious is that he works in a hospital ward.
- (138) The reason John *must* be contagious is that he works in a daycare center.

(139) The reason the coin *might* be heads is that the toss was fair.

(140) The reason the coin *must* be tails is that the coin was weighted.

With (137) and (138), for example, we suggest that certain circumstances make John's being contagious either a possibility or a necessity. Insofar as he works in a hospital ward, we suggest that it is possible that John has contracted a contagious disease. In an alternative scenario in which John did not work in a hospital ward, his interactions with others might not have left open this possibility. And insofar as he works in a daycare center, we suggest that it is impossible that John has *not* contracted a contagious disease. In any scenario in which John works in a daycare center, he will have contracted one.<sup>22</sup>

These non-epistemic explanations of why John might or must be contagious, even if in practice they are rejected, are in principle the right sort of explanation to be giving. There is nothing anomalous about suggesting that the reason that John might be contagious is that he works in the worst hospital ward, even if we ultimately conclude otherwise. Yet there is something anomalous about suggesting the reason he might be so is that our information is incomplete. The source of this anomaly, if I am correct, is that *might* and *must* function to make the content of a speech act truth-conditionally dependent on the prejacent's relation to the circumstances, rather than on its relation to the content of some information state. Epistemic explanations of what makes *might*- or *must*-claims true are thus always misguided. They always commit a sort of category error that non-epistemic explanations do not.<sup>23</sup>

Let me emphasize that there is nothing untoward in the idea that *might* and *must* could make a non-epistemic contribution to meaning even though certain epistemic states will typically be implied when they are used. For with many terms that make non-attitudinal contributions to meaning, we can nonetheless imply specific attitudes when we use them. All else being equal, this is probably what we should assume happens *whenever* the truth of a speaker's claim cannot be explained with regard to the contingent etiology of whatever attitude(s) she may have implied in making it. Hence, this is probably what we should assume happens in the case of *might* and *must* too.

In fact, in the default case of uttering a declarative sentence, a speaker will imply an attitude—namely, belief—that is not itself part of the content of that utterance (whence Moore's Paradox). This implication without assertion of an attitude explains why ((141)b') is fine as a response to ((141)a) even though ((141)b)—which purports to explain what makes the asserted content of that utterance true—is not.

(141) a. John is contagious.

b. #Only because you haven't asked him his diagnosis.

b'. You only believe that because you haven't asked him his diagnosis.

But if it is standard for the attitude of belief to be implied without being asserted, it should not surprise us if specific lexical items enabled other attitudes to be implied

---

<sup>22</sup> And something similar holds for (139) and (140), where we suggest that certain circumstances of the coin toss make it either possible or necessary that the coin has landed in a certain way.

<sup>23</sup> Again, this is not to say that non-epistemic explanations may never be rejected for other reasons.

without being asserted either. This especially should not surprise us in the case of modals. For it is natural to think of sentences like *p*, *might-p*, and *must-p* as part of a single system with which we assert, respectively, (142)-(144).

- (142) *p* is actually true.
- (143) *p* is possibly true.
- (144) *p* is necessarily true.

We have learned to separate the attitudinal state one must be in to make the first sort of judgment from the content of that judgment, and the truth conditions for that content (cf. Frege 1956). We can, and should, do so for the second and third sorts of judgments about the truth value status of a proposition as well.

### 3.3 The framework for a non-epistemic analysis

In this section, I am going to show how we can separate the state that one is generally in when one makes a possibility judgment with a sentence like *might-p* from both the content of that judgment and the truth conditions for that content. To effect the separation, I propose we do four things. The first is to recognize that circumstantial possibility is relational. The second is to remove one of the *relata* of this possibility relation—namely, sets of circumstances—from content to the evaluation matrix for content. The third is to allow that discourse contexts may not always antecedently determine unique values for all the parameters of the evaluation matrix. And the fourth is to allow that speaker intentions may not always do so either.

These four points are not novel. The second is analogous to one of the two key components of epistemic relativist proposals (cf. Egan et al. 2005; MacFarlane 2011; 2014). And the other points are more or less standard fare in the Kratzerian analysis of circumstantial modalities, or are adaptations of that standard fare to the removal of the modal base from content to the evaluation matrix (cf. Kratzer 1977, especially 342-343; 1981, 294–95). What *is* novel is their combination and their application to modalities that have been analyzed as epistemic.<sup>24</sup> Combined, the fundamental effect of these four points is to allow the contents of possibility judgments to be independent of *specific* sets of circumstances. And this opens up the analytical possibility that speakers could typically make such judgments only when in a certain state—for example, the state of not being aware of circumstances that close off the possibility of the prejacent's being true—without the content of those judgments being about or truth-conditionally dependent on that state.

I expand upon these four points below, and then develop them within an adapted version of the standard Kratzerian framework for modal semantics (cf. Kratzer 1977; 1981; 1991).

---

<sup>24</sup> I will also argue, here and in the following chapter, that the four points should be combined for *all* circumstantial modalities, not just those currently analyzed as epistemic. That is, the modal base should be removed from content to the evaluation matrix for circumstantial modals in general, and the Kratzerian points about the relation of discourse contexts and speaker intentions to modal bases should be adapted to this removal. This makes my treatment of so-called epistemic uses of modals conservative, in the sense that I do not appeal to machinery that is not also required for other modals.

### 3.3.1 The conceptual framework

First, relationality. We have long been comfortable with the relationality, or relativity, of modal concepts, including circumstantial ones (cf. Kratzer 1977; White 1975, 175–79). Relationality does not generally lead us to adopt an epistemic analysis of a particular modal concept either. So I think there should be no obstacle to recognizing that *might* could express a sort of circumstantial possibility that is also relational. John’s being contagious would not be a circumstantial possibility from any absolute standpoint, then, but instead relative to sets of circumstances that are compatible—in some sense—with his actually being so.<sup>25</sup>

The sense of compatibility relevant to possibility is often left undefined, even for epistemic possibility. A reason for this, as Bach (2011) notes, is that if compatibility is understood strictly in terms of logical consistency, then many things will count as possibilities that in most contexts we do not want to count as such. Authors often talk of compatibility in terms of “not being ruled out”, but what it is for something to not be ruled out is also left undefined (cf. Bach 2011, 24–25; MacFarlane 2011, 145).

In this chapter, I thus remain agnostic on the sense of compatibility that is relevant to circumstantial possibility. While I allude to an analysis of circumstantial possibility in Chapter 7, for now I simply assume it is some sort of relation that holds between sets of circumstances and propositions, and one that is paralleled in a relation of those same propositions to information states that *represent* those circumstances. What is at stake in this chapter is whether the relevant possibility relation for analysis of *might* is epistemic or circumstantial, and not what the relevant notion of compatibility for the analysis of any possibility relation will be. When I need to denote the circumstantial possibility relation, I use  $\mathcal{R}$  to do so.

Second, content and the evaluation matrix. The reason for removing the circumstantial *relata* of the possibility relation from content to the evaluation matrix is two-fold (cf. analogous proposals in Egan et al. 2005; Egan 2007; Beddor and Egan 2018; Lasersohn 2005; 2017; Stephenson 2007; MacFarlane 2011; 2014; Richard 2015). First, it respects the appearances. In order to have made the same claim with (145), speakers need to have in mind the same task for which John is ready.

(145) John is ready.

However, in order to have made the same claim with (146), it does not seem that speakers need to have in mind the same set of circumstances relative to which there is a possibility that John is contagious (cf. Braun 2013, 501–2).

(146) John might be contagious.

Similarly, a person who utters (145) will often seem to have said that John is ready *for* such-and-such a task. But the person who utters (146) will rarely seem to have

---

<sup>25</sup> Failure to recognize this relativity, if I am correct, leads Yalcin (2011) to his *non-factualism* about what he calls epistemic modality (and what I would call circumstantial modality). When I tell you that it might be raining, he asks skeptically, “am I ... to be understood as describing a way the world is?” (2011, 295). No, he thinks, and of course he is right, to an extent. But you are describing a way that a *part* of the world is.

said that John might be contagious *relative to* such-and-such circumstances.<sup>26</sup> And it turns out that *might* is not distinct from other circumstantial modal terms in this regard. For example, a sports analyst and a layperson can be judged to have made the same claim with (147), even though the set of circumstances relative to (and in virtue of) which they think that John can win will obviously not be the same.

(147) John can win tomorrow's race.

Removing circumstances from content to the evaluation matrix also opens up a certain desirable under-specificity. We can make attributions of content without having to find specific sets of circumstances to be part of that content. In some cases, of course, specific sets of circumstances may be part of content, as when they are made explicit with phrases like *given John's symptoms* or *with respect to John's medical history*. These phrases, I assume, act as intensional operators, taking contents which have truth values at worlds only relative to sets of circumstances within them and yielding contents which have absolute truth values at those worlds. Apart from phrases like these, however, I assume that the circumstantial relatum is always part of the evaluation matrix, never part of content.

Third, discourse contexts. Discourse contexts provide specific worlds, times, and locations for the contents of the claims made in those contexts to be evaluated against: namely, the worlds, times, and locations of those contexts. Yet it is not clear that there is any good candidate for "the circumstance-set" of a context (cf. von Fintel and Gillies 2011 for an analogous proposal about information states). Perhaps the only plausible candidate for the presumed circumstance-set of a context is the common ground (cf. Stalnaker 1978), supplemented by the physically salient circumstances. However, while this will certainly *constrain* what sets of circumstances the content of a possibility claim may be evaluated against, it seems not to antecedently *determine* any specific, unique set. This for two reasons.

One, the contents of the possibility claims themselves seem to have a role to play in determining what sets of circumstances are relevant to their evaluation. Within a given discourse context, for example, the different contents of the different claims made with (148) will naturally be evaluated against different sets of circumstances.

(148) John might be {contagious/a serial killer/at home/etc.}.

In this regard, locations provide an instructive contrast. Independently of the different contents of the different claims made with (149), these claims will naturally initially be evaluated against the location of the discourse context (apart from extra-contextual cues to evaluate them otherwise).

(149) It's {raining/cloudy/breezy/etc.}.

---

<sup>26</sup> In the literature, speakers are sometimes glossed as having said some such thing, but with information states replaced for circumstances (cf. Braun 2012, 461). This seems to me tendentious. The content of this gloss is certainly typically inferable from a speaker's claim. And we may be led for theoretical reasons to posit it as the content of that claim too. However, this is not *obviously* the intuitive content of that claim (cf. Yablo 2011, 271–72). Instead, it seems to involve an attribution of overly determinate content.

Two, discourse contexts are often compatible with multiple interpretations of a circumstantial modal claim, even of a given flavor (cf. Kratzer 1977; 1981, 294–95). As Kratzer writes, “the kind of facts we take into account for circumstantial modality are a rather slippery matter [which] may give rise to misunderstandings and jokes” (1981, 303). This seems correct, and it suggests that the common ground does not antecedently determine a specific, unique set of circumstances for the contents of those claims to be evaluated against. The phenomena of disagreement with modal claims also suggests as much (cf. Chapter 4). Disagreement does not typically center around a set of circumstances antecedently determined—somehow—by the discourse context, and on whether that set bears a particular modal relation to the prejacent. Instead, disagreement often seems to arise insofar as interlocutors are taking into account different sets of circumstances as they evaluate the content of the modal claim.

Finally, speaker intentions. These may sometimes override the features of a discourse context in determining a value for a parameter of an evaluation matrix, as in free indirect discourse (cf. Eckardt 2015). However, in the case of circumstance-sets, I see no reason to think that the speaker’s intentions, any more than the discourse context, need determine a specific set of circumstances for the content of her possibility claim to be evaluated against. Speakers may have certain *negative* intentions that constrain how their claims are to be understood. But I see no reason to think that the lack of a determinate positive intention is any more an obstacle to understanding than the lack of a determinate positive intention for the interpretation of *people* is an obstacle to understanding the claim made with (150).

(150) People are contracting contagious diseases.

Which people are contracting contagious diseases? Does it have to be people in general, people at random, people of a specific type, people *x*, *y*, and *z*? I do not think speakers need have answers to these questions in advance of uttering (150). Similarly, I do not think speakers need to have answers in advance of uttering (146) or (151) as to the question of what circumstances are compatible with John’s being contagious.

(151) Circumstances are compatible with John’s being contagious.

Perhaps it will be suggested that the speaker does have a specific set in mind, namely, the *relevant* circumstances. In a sense this is correct, of course, but not in any sense that is relevant. For decisions of relevance are made on the fly. If the phrase is used to pick out the specific set of circumstances that a speaker *has* determined so far to be relevant, I see no reason to think that the speaker intends the content of her possibility claim to be evaluated against only this set. And if the phrase is used to pick out the set the speaker *would* determine to be relevant, we create a justificatory problem. For often a speaker could not know in advance what (sort of) circumstances she would take to be so. Appeal to relevance seems to be more the tool of a theoretician than the content of a speaker’s intentions. Better to hold that speakers can

have general, indeterminate intentions that are guided by considerations *of* relevance than to hold that they have specific, determinate ones *about* relevance.

Sometimes it is suggested that the relevant set of circumstances for evaluating the content of a possibility claim are the known ones (cf. Kratzer 1981; 1991). In a sense this is also correct, in that speakers cannot evaluate the content of possibility claims on the basis of circumstances unknown to them (though they can withhold evaluation on that basis). Yet there are problems with attributing as the content of a speaker's intentions what happens *de facto*, one of which again is justificatory. For the considerations that lead us to assume that a speaker intends the content of her possibility claim to be evaluated against the known circumstances lead us to extend the group of knowers indefinitely (cf. MacFarlane 2011), and even sometimes to include *unknown* circumstances, supposedly on the basis of their knowability (cf. Hacking 1967; Teller 1972; DeRose 1991). But surely it is the antecedent relevance of circumstances that makes their subsequent knowability matter, and not their antecedent knowability that makes them subsequently relevant. We again do better to simply assign general, indeterminate intentions to speakers than to assign specific, determinate ones.

We should avoid the trap for theorists of thinking that just because the content of a possibility claim has a truth value assigned only relative to an extended evaluation matrix, then either a discourse context or a speaker's intentions must determine a unique evaluation matrix for that content to be evaluated against. Sometimes neither may do so.

These four points I have made allow us to separate the content of a possibility judgment, and the truth conditions for that content, from the state that a speaker is typically in when she makes such a judgment. Speakers may typically entertain the content of a possibility judgment only when they are uncertain whether its prejacent is true. But this can be attributed to interest in possibility typically being ancillary to unsatisfied interest in actuality (or to interest in the compatibility of a proposition with the circumstances typically being ancillary to interest in its truth). It need not—and if the argument of this dissertation is correct, *should* not—be attributed to a truth-conditional dependency on that uncertainty.

These points I have made are in some ways not novel. Instead, they combine select features of several epistemic analyses, including relativism (for the point about evaluation matrices), cloudy contextualism (for the point about discourse contexts), and Yalcin's expressivism (for the point about speaker intentions). Several of them are also already standard fare in the Kratzerian analysis of circumstantial modalities.

The novelty of these points—and also, I would argue, their success—derives from their combination in a non-epistemic analysis of *might*. For it is natural to think that discourse contexts *could* provide determinate information states for contents to be evaluated against: for example, the joint information state of the interlocutors in that context. It is also natural to think that speakers' intentions about information states would be determinate too. That said, if unaccompanied by these indeterminacy proposals, the proposal to remove information states from content to the evaluation matrix leads to a number of incorrect predictions, as will be shown in the following chapter. The proposals thus need to be combined, and it is by replacing an epistemic

analysis with a non-epistemic one that we can do so, finally effectively severing the contents of possibility judgments from the states of those who make them.

### 3.3.2 The formal framework

The proposals I have advocated can be developed formally within the standard Kratzerian framework for modal semantics (cf. Kratzer 1977; 1981; 1991), albeit with some adaptations. Formalization can allow for greater precision, it can be hoped. Argumentatively, nothing should depend upon it. Alternative formalizations could be chosen, and in the final chapter I allude to one such alternative.

To model modal relativity, I follow Kratzer in adding what I will call a *fact*-parameter to the evaluation matrix for the contents of modal sentences (Kratzer also adds a *norm*-parameter, relevant only in the final chapters). This parameter will range over functions from world-time pairs to sets of facts, or circumstances. The reason for the parameter to range over functions, instead of over sets of facts directly, is to formally ensure contingency for modal contents (cf. Kratzer 1991). If the world-parameter shifted independently of the fact-parameter, we could not do so.<sup>27</sup>

The functions the fact-parameter ranges over are subject to two constraints. First, the set of facts a function assigns to a world-time pair must be a (potentially proper) subset of the facts of that world up till that time. That is, there is no assignment of “future facts”. Second, the set of facts assigned must exclude purely logically constructed ones. There may well be primitive conjunctive, disjunctive, and conditional facts: that is, facts that are best represented by conjunctions, disjunctions, and conditionals. Just as well, there are conjunctions, disjunctions, conditionals, etc., that do not represent primitive facts at all: for example, disjunctions formed randomly by iteratively adding disjuncts to true propositions. These purely logically constructed facts, I assume, do not play any real role in circumstantial modal relations and so should be excluded from artificially doing so either by being categorically excluded from assignments of facts to world-time pairs.

So much by way of preface. The contents of modal sentences will now be assigned truth values—“1” for true, “0” for false—relative not simply to possible worlds (or world-time pairs) but also to sets of facts. To represent the assignment, we enclose sentences in double brackets superscripted with the relevant parameters of the evaluation matrix, followed by an indication of the conditions for truth at those parameters. To represent the content itself, defined in terms of its truth conditions, we use double brackets subscripted with the cent sign. Given these conventions, we use (152) to begin to formalize the proposal from the previous section.

$$(152) \llbracket \textit{might-p} \rrbracket^{w,t,f} = 1 \text{ iff } \mathcal{R}(f(w,t), \llbracket p \rrbracket_c)$$

In (152), what  $\mathcal{R}$  denotes is the circumstantial possibility relation, however this relation is ultimately to be analyzed.

---

<sup>27</sup> Yalcin (2007) proposes a world-independent parameter to capture the apparent world-independence of allegedly epistemic modals under attitude verbs. But, first, what is key to Yalcin’s proposal is not the world-independence of the parameter but instead the attitude verb’s shifting of its value (cf. Harr 2014; Ninan 2018). And, second, the apparent world-independence is an illusion that arises from the assumption of epistemicity anyway (cf. Chapter 5 of this dissertation).

While I will not attempt to provide an analysis of the circumstantial possibility relation here, I do want to make two comments about it vis-à-vis the entry in (152).

First, the entry in (152) implicitly assumes that this relation holds between sets of circumstances and *eternal* propositions, or propositions that do not vary in truth-value across times. This is a simplifying assumption that is standard in the philosophical literature but that will also be revisited in Chapter 7. For within the linguistics literature, there *is* significant debate whether modals like *might* combine with tensed complements or instead untensed ones (Condoravdi 2002; Hacquard 2010; Rullmann and Matthewson 2018).

Second, the entry in (152) does *not* assume that the possibility relation must be analyzed in terms of truth at some possible world. Braun (2013) has argued that what he calls the epistemic possibility relation (and what I would call a circumstantial one) cannot be analyzed in this way (2013, 489–90). While I will not enter into this debate here, it is worth noting that my entry is compatible with Braun’s argument. There is no reference in (152) to the prejacent’s being true at some world accessible from the evaluation world (though nothing precludes this analysis of the  $\mathcal{R}$  relation either). This, then, is my first major adaptation to the Kratzerian framework.<sup>28</sup>

But set aside questions about the correct analysis of the circumstantial possibility relation. With (152), we have a semantic entry: the conditions for the content of a *might*-sentence to be true at a given evaluation matrix. We do not, however, have what is sometimes called a *post-semantic* entry (cf. MacFarlane 2003): the conditions for a *claim* with that content to be true in a given discourse context. I have argued that discourse contexts do not antecedently determine specific sets of facts (or functions to them) in the way that they determine specific worlds, times, and locations. However, this does not mean that a specific set cannot be settled upon as relevant as a result of discourse. The value of the fact-parameter can be resolved through conversational negotiation, that is, in much the same way that relevant standards of taste, precision, etc., can also be resolved through such negotiation. Once this has been done in a discourse context, we can assign a truth-value to the possibility claim’s content in that context and so can also assign a truth value to the claim itself.<sup>29</sup>

---

<sup>28</sup> I am grateful to Paul Pietroski for questions that encouraged me to make this point explicit. One argument for thinking that the relevant possibility relation is not to be analyzed in terms of truth at a possible world comes from claims like *Hesperus might not be Phosphorus* (Paul Pietroski, p.c.). Intuitively, this claim seems as though it could be true if made by some astronomer of the past (cf. Frege 1948; Kripke 1980). However, if Kripke is right that *Hesperus* and *Phosphorus* pick out one and the same object across all possible worlds, then at *no* world is the prejacent true, a problem if we are analyzing possibility in terms of truth at possible worlds. That said, I am not entirely sure that this argument shows anything about the relevant possibility relation itself, as opposed to about the function of proper names in the relevant possibility claim. If metaphysical claims like *Hesperus didn’t have to be Phosphorus* are necessarily false, while so-called epistemic claims like *Hesperus might not be Phosphorus* are possibly true, this could be due to proper names somehow functioning semantically or pragmatically differently in the two cases: as Kripkean rigid designators in the former case, for example, but crucially not the latter. How this would work, and the theory of proper names in general, is not my concern here (cf., among others, the two-dimensionalist proposal of Chalmers 2006). My point is simply that the relevant case could lead us to revisit *either* our theory of the relevant possibility relation *or* our theory of proper names.

<sup>29</sup> Note that I keep the assignment of truth values to contents distinct from the assignment of truth values to claims with those contents. Relativists do not, unhappily so as we will see in Chapter 4, Section 4.5.

The set of facts settled upon as relevant in a given discourse context may be a subset of the known facts. Then again, it may not be. For sometimes we are unwilling to accept a *might*-claim as true until some unknown fact has been established, as in (153) (cf. also the examples in DeRose 1991; Yanovich 2014).

- (153) If John's been working at the hospital this week, *then* he might be contagious. But he may also still be on retreat in isolation at the hermitage.

Yanovich (2014) has some nice suggestions for how we settle on a set of facts as relevant in a given discourse context, as well as for how we set the bounds of a discourse context. These matters are not my concern here. For now, I would simply note that, once a given evaluation matrix *has* been settled on as relevant, we can reason about the contingent facts that would make the content of a claim true at that evaluation matrix. We can reason about how the facts, our information state, the laws, etc., could have been different. (This is crucially not the same as reasoning about how different facts, information states, laws, etc., could have been relevant.) And this allows us to reason about what makes a claim true, namely, by reasoning about what makes its content true at this matrix.

Before closing this section, I should note that, in recent work, Kratzer has rejected the use of the sort of parameter that I am here depending on (cf. Kratzer 2012; 2013). Let me briefly explain why she has done so and why I do not follow her lead.<sup>30</sup>

### 3.3.3 The formal framework revisited

Kratzer had hoped to intensionally characterize the functions her parameter ranged over as being either epistemic or circumstantial, and to use this intensional difference to characterize the meaning difference between sentences like (154) and (155).

- (154) John might easily be a military man.  
(155) John can easily be a military man.

By characterizing the meaning difference in this way—that is, as a difference in the intensional characterization of the value of a contextual parameter—Kratzer was able to avoid having to lexicalize the difference between the different flavors of modality expressed (cf. Kratzer 1977). While individual lexical items like *might* and *can* place idiosyncratic selectional restrictions on admissible values of the contextual parameter (cf. Kratzer 1991, 649–50), it would be possible in principle for a single lexical item to express *both* flavors without having to posit an ambiguity. Yet Nauze (2008) convinced Kratzer that the values of her parameter were not able to be intensionally

---

<sup>30</sup> I should also note that, while Kratzer had used a parameter to represent modal relativity, she seems to have viewed its value as part of content (Kratzer 1991, 640–41). As far as I can tell, this is because parameters may be used not only to represent that some truth-conditional contributor is part of the evaluation matrix, rather than of content, but also to represent that a truth-conditional contributor is not syntactically represented (cf. Hacquard 2010, 84–85; von Stechow and Heim 2011, 110), a view that Kratzer herself seems to hold (cf. Schaffer 2011, 202). I will not consider questions of syntactic representation here but will continue to use parameters exclusively to represent the content/evaluation matrix distinction.

characterized in the way she required for her non-ambiguity project to succeed (cf. Kratzer 2012, 24). Hence, Kratzer rejected the use of this parameter altogether.

Nauze's basic argument was this. The functions that Kratzer's parameter ranges over are all of them functions from possible worlds to sets of propositions. But no set of propositions is inherently epistemic or circumstantial, and so no function from possible worlds to sets of propositions is inherently such either. That is, no set of propositions construed as the content of a knowledge state could not also, in principle, be construed as the circumstances, and vice versa. Hence, Kratzer concludes, nothing allows us "to single out some . . . functions as epistemic, but not circumstantial, or the other way around" (2012, 24). Nothing about a given function *itself* allows us to characterize it as either epistemic or circumstantial: a function is "epistemic" or "circumstantial" insofar as *we* use it to represent an epistemic or circumstantial meaning, not vice versa. As a result, the differences between the modal flavors expressed by *might* and *can* cannot be attributed to any contextual variation in the value of the fact-parameter.<sup>31</sup> The values of the fact-parameter cannot themselves be distinguished in the way necessary for this to happen. Kratzer thus abandoned the use of her fact-parameter, and looked elsewhere to generate the meaning difference between sentences like (154) and (155), which she still took to be a difference in "the kinds of facts [the two modals] depend on" (2012, 24).

Neither of Kratzer's hopes for her fact-parameter are mine. I thus do not see any reason, in the light of Nauze's arguments, to abandon its use. Obviously, I do not think that we should intensionally characterize the values of her parameter for *might* as epistemic ones, or as functions to the *known* facts. However, I also do not think that the difference between (154) and (155) is best analyzed as a difference in intensional characterization of facts either. I think the difference is not in the kind of *facts* the modal relates the prejacent to, but instead in the kind of *relation* to those facts. This is suggested by the persistence of that meaning difference even when the relevant facts appear to be specified, and to be the same, as in (156) and (157).

(156) Given only his psychiatric profile, John might easily be a military man.

(157) Given only his psychiatric profile, John can easily be a military man.

Sharing neither of Kratzer's hopes, I thus see no objection to continuing to use her fact-parameter.

That said, if we are to do so, it is important to understand Nauze's criticism well. This criticism is not of the use of intensionally characterized functions to represent meaning *per se*. Rather, it is a criticism of the attempt to do so via context rather than the lexicon. Any intensional characterization intended to distinguish between modal flavors, Nauze argues, belongs there (Nauze 2008, 154–58).

Understanding this point is important, for it allows us to recognize that phrases like *given John's symptoms* can also be used to intensionally characterize the relevant function for a *might*-claim. However, the intensional characterization of this function does not determine a modal flavor, as Kratzer would have it do, but instead a specific modality of an antecedently given flavor. Adapting terminology from Kratzer (1981),

---

<sup>31</sup> Though such variation can be held responsible for differences between claims of the *same* flavor (cf. Viebahn and Vetter 2016).

such an intensionally characterized function may usefully be described as providing a *specific modal background* for a modal claim of a given flavor to be made against. Speaker intentions may sometimes do the same: for example, when it is clear, without being explicitly stated, that the content of a modal claim is to be evaluated against a specific type of facts. When there is no specific modal background, we may say that there is a *general*, or *generic*, one.

Say we refer to the set of facts against which the content of a modal claim is to be evaluated as the *modal base*, again adapting terminology from Kratzer (1981). Then we can say that in a discourse context with a specific modal background, the modal base is determined by the modal background. However, in a discourse context with a generic modal background, the modal base is constrained, but underdetermined, and it has to be resolved as the result of conversational negotiation.

At this point, it should be abundantly clear that I have made significant, further adaptations to Kratzer's framework. In particular, I am not deploying it as a way of avoiding ambiguity when accounting for modal flexibility, or the ability of a single modal to be used to express multiple flavors of modality. I am also not assuming that the meaning differences between modal flavors are always such as to be susceptible to a non-ambiguity account anyway. If *might* and *can* differ in the kind of modal relation that they express, and not in the kind of facts that they express that relation to, then it is not immediately clear that they could meaningfully be given the same semantic entry.

That said, it is important to be clear that these adaptations, like the previous one, are tangential, both to my argument that *might* is non-epistemic and to the conceptual framework developed in Section 3.3.1 for explaining how *might* could express a non-epistemic meaning. I do think, and will argue in Chapter 6, that very likely we will have to substantively rethink the project of attempting to account for modal flexibility without having to posit modal ambiguity. But this is due to Nauze's argument about the intensional characterization of functions, not to my argument about *might* being non-epistemic. If Kratzer (2012; 2013)'s attempt to account for modal flexibility via syntax rather than context could be made to work (cf. Hacquard 2006; 2010), then all my argument would require is that the alleged minimal difference between *might* and *can* in the kind of facts they depend on not be an epistemic/non-epistemic one. In her recent work, Kratzer herself sometimes even seems to concede this point, arguing on the basis of new data for locating a contentful/factual distinction *within* the traditional class of epistemics (cf. Kratzer 2012; 2013; Matthewson 2016).<sup>32</sup> What is important in all this, and what I continue to use Kratzer's standard fact-parameter to represent, is only that *might*, like *can*, can be used in different contexts to make claims that depend for their truth on the potentially different circumstances that are settled upon as relevant in those contexts.

### 3.4 Do we need an epistemic overlay for the analysis?

The analysis of *might* developed so far has been entirely in terms of circumstantial possibility. Even proponents of epistemic analyses should recognize such a possibility

---

<sup>32</sup> I say *seems* because the support rests on a distinction between evidence and content that I am not sure does the work it is required to (Kratzer 2012, 34–36; 2013). For further discussion, cf. Chapter 6, Section 6.2.2.2.

relation, I believe. For say that people judge that John might be contagious. Invariably, they do so on the basis of their judgments about the circumstances and how they stand with regard to the proposition that John *is* contagious.

Yet it may be wondered whether *might*-claims can be analyzed entirely in terms of the circumstantial possibility relation, or whether we also need to appeal to what I will call the *epistemic overlay* of this relation. That is, it may be wondered whether at some point we must appeal to the *known* circumstances. I consider two common reasons for thinking that an epistemic analysis of *might*-claims is required below and argue that neither forces us to such an analysis. In fact, in both cases, a non-epistemic analysis fares just as well, if not better, in accounting for the relevant phenomena.

### 3.4.1 Non-epistemic truth conditions and triviality

The speaker who utters (158) does not typically make a trivial claim.

(158) John might be contagious.

But say that the content of that claim had circumstantial truth conditions. That is, say that its truth depended on how John's being contagious stood with regard to the circumstances. Well, the circumstances in fact *settle* whether John is contagious. In the circumstances, that is, John either is contagious, or he is not. Hence, his being contagious could be a possibility with regard to the circumstances if and only if it was an actualized possibility: that is, if and only if John actually was contagious. But then our modal claim would describe a trivial possibility, and would turn out to be truth-conditionally equivalent to the non-modal claim made with (159).

(159) John is contagious.

However, the addition of *might* is not trivial in this way. Therefore, so the objection goes, the content of our claim could not have circumstantial truth conditions but must instead have epistemic ones. For John's being contagious *could* be a non-trivial possibility with regard to our state of knowledge even if it could not be so with regard to the state of the world (cf. Condoravdi 2002, a.o.).<sup>33</sup>

This objection assumes an absolute view of circumstantial possibility. It assumes that whether a proposition counts as a circumstantial possibility depends on how it stands with regard to the *total* circumstances. Hence, it assumes that only propositions whose truth values have not yet been metaphysically settled (or determined) can count as non-trivial circumstantial possibilities. However, I see no reason to assume that circumstantial possibility must be absolute. Why not think a settled proposition could count as a non-trivial possibility relative to a subset of the total circumstances, provided they excluded the relevant fact of the matter about that proposition?

---

<sup>33</sup> Condoravdi writes of sentences like (158), "where reference is to the present, [that] the modals have only an epistemic reading: [John's being contagious] is compatible with what the speaker takes the actual world to be now. The issue itself of his [being contagious] is settled by the course of events in the actual world though it may not be known which way it is settled. He either [is contagious] now or [is] not, and nothing that happens from now on can change whatever is in fact the case" (Condoravdi 2002, 21).

Here is an example. Say that circumstantial possibility claims are claims that a set of circumstances increases the objective probability of a prejacent above a certain low (but contextually variable) threshold. While the total circumstances of course trivially increase or decrease the objective probability of a settled proposition, subsets of the circumstances could non-trivially do so, provided they excluded the relevant fact of the matter. That John works in the worst hospital ward, for example, non-trivially increases the objective probability that he is contagious. That he has a powerful autoimmune system non-trivially decreases it. And so on. Relative to any of these relevant subsets of the total circumstances, we could say that John's being contagious was a non-trivial possibility, provided they had sufficiently increased the objective probability of his being so.

Say that a settled proposition *could* count as a possibility relative to a subset of the total circumstances. Then all we require in order to interpret a speaker's claim as non-trivial is to assume that she intends its *modal base*—or the set of circumstances relative to which its prejacent is to be evaluated as a possibility—to exclude the relevant fact of the matter.<sup>34</sup> All we require, that is, is the assumption that the speaker has a certain negative intention to not be talking about a trivial possibility, even if her positive intentions are general and indeterminate (cf. Section 3.3.1).

Is this cheating? Can the one fact that would make a *might*-claim trivial simply be excluded from the modal base? Do we not instead need to appeal to a specific modal background that could ensure this exclusion—for example, an epistemic background that references the *known* facts?

It is not cheating.<sup>35</sup> For it is safe to assume that speakers are generally unlikely to be interested in trivial instances of possibility relations. Take the objective probability example. Speakers are generally more likely to be interested in—hence more likely to be making claims about—the non-trivial effects of subsets of the circumstances on the objective probability of John's being contagious than in the trivial effect of the total circumstances on the objective probability of his being so. In fact, *whatever* the correct analysis of circumstantial possibility, whether in terms of probability or in terms of something else, the point remains the same. If the relations circumstances bear to settled propositions have both trivial and non-trivial instances, speakers are more likely to be interested in the latter than in the former. Hence, they are more likely to be making general, non-trivial claims like (160) than to be making trivial claims like (161).

(160) There's a non-trivial possibility that John is contagious.

(161) There's a trivial possibility that John is contagious.

The assumption that speakers are not only more likely to be interested in, but also more likely to be making claims about, non-trivial instances of circumstantial relations is one we have the latitude to make. For, in most cases, the modal base for a

---

<sup>34</sup> The same sort of principle holds for necessity claims, as will become significant later on.

<sup>35</sup> And even if it were, an epistemic background would not help. For as we will see from the following chapter, the considerations that could lead us to embrace an epistemic background would lead us to embrace one that referenced the facts known by potentially anyone. This would result in triviality in any case in which we did not think that the prejacent itself was inherently unknowable—that is, in almost every case!

claim is not explicitly specified. There are exceptions, to be sure, as when we delimit a modal base by using a modifier, as in (162), thereby providing a specific modal background. We can also constrain a modal base, without delimiting it, by adding a proposition to the common ground, as in (163).

- (162) Given his {work environment/medical history}, John might be contagious.  
 (163) John works in the worst hospital ward. He might be contagious.

In most cases, however, the modal base will at best be underspecified. And given an underspecified claim with competing trivial and non-trivial interpretations, it makes sense to prefer the latter (cf. Condoravdi 2002). People do make trivial claims, but generally only to indirectly make non-trivial points. But if an underspecified claim could be used *directly* to make a non-trivial point, there is no reason to think it is being used *indirectly* to do so, unless a speaker has reason to be coy.<sup>36</sup> The relevant fact of the matter is thus excluded from a circumstantial modal base by general pragmatic principles and not, as it may have seemed, by *fiat*.

There is an important connection to be made here to an attempt to ensure non-triviality via a semantic constraint on modal bases, the Diversity Condition of Condoravdi (2002). On my account, what is required for non-triviality is that the modal base exclude the relevant fact of the matter. On Condoravdi's account, what is required is that the set of worlds *determined* by the modal base—the worlds in which all the circumstances of the modal base obtain—not settle it (2002, 25–26).<sup>37</sup> (This is the constraint imposed by the Diversity Condition.) Condoravdi's requirement entails mine, but not vice versa. It is possible, after all, that circumstances excluding the relevant fact of the matter nonetheless settle it. A claim could thus be non-trivial on my account without satisfying the Diversity Condition. This is a good thing, for a limitation of the Diversity Condition is that it does not apply naturally to necessity claims, which *do* require the set of worlds they quantify over to settle the relevant fact of the matter (cf. Klecha 2016; Werner 2006), whereas my requirement applies naturally to necessity and possibility claims alike. To be non-trivial, both simply need the relevant fact of the matter to be excluded from the modal base, even if ultimately settled by the other circumstances in it. My approach to the triviality problem may thus be preferable to the standard approach independently of the argument against epistemic truth conditions from Chapter 2.<sup>38</sup>

<sup>36</sup> In some cases, a speaker does have reason to be coy, as when she is prohibited from making a public commitment about the relevant fact of the matter. In such a case, an underspecified possibility claim that could be interpreted as trivially true but defended as non-trivially so could allow a sympathetic speaker to communicate information she otherwise could not.

<sup>37</sup> Condoravdi uses the term *modal base* to refer to what I call the set of possible worlds determined by the modal base. For related ways in which the term is used, cf. von Stechow and Heim (2011, 42).

<sup>38</sup> It may also help explain why felicitous necessity claims generally require the speaker's evidence for the prejacent to be indirect (von Stechow and Gillies 2010; Matthewson 2015). Say direct evidence for circumstances that non-trivially settle the prejacent is indirect evidence for the prejacent itself. And say that direct evidence is stronger than indirect evidence, and that speakers are generally expected to make claims that their strongest evidence about the prejacent supports (Faller 2012; Mandelkern 2017, 120–21). A speaker making a necessity claim can thus often be expected not to have direct evidence for the prejacent itself, but only for circumstances that non-trivially settle it. Yet this expectation is waived in contexts in which a speaker's explicit concern is to show *how* certain circumstances settle

### 3.4.2 Non-epistemic truth conditions and speaker uncertainty

Say that a speaker knows whether John is contagious. It will then generally seem inappropriate for her to utter (158), repeated here as (164).

(164) John might be contagious.

This is predicted on an epistemic analysis of (164), for the speaker will have implied with her claim that she does not know whether John is contagious, which is false. However, it might be objected, there should be no such misleading implication on a non-epistemic analysis. For a speaker, the objection goes, should be able to claim that a prejacent bears some non-trivial relation to the circumstances without implying, misleadingly or not, that she is uncertain about it. If this is correct, then, to make sense of uncertainty generally being a condition for appropriate utterances of (164), we should take these utterances to have epistemic truth conditions, not circumstantial ones (cf. DeRose 1991).<sup>39</sup>

A problem with this objection is that it overlooks the exceptions to the uncertainty generalization. These are difficult to make sense of given epistemic truth conditions. And they also help to characterize the unexceptional cases in a way that makes sense of how uncertainty implications could be pragmatically generated from circumstantial possibility claims, in essentially Gricean ways.

The first sort of exception involves evasion. Even if it is common knowledge that the prejacent is false, a speaker may appropriately make a *might*-claim, as in (165), if her concern is to avoid public commitment on whether it is true (cf. Yablo 2011).

[Context: *S1 has been caught looking without permission at medical records indicating that John is not contagious. He tries to take control of the situation.*]

- (165) a. [S1] Do you know whether John is contagious?  
b. [S2] I do.  
c. [S1] Is he?  
d. [S2] He might be. But it's not my place to say.

The second sort of exception involves instruction. Even if it is common knowledge that the prejacent is false, a speaker may again appropriately make a *might*-claim, as in (166), if her concern is with whether it is a possibility relative to some more or less well-defined set of circumstances (cf. von Stechow and Gillies 2008; Bach 2011).

[Context: *Sue knows her cadets have a copy of her legend indicating that the practice field about to be charged is not one in which fake landmines are planted.*]

---

the prejacent. And, in these deductive contexts, the requirement for indirect evidence is waived too (Goodhue 2017). Goodhue argues that epistemic analyses require a distinctive deductive modal to account for these contexts. A circumstantial analysis, however, can account for them pragmatically.

<sup>39</sup> DeRose writes: the “prospects for explaining why [speakers] say what they do seem much brighter if we suppose that the content of their . . . modal statements somehow involves something about their own epistemic positions with respect to the [prejacent]” (DeRose 1991, 583–84).

- (166) Are you sure you want to charge the enemy? Look at that field. There might be landmines. Think of how many troops you could lose if there were.

To explain cases of evasion and instruction, epistemic theorists have usually assumed that the speaker's claim is not about her own knowledge state, but instead an interlocutor's (cf. von Fintel and Gillies 2008; Bach 2011; Dowell 2011). However, this will not work in cases involving common knowledge, like (165) and (166), since a speaker no more takes her interlocutors to be uncertain about the prejacent than she takes herself to be. Further, in these cases, there is also no other contextually salient individual whose knowledge state the speaker's claim could plausibly be about.

Instead of trying to provide epistemic interpretations for the *might*-claims in (165)-(166), the epistemic theorist could hold that *might* has a different type of reading in these cases. However, the usual alternative candidates to epistemic readings will not work. The claims are not metaphysical ones about ways in which the world could yet develop. Nor are they counterfactual ones about ways in which the world could have developed. Instead, they seem to be claims about possible ways in which the world *has* developed to this point. But if these claims about the space of present possibilities are not epistemic ones, as I have argued they do not seem to be, then presumably they must be circumstantial ones. They must be claims about the present possibilities left open by the circumstances: not the total circumstances, of course, but instead some subset of them.

Once we have granted circumstantial readings to make sense of the exceptions to the uncertainty generalization, can we also make sense of the generalization itself with only such readings? We can, given reasonable expectations about speakers making underspecified modal claims. Let me explain.

With underspecified claims, there seems to be a general expectation that a speaker is making what we might call an *all-things-considered* judgment, a judgment that takes into account all of the circumstances known to her.<sup>40</sup> Speakers can take into account fewer circumstances in making a modal judgment, of course, but to do so they generally need to make this explicit: to use a modifier that provides a specific modal background to delimit the modal base, as in (167).

- (167) Given his *symptoms*, John might be contagious. Given his *test results*, he can't be.

In cases of instruction, however, a speaker often does not need to use a modifier to delimit the modal base. Her evident concern with a specific relative possibility (the possibility of landmines relative only to the terrain, for example) allows her claim to support the specific interpretation she intends without the use of such a modifier.

And in cases of evasion, a speaker does not want to commit herself to any specific delimitation of the modal base. Her evident concern to avoid public commitment on the prejacent's truth value makes her underspecified claim, which does not support any specific interpretation, but is bound to be true on at least one, the perfect vehicle

---

<sup>40</sup> Cf. Mandelkern (2017, 120–21) for the related point that speakers are expected to take into account their strongest evidence.

for evasion. In both cases, *context* makes it clear that a speaker may not be making an all-things-considered judgment in making her underspecified claim.

These exceptional cases aside, speakers can generally be expected to be making all-things-considered judgments when making underspecified modal claims. That is, they can generally be expected to be making their modal judgments—just like their non-modal ones—on the basis of all the relevant circumstances known to them. Given this expectation, we can explain why it should generally seem inappropriate for speakers to make *might*-claims unless they are uncertain about their prejacent. For the all-things-considered judgment of speakers who *were* certain would be either that the prejacent was a trivial possibility or that it was a trivial impossibility. And these are judgments that speakers are either unlikely or unable to make with *might*-claims.

Perhaps it will be objected that if a speaker is expected to take into account all the circumstances known to her in making an underspecified modal claim, then this must be because that claim's truth depends, somehow, on *which* circumstances are known. Yet this objection proves too much. (It also conflates the grounds for making a claim with what makes that claim true.) For it is not just with underspecified *might*-claims that speakers are expected to take into account all the circumstances known to them, and to make all-things-considered judgments, but with underspecified *can*-claims too. And here there is no serious question of adopting an epistemic overlay for the circumstantial analysis.

Take the ability claim made with (168), for example.

(168) John can speak Finnish.

Say I know that John has the right anatomy to speak Finnish but not the right training to do so (cf. Lewis 1976, 150). It would then generally be inappropriate for me to utter (168). For while John counts as having the ability to speak Finnish relative to some of the circumstances I am aware of, he does not count as such relative to all of them. To take into account fewer circumstances, I should generally use a modifier to provide a specific modal background that delimits the modal base, as in (169).

(169) Given his *anatomy*, John can speak Finnish. Given his lack of *training*, he can't.

Here is another case, this one inspired by Kratzer (1981)'s example of a trombone player (cf. Vetter 2013, 6–7). Say I know that John has the right training to swim but not the right frame of mind to do. (Imagine he is currently inebriated.) Absent this second piece of information, I could have used (170) to express my modal judgment about John. Yet given this second piece of information, it would generally be more appropriate to use (171) or (172) instead.

(170) John can swim.

(171) In general, John can swim. Today he can't.

(172) Given his *training*, John can swim. Given his current state of *inebriation*, he can't.

With both underspecified *might*- and *can*-claims, then, a speaker is generally expected to be making all-things-considered judgments: judgments that take into account, or are made on the basis of, all of her information (or all the circumstances known to her). This expectation follows from the basic pragmatics of assertion (cf. Grice 1989), and is not limited to the modal case. However, making an all-things-considered modal judgment is crucially *not* the same as making a judgment about a modal relation that involves one's information (as will become important in Chapter 5 when we turn to embedded occurrences of *might* used to report assertions or beliefs). Our all-things-considered judgments always reflect the limits of our knowledge. But this does not mean, of course, that they are ever in any way about that knowledge.<sup>41</sup>

To summarize: circumstantial truth conditions allow us to explain both the uncertainty generalization and the exceptions to it, given general expectations about speakers making underspecified claims. Epistemic truth conditions do not. Instead, we would have to posit epistemic readings to make sense of the generalization and circumstantial ones to make sense of the exceptions to it. Hence, my approach to the uncertainty problem may again be preferable to the standard approach independently of the argument against epistemic truth conditions from Chapter 2.

### 3.5 Conclusion

Given the (in)admissible explanations for *might*-claims, a non-epistemic interpretation of *might* seems to be necessary. However, it may also have seemed to be impossible, given standard objections to any such interpretation. In this chapter, I have thus shown not only how a non-epistemic interpretation does indeed make ready sense of the data from Chapter 2, but also how the standard objections to it are misplaced. In explaining the data and responding to these objections, my focus has been on the general form of a non-epistemic analysis, as opposed to on the details of any particular analysis of this form. This is appropriate, given that the central argument of this dissertation is against epistemic analyses in general rather than in favor of any non-epistemic analysis in particular. That said, I will allude to a particular non-epistemic analysis of *might* in the concluding chapter. Before that, however, I want to show the promise of non-epistemic analyses in general in making sense of some well-known data that have proven troublesome for epistemic analyses.

---

<sup>41</sup> For an example of a *non*-modal, all-things-considered judgment, consider a moral one. Say that I rank one standard of behavior higher than another. And say that I judge some action acceptable relative only to one standard. My all-things-considered judgment about the acceptability of that action is then based on the standard I rank higher. And an underspecified claim about that action's acceptability—one that does not specify the relevant standard of behavior—can also generally be expected to express that judgment. For example, it will generally be inappropriate for me to utter (i) if I think that only relative to some inferior standard of behavior is it good to care for the quarantined.

- (i) It is good to care for the quarantined.

If I judge an action to be acceptable only relative to some lower-ranked standard known to me, then I should make this explicit, unless it can be easily inferred. That said, even if an all-things-considered judgement is made on the basis of the highest-ranked standard known to me, it is not a judgment about that standard *qua* known. All-things-considered moral judgments reflect the limits of our knowledge without necessarily being about that knowledge.

## **Part Two: Applications**

## Chapter Four: Applying the Analysis to Modal Disagreement

[T]wo speakers can disagree about a statement containing an epistemic modal simply because they have different knowledge states.

~ Stephenson (2007, 492)

### 4.1 A puzzle about disagreement

A central puzzle of discourse dynamics for epistemic analyses of *might* is how to make sense of apparent disagreement about what might be the case (cf. Stephenson 2007; MacFarlane 2014; Willer 2015; a.o.). It seems that such disagreement could sometimes be both genuine and warranted, even when it occurs between strangers, as in (173).

- (173) [*Context: Two hotel guests have struck up a conversation in the lobby.*]
- a. That concierge looks pretty suspicious. He might be a spy.
  - b. You're wrong! He can't be a spy! This hotel is more rigorous in vetting their employees than even the top government agencies!

Yet it is not immediately obvious on an epistemic analysis of *might* how the disagreement in (173) could ever be both genuine *and* warranted. For say the speakers are making claims about their respective information states. Then their disagreement would seem to not be genuine. But say instead that they are making claims about the same information state: their joint one, for example. Then their claims would seem to not be warranted. For, having just met, neither speaker is in a position to make a warranted claim about the other's information state. The disagreement in (173) *does* seem to be legitimate, that is for sure, and yet the basis for this legitimacy is not what we would expect for an epistemic claim. Instead of being divergent views about some information state, the basis for legitimate disagreement, as Stephenson suggests, seems simply to be the speakers' divergent information states themselves. This is puzzling.

In this chapter, I argue that the best way to solve this apparent puzzle is to *dissolve* it, by adopting a circumstantial analysis of *might* in place of an epistemic one. On this analysis, there is no more puzzle about how the disagreement in (173) could ever be both genuine and warranted than there is about how the disagreement in (174) could be so.

- (174) [*Context: Two sports fans have struck up a conversation at the local bar.*]
- a. Phelps can win yet another Olympic medal this year. He's so good.
  - b. You're wrong! He can't win another medal! He's got plantar fasciitis!

In both cases, the disagreement is about what possibilities or potentials the circumstances leave open. Speakers' information states will obviously affect what possibilities and potentials they think these are. However, their information states do not contribute in any way either to the contents of the modal claims made with (173) or (174), or to the truth conditions for those contents. Speakers thus do not have to

converge on a single information state for their disagreement to be genuine, and so also do not have to be in privileged position with regard to each other's information states in order for such disagreement to be warranted. Making sense of disagreement between strangers about what might the case is thus no more difficult, if we adopt a circumstantial analysis for *might*, than is making sense of disagreement between strangers about what can happen (cf. Willer 2013 for a similar comparison to non-modal disagreement).<sup>42</sup>

The same cannot be said, unfortunately, for an epistemic analysis of *might*. Faced with disagreement that unpredictably appears to be both genuine and warranted, proponents of epistemic analyses have had to go to great lengths to deny either the appearances or their unpredictability. Support for these lengths is supposed to come from data about the range of responses available to apparent disagreement. Yet the exact same range is available for apparent disagreement about what can happen, and this latter sort of disagreement clearly does not concern a proposition with epistemic truth conditions. Without the presupposition that the former does so either, the innovations of epistemic analysts seem as unnecessary and undesirable in the one case as they clearly are in the other.

Here is the structure my argument in this chapter will take. In Section 4.2, I critically review the arguments of those who would question the genuineness of the disagreement in (173), or would redefine what is required for disagreement to be genuine. In Section 4.3, I turn to the arguments of those who would maintain that the disagreement is genuine, as is standardly defined, but would deny that there is any special problem of warrant, or would redefine their semantics and/or pragmatics to avoid such a problem. In Sections 4.4 and 4.5, I present the circumstantial analysis of disagreement and related phenomena. And in Section 4.6, I conclude by responding to an objection.

Before anything else, I should note that the basic puzzle about disagreement, if not carefully formulated, is easily obscured. That is why I have focused on a case in which the apparent disagreement occurs between complete strangers who have struck up a conversation, rather than on cases in which it occurs between friends or between strangers who have butted into one. For in cases like these, it has been argued, the conditions *could* be met for speakers to make warranted claims about others' information states, and so also for there to be warranted and genuine disagreement about those information states too (cf. Dowell 2011, 8–9, 12–13). But the puzzle about disagreement is not really a puzzle about whether, or how, there could be warranted claims about others' information states, despite passing comments that could erroneously suggest otherwise.<sup>43</sup> Instead, the puzzle is whether the requirements for legitimate disagreement about what might be the case are the same as those for legitimate disagreement about an information state. And the puzzle is best brought out by cases, like (173), in which they intuitively do not seem to be the same.

---

<sup>42</sup> My proposal is similar to Willer's in taking there to be a single, non-epistemic content that is the object of disagreement. However, Willer takes this content to be dynamic, a context-change potential. A consequence of this view is that disagreeing with a *might*-claim requires believing its prejacent to be false (2013, 60–61). This seems too strong.

<sup>43</sup> von Fintel and Gillies (2011), for example, claim of a generic case involving individuals *A* and *B* that "it is obvious that *A* has no business making assertions about *B*'s information state" (2011, 115, fn. 13). In fact, this is not obvious, as Dowell rightly points out (2011, 8–9, 12–13).

## 4.2 Some attempted epistemic solutions to the puzzle

### 4.2.1 Denying the genuineness of disagreement

The first response of many in the literature to cases of apparent disagreement like (173) is to try to maintain that the disagreement is *only* apparent (cf. Braun 2012; Bach 2011; Dowell 2011; Schaffer 2011; Portner 2009; Dietz 2008; von Fintel and Gillies 2008; Wright 2007). On this line of response, it is assumed that the second speaker may have denied a proposition involving his own information state, or perhaps even the prejacent, but that neither of these is what the first speaker has asserted (if she has asserted anything at all; cf. von Fintel and Gillies 2011). If this assumption is correct, then there is no single content that has been both asserted and denied in a case like (173). Hence, the apparent disagreement would be *only* apparent.

#### 4.2.1.1 Evidence from entrenchment to one's own information state

Evidence for thinking that disagreement is only apparent is supposed to come from the so-called entrenchment data: responses like (175) and (176), which are available as continuations of the dialogue in (173), and which are supposed to show that the second speaker has misunderstood the first (cf., respectively, Schaffer 2011, 213; von Fintel and Gillies 2008, 81).

- (175) Well look, I just meant that *I* thought he might be a spy. I didn't know about this hotel's vetting process.
- (176) Well look, I wasn't saying that he *was* a spy. I was just saying that he *might* be. And he might have been one.

How exactly do (175) and (176) show the first speaker has been misunderstood? The thought is that with both of these responses a speaker is asserting an epistemic possibility proposition involving her *past* information state. The details of how she does this vary from response to response,<sup>44</sup> but the important point is that in making such an assertion a speaker is supposed to be re-asserting the truth of the proposition she originally committed herself to with ((173)a). Her past information state *was* compatible with the concierge's being a spy. And her interlocutor, whatever he has shown, has not shown the opposite. Instead, he has denied some content other than the one the speaker committed herself to, as her responses in (175) and (176) signify.<sup>45</sup>

The evidence from the so-called entrenchment data, if genuine, would prove too much. For entrenchment occurs not just in debates about what might be the case but also in debates about what can happen. (177) and (178), for example, are exactly parallel to (175) and (176), and are also available as continuations of the dialogue in (174).

---

<sup>44</sup> In (175), this past information state is explicitly referenced with the propositional attitude vocabulary, and the interpretation of the embedded modal is presumed to be indexed to it. In (176), this past state is not explicitly referenced, but the aspectual morpheme is presumed to induce reference to some past information state, for which the speaker's is the obvious contextual choice.

<sup>45</sup> Or, more cautiously, as they purport to signify (cf. Dowell 2011, 11).

- (177) Well look, I just meant that *I* thought he could win. I didn't know about his plantar fasciitis.
- (178) Well look, I wasn't saying that he *would* win. I was just saying that he *could*. And he could have won.

What is a speaker doing with (177) or (178)? One thing is clear: she is not reasserting some epistemic proposition that she originally asserted with ((174)a). For while *Phelps can win* is used to make an ability claim, *I thought Phelps could win* is used to make a propositional attitude claim and *Phelps could have won* to make a counterfactual one. Responses like (177) and (178) thus give us no reason to think that genuine disagreement has not occurred in (174). If anything, they give us reason to think the exact opposite. For it seems that with (177) and (178), a speaker has already implicitly conceded the falsity of her past claim and is choosing instead to defend its warrant, or reasonableness, given the information that was available to her. This is clearest with the propositional attitude claim, yet also seems to happen with the counterfactual one too. (179), for example, makes explicit how asserting a counterfactual can play a role in defending a claim as warranted.

- (179) Ok, so Phelps can't win. But he could have won if he hadn't had plantar fasciitis. And I didn't know that he did.

If entrenchment is used in the case of (174) to defend a past claim not as true, but instead as warranted, can we conclude anything about how it is used in the case of (173)? By far the simplest hypothesis is that (175) and (176) are used in the same way as (177) and (178), allowing for a uniform analysis both of the examples as a whole and of the individual constructions that make them up. However, even if we cannot rule out the alternative hypothesis that they are *not* being used in the same way, we also cannot rule out the simpler hypothesis either. And this is all that matters for our purposes here. The basic point is that a speaker's appeal to her past information state could be used to defend the warrant of her claim, rather than its truth. The so-called entrenchment data showcase the ability to make such an appeal, but this ability by itself provides no evidence that genuine disagreement has not occurred (cf. also MacFarlane 2014, 258–60).

#### 4.2.1.2 Evidence from dismissal of others' information states

Better evidence for whether genuine disagreement has occurred is a speaker's ability to dismiss others' information states as irrelevant to her claim, as with the response to the eavesdropper in ((180)c) (cf. Egan 2007; MacFarlane 2011).

- (180) [*Context: A guest is addressing his friend in the hotel lobby.*]  
a. [S1] Everything we know is compatible with that concierge's being a spy.

[*An eavesdropping hotel guest interjects.*]

- b. [S2] No, it's not! *I* happen to know this hotel only hires confirmed loyalists!
- c. [S1] Hey now, I wasn't talking about what *you* know! I was talking about what me and my friend do.

With ((180)c), S1 makes clear that her original claim was not about S2's information state, and so is not contradicted by the facts about that state that S2 has presented. S2 has misunderstood the topic of S1's claim, and S1 can appropriately dismiss S2's information state as irrelevant to it. The apparent disagreement between them is only apparent.

When it comes to possibility claims, however, a speaker does not seem to have the same ability to dismiss others' information states as irrelevant. The response to the eavesdropper in ((181)c) seems to change the topic of the original claim rather than to clarify it.

(181) [*Context: A guest is addressing his friend in the hotel lobby.*]

- a. [S1] That concierge might be a spy.

[*Another hotel guest interjects.*]

- b. [S2] No, he can't be! *I* happen to know this hotel only hires confirmed loyalists!
- c. [S1] #Hey now, I wasn't talking about what might be the case given what *you* know! I was talking about what might be the case given what me and my friend do.

Say that, in attempting to defend her claim, a speaker appears to change its topic. This would suggest that her interlocutors have not misunderstood her claim in their apparent disagreement with it. Say further that the way in which a speaker appears to change the topic is by dismissing others' information states as irrelevant. This would suggest that those other states *are* relevant, and that her interlocutors have certainly not misunderstood her claim in virtue of thinking as much. But if a speaker's interlocutors have not misunderstood her claim in their apparent disagreement with it, then this would suggest that that disagreement is in fact genuine. Thus, not only do the entrenchment data provide us no reason to think that genuine disagreement has not occurred, if I am correct, but the dismissal data also provide us reason to think that it has.

Perhaps the apparent topic-changing effect of ((181)c) may be attributed to the speaker's failure to reference any particular group, or their information state, in making her original claim. And perhaps subsequent explicit reference to what was left implicit will always appear to change the topic of the original claim. If so, one might object, then the apparent topic-changing effect that arises from dismissing others' information states as irrelevant in ((181)c) provides us no reason to think that genuine disagreement has occurred in (181).

This objection is flawed. For it is not in general true that making explicit material that had earlier been left implicit has the effect of apparently changing the topic, as (182) shows.

- (182) [*Context: A make-up artist is reporting to a campaign manager.*]
- a. [*Artist*] Your candidate's ready.
- [*A speech-writer interjects.*]
- b. [*Writer*] No, she's not! She's still making last-minute tweaks to the closing!
  - c. [*Artist*] Hey now, I wasn't talking about whether she was ready for her speech! I was talking about whether she was ready for the cameras.

It may be natural in ((182)b) for the speech-writer to interject as she does, given the potentially disastrous consequences of the manager's misunderstanding the artist. Still, there is nothing about ((182)c) to suggest that the artist is changing the topic of her claim rather than clarifying it. Making explicit a part of content that had earlier been implicit does not in general have a topic-changing effect.

Indeed, it is not clear that there is always even this effect when what is made explicit is an information state, as (183) shows.

- (183) [*Context: A guest is addressing his friend in the hotel lobby.*]
- a. [S1] It's plausible that that concierge is a spy.
- [*Another hotel guest interjects.*]
- b. [S2] No, it's not! *I* happen to know this hotel only hires confirmed loyalists!
  - c. [S1] Hey now, I wasn't talking about what's plausible given what *you* know! I was talking about what's plausible given what me and my friend do.

To my ears, at least, ((183)c) does not have the same topic-changing effect that ((181)c) does. Matters may not be as clear-cut as with ((182)c), but ((183)c) seems closer to clarifying the topic than to changing it.<sup>46</sup>

The difference between the topic-changing effect for *might* of dismissing others' information states, and the topic-clarifying effect for *plausible* of doing so, is even more pronounced in examples like (184) and (185).

- (184) [*Context: A British politician is addressing his fringe, Eurosceptic caucus.*]
- a. [*Pol.*] The prime minister might be a mole for the EU.

---

<sup>46</sup> If matters are *not* as clear-cut, this may be because the relevant group or information state should be taken to be part of the intended evaluation matrix for the claim made with ((183)a) and not, unlike in ((180)a), as part of its content (cf. Recanati 2007; MacFarlane 2014; Richard 2015). The thought would be that, even in specifying the intended evaluation matrix for your claim, you cannot help but change its content. But then if this explains any difference between (180) and (183), what explains the further difference with (181)?

[A government minister bounds onto the stage.]

- b. [Min.] Rubbish! The prime minister simply *cannot* be a mole for the EU!
- c. [Pol.] #Hey now, I wasn't talking about whether she might be a mole from *your* standpoint. I was talking about whether she might be one from *ours*.

(185) [Context: A British politician is addressing his fringe, Eurosceptic caucus.]

- a. [Pol.] It's plausible that the prime minister is a mole for the EU.

[A government minister bounds onto the stage.]

- b. [Min.] Rubbish! It's simply *not* plausible that the prime minister is a mole for the EU!
- c. [Pol.] Hey now, I wasn't talking about whether it's plausible that she's a mole from *your* standpoint. I was talking about whether it's plausible that she's one from *ours*.<sup>47</sup>

In both cases, the intended audience for the speaker's comments is a fringe group that finds certain things plausible and/or epistemically possible that the public at large will not. Yet in ((184)a), unlike in ((185)a), the speaker's claim requires him to engage with any who would challenge it, even if clearly *not* part of his intended audience. In defending the former, unlike the latter, he does not get to choose whose information state is relevant.

*Might*, it seems, is thus more similar to *can* than to *plausible*. In defending a *can*-claim, as in (186), speakers do not get to choose whose information is relevant to it either.

(186) [Context: A guest is addressing his friend in the hotel lobby.]

- a. [S1] Phelps can win yet another Olympic medal this year. He's so good.

[Another hotel guest interjects.]

- b. [S2] No, he can't! *I* happen to know that he's got plantar fasciitis!
- c. [S1] #Hey now, I wasn't talking about what can happen given what *you* know! I was talking about what can happen given what *me* and my friend do.

In attempting to dictate whose information is relevant, as in (186), a speaker seems to change, rather than to clarify, the topic of her original claim. In this case, this is obviously because the speaker's original claim was not about an information state to begin with. Thus, information states are relevant to that claim, not in terms of making it true, but instead in terms of providing the grounds for determining whether it is. And a speaker does not get to choose whose information states may be used to make *that* determination.

---

<sup>47</sup> Alternatives to *from {you/our} standpoint* include *as far as {your/our} information is concerned*, *given {your/our} knowledge base*, *given what {you/we} know*, etc.

Now a simple and compelling conclusion to draw is that the case of *might* is like the case of *can*—that a speaker does not get to choose whose information states are relevant because information does not play a metaphysical role in determining the truth of a claim but instead an epistemic one.<sup>48</sup> However, there is also no lack of ingenuity that has been motivated by the unchallenged presupposition of epistemicity. So the more conservative conclusion to be drawn from the dismissal data for the time is the following: if a *might*-claim is about information, a speaker seems not to get to decide *whose* information it is about, given the possibility of genuine disagreement with potentially anyone about that claim. And this makes the problem of warrant that much more difficult than it was to begin with (cf. MacFarlane 2011, 150–52).

#### 4.2.2 Redefining what is required for genuine disagreement

Several authors have recently made proposals that would skirt the problem of warrant altogether by redefining what is required for genuine disagreement (Montminy 2012; Huvenes 2015; Khoo 2015). These authors assume that a speaker performs multiple speech acts in making a *might*-claim (cf. Swanson 2006; von Stechow and Gillies 2007; Portner 2009, 172–77), and they suggest that disagreement could center around some secondary speech act. In this way, hearers’ information states could be relevant to a speaker’s claim even if it has as its asserted content a so-called solipsistic proposition: a proposition involving *only* the speaker’s information state. As a result, no problem of warrant in cases of genuine disagreement need ever arise.

As an example of this sort of proposal, consider Huvenes (2015). Huvenes suggests that in making a *might*-claim, a speaker not only asserts that the prejacent is compatible with her own information state but also “conveys the advice not to overlook the possibility that [the prejacent] is true, and to have a certain positive credence in [it]” (2015, 997).<sup>49</sup> Disagreement between this person and the person who makes a *can’t*-claim centers around this second speech act, and, specifically, on the different credence that each individual places in the prejacent (2015, 991–95).

Huvenes is correct that there is a sense in which people disagree if they place different levels of credence in the same proposition, as in (187), or even if they give different advice about the level of credence to place in that proposition, as in (188).<sup>50</sup>

- (187) a. For me, it’s plausible that that concierge is a spy.  
 b. For me, it’s implausible.

---

<sup>48</sup> Cf. Bach (2001, 29–30; 2005).

<sup>49</sup> It is striking, by the way, that Huvenes alludes to the possibility that the prejacent is true in developing his analysis. On my analysis, a speaker is *asserting* that there is such a possibility, not merely conveying the advice not to overlook it. But so strong is the unchallenged presupposition of epistemicity in the literature that authors like Huvenes—and he is not alone—can freely allude to the non-epistemic possibilities that would simplify their analysis without actually using them to do so.

<sup>50</sup> Huvenes focuses only on the former sort of disagreement, oddly enough. Indeed, along with most authors who advocate a multiple-speech-act proposal, he does not seem to entertain the possibility that a speaker who utters *can’t-p* also performs a secondary speech act, despite suggesting that it is a general feature of epistemic modals that they “can be used to perform two distinct speech acts” (2015, 997). I am not sure what explains this oversight.

- (188) a. I advise that we place some credence in that concierge’s being a spy.  
 b. I advise that we place none.

However, neither sort of disagreement licenses the standard markers of disagreement that we have seen in the modal cases. Inserting a sentence-initial *no* or *you’re wrong* in ((187)b) or ((188)b) makes each dialogue marked, sometimes quite noticeably so.

Recognizing this as a problem for his proposal, Huvenes claims that, in the modal cases, the disagreement markers could be “targeting” the prejacent. What he has in mind can be illustrated with examples like (189) (cf. Huvenes 2015, 995–96).

- (189) a. That concierge might be a spy.  
 b. {No / You’re wrong}, the concierge is not a spy.

In (189), the second speaker has denied the proposition that the first speaker claimed to be possible. This denial of the prejacent, not the denial of the asserted modal content, is what licenses the disagreement markers in (189). Something similar is at work in (190), Huvenes proposes.

- (190) {No / You’re wrong}, that concierge can’t be spy.

Here the speaker *does* deny a modal content, though not for Huvenes the same as the one that was asserted. Hence, this denial is not what licenses disagreement markers. Instead, they are licensed by the denial of the prejacent that is implied in the denial of this content. (If the speaker’s knowledge is incompatible with the concierge’s being a spy, then, given that knowledge is veridical, it follows that the concierge is not a spy.) In (190), as well as in (189), according to Huvenes, the disagreement markers are thus licensed because they “target” the prejacent.

There are two problems with this explanation for when disagreement markers are licensed. The first is that it does not make sense of the contrast between (191) and the modal cases we have considered so far.

- (191) a. It’s compatible with what I know that that concierge is a spy.  
 b. {#No / #You’re wrong}, it’s incompatible with what I know that that concierge is a spy.

With (191), we have what Huvenes would take to be the asserted content of the speakers’ modal claims in their apparent disagreement in (173) (cf. Huvenes 2015, 989–90). Yet unlike in that case, in this case disagreement markers are *not* licensed. The explanation of licensing in terms of prejacent-targeting would predict otherwise. For in both cases, the content that is allegedly denied entails that the prejacent is not true. The prejacent is thus equally “targeted” in both cases by the speaker’s denial.

The second problem with Huvenes’ explanation is that it does not make sense of modal disagreement that goes in the opposite direction, as in (192).

- (192) a. That concierge can’t be a spy.  
 b. {No / You’re wrong}, that concierge might be a spy.

With ((192)b), the content of the speaker's claim does not entail that the concierge *is* a spy. Nor, if it is about her own knowledge, does it entail that the previous speaker's knowledge must be compatible with the concierge's being a spy either, contrary to what that speaker is alleged to have claimed.<sup>51</sup> Hence, there is no clear way in which the disagreement markers could be licensed in this case, unless by disagreement about a single, non-epistemic modal content.

A different problem from Huvenes' for those who would make his same denial is not the identification of a secondary speech act that supports disagreement markers but rather the derivation of this act from the content of the primary one. Khoo (2015) provides an example of this problem, with his suggestion that the relevant secondary speech act in making a *might*-claim is a proposal that it "not be common ground that [its] prejacent is false" (2015, 528). In rejecting this proposal by virtue of asserting what is not supposed to be presupposed, as in (193), one does indeed find that disagreement markers seem to be licensed.

- (193) a. Let's not presuppose that that concierge is a spy.  
 b. {No / ?You're wrong}, that concierge is a spy.

Yet there is a problem with the derivation of this proposal from an assertion that includes an information state as part of its content. For, depending on the information state, either that content itself will not license disagreement markers when the prejacent is denied, or it will, but will also be able to be denied.

The problem is best brought out by Khoo's preferred identification of the relevant information state as a group one, which he identifies using the phrase *the best available evidence* (2015, 529–30). The phrase is ambiguous, as Price (1983) points out, and as Khoo also notes, but its most natural interpretation is a distributive one, in which the best available evidence includes evidence that is available to anyone in the group, even if not to everyone. This is the interpretation I take to be prominent in (194), and it seems to license disagreement markers when the prejacent is denied.<sup>52</sup>

- (194) a. It's compatible with the best available evidence that that concierge is not a spy.  
 b. {No / You're wrong}, that concierge is a spy.

However, on this interpretation the speaker who denies the prejacent is also in a position to deny the asserted content of the modal claim. So the disagreement markers in the modal case could just as well be licensed by disagreement centering on the primary speech act itself, the default hypothesis. The alternative interpretation, which

---

<sup>51</sup> This is the explanation that Huvenes proposes for dialogues like (ii), in which there is no disagreement with the prejacent of the first speaker's claim, but in fact the opposite (2015, 1004–5).

- (ii) a. That concierge can't be a spy.  
 b. {No / You're wrong}, he is a spy.

<sup>52</sup> If it does not, then the problems facing Khoo are at once more immediate and yet also easily inferred from the ones I describe below.

Khoo intends, is a collective one, and is made more prominent by phrases like *the shared evidence* or *the publicly available facts*. On this interpretation, however, disagreement markers are not licensed when the prejacent is denied, as (195) shows.

- (195) a. It's compatible with {our shared evidence / the publicly available facts} that that concierge is not a spy.  
 b. {#No / #You're wrong}, that concierge is a spy.<sup>53</sup>

The point about disagreement markers is even clearer when the proposition asserted is a solipsistic one, as in (196).

- (196) a. It's compatible with my evidence that that concierge is not a spy.  
 b. {#No / #You're wrong}, that concierge is a spy.

The problem facing Khoo is that both *might*-claims and compatibility claims should give rise to the proposal whose rejection—via denial of the prejacent—he takes to license disagreement markers (cf. Khoo 2015, 528–31). Yet while denial of the prejacent always licenses disagreement markers in the former case, in the latter it does so only when the speaker could disagree with the asserted content of the claim itself. Why is this? The simplest explanation is that in both cases the disagreement markers are licensed by disagreement with the asserted content itself, rather than by rejection of the proposal, or secondary speech act.

There may well be disagreement, as both Khoo and Huvenes propose, that centers around a secondary speech act performed in making a *might*-claim. There is no evidence, however, that *that* sort of disagreement licenses the same disagreement markers in all the same cases that disagreement with a *might*-claim does. In fact, there is every evidence to the contrary. Perhaps some proposal can be developed to avoid this mismatch. Until then, the simplest explanation for it is the one suggested above.

To summarize: in the apparent disagreement dialogue in (173), there is no clear evidence that genuine disagreement has not occurred, and also no clear evidence that it has centered around some secondary speech act. But if genuine disagreement *has* occurred, and if a speaker's interlocutors have not misunderstood her in thinking their information states relevant to her claim, then the speaker's warrant for making that claim is unclear. For any of a speaker's interlocutors, even her unintended ones, can disagree with her claim solely on the basis of their own information state (and without having to have any beliefs about hers). Thus, if the speaker's claim includes some sort of information state as part of its content, this gives the speaker's claim far wider scope than it would seem she could ever have warrant for. I turn to this problem of warrant now.

---

<sup>53</sup> Perhaps you hear the disagreement markers in ((195)b) as acceptable. If so, be careful that you are not interpreting phrases like *the shared evidence* as *everything known to any one of us* but instead as *anything known to every one of us*.

### 4.3 Other attempted epistemic solutions to the puzzle

#### 4.3.1 Denying the problem of warrant

Some theorists have maintained that—on closer inspection—there is no special problem of warrant that arises from taking an information state to be part of the content of a *might*-claim. Yanovich (2014), for example, argues that what is required to be warranted in asserting that it might be that  $p$  is the same as what is required to be warranted in asserting that  $p$ : namely, to think that one is “as good an authority” on whether  $p$  as anyone else (2014, 78–80). However, *pace* Yanovich, the two cases are clearly not the same on an epistemic analysis of *might*.

To assert that  $p$ , you must assume that you have a reasonably justified opinion about a settled matter. To assert that no one’s information rules out that  $p$ —which is what you are doing if you are making a claim about information to which potentially anyone’s information might be relevant (cf. Yanovich 2014, 78–80)—you must assume that *no one* has such an opinion. But often this assumption will not be warranted when you are agnostic about a matter, even though asserting that it might be that  $p$  will be. In the case of the concierge, for example, the first hotel guest need not assume that the second guest could have no privileged information in order to be justified in asserting that the concierge might be a spy, or even that he is. Yet he would need to make such an assumption if he claimed that their joint information did not rule out that the concierge was a spy. Asserting that  $p$  and asserting that it might be that  $p$  may indeed be equivalent in their requirements for warrant, as Yanovich suggests, but they would not be so if the modal claim were about information.

I should add, for what it is worth, that Yanovich distinguishes his particular epistemic analysis, which he labels *practical contextualism*, from an analysis that he labels *group contextualism* (2014, 69–82). On the former, the modal claim is supposed to be that no relevant piece of information rules out the prejacent; on the latter, that no information of anyone in the relevant group does. Yanovich claims that only on practical contextualism is the warrant for making a modal claim the same as for asserting its prejacent. On group contextualism, the warrant for the modal claim is supposed to be assuming that you are as good an authority as anyone else on your group’s *information state* (2014, 78–82).

As far as I can tell, Yanovich’s reason for thinking practical contextualism escapes the problem of warrant facing group contextualism is that it references information without referencing *whose* information it is (cf. the definition of knowledge in Yanovich 2014, 77). If this is correct, it is also irrelevant. For if I claim that no relevant piece of information rules out that  $p$ , and I allow that relevant pieces of information are able to be contributed by my interlocutors, then I must assume that my interlocutors have no privileged pieces of information just as much as if I had made a claim about our group’s information. So the problem of warrant persists no matter what epistemic contextualist analysis we adopt.

#### 4.3.2 Redefining what is required for warrant

##### 4.3.2.1 Cloudy contextualism

Some authors have proposed that the solution to the problem of warrant is to redefine what is required for it. One example of such a proposal can be found in von Fintel and Gillies (2011). These authors, you will recall, argue that the discourse context does

not typically provide a single, determinate context of evaluation for the interpretation of a *might*-claim. Instead, it leaves open a range of information states that could be relevant, and the *might*-claim is thus made against a “cloud of admissible [evaluation] contexts” delineated by these states (von Fintel and Gillies 2011, 117–24). Given this indeterminacy about the relevant state, the requirements for warranted *might*-claims are not the same as for claims made against a single, determinate evaluation context. Instead, a claim that may permissibly be interpreted as being about an information state on which the speaker does *not* have authority to pronounce can be warranted just as long as it may also permissibly be interpreted as being about an information state on which she *does* have the authority to pronounce—for example, her own.

von Fintel and Gillies’ solution to the problem of warrant does not seem promising to me, for three reasons. First, the solution rests upon the assumption that a speaker’s claim may permissibly be interpreted as being exclusively about her own information state, as the entrenchment data from Section 4.2.1.1 is supposed to have shown. Yet that interpretation of the data has been called into question here. Further the dismissal data in fact seem to have shown the exact opposite. But then why think that the only warrant required for a speaker to make a claim is that she takes some interpretation which that claim cannot have to be true?

Second, what creates the problem of warrant is the combination of an assumption and an observation. The assumption is that *might*-claims are to be interpreted as *de re* claims about information states. The unexpected observation, given this assumption, is that hearers systematically and appropriately tend to assess *might*-claims on the basis of their own information states. From this observation comes the problem of warrant. How could speakers ever be warranted in making claims about these states? In response, von Fintel and Gillies propose that claims that systematically and permissibly receive one interpretation are warranted (only) on the basis of another. But this is unnecessary. Better to drop the assumption that makes the observation unexpected in the first place than to try to accommodate it in this way.

By way of analogy, consider the case of *can*-claims, which hearers also tend to systematically and appropriately assess on the basis of their own information states, or on the basis of the sets of circumstances they believe to obtain. If we took such claims to be *de re* claims about those sets, we would also have a problem, given these assessment facts, of how speakers could ever be warranted in making such claims. We could, in response, adopt a von Fintel and Gillies style account of warrant. In this case, however, it is clear that the problem is only apparent: that it is generated entirely by the *de re* assumption and is best addressed by that assumption’s being abandoned. The same holds, if I am correct, in the previous case too. The von Fintel and Gillies solution is as needlessly complicated a way to make sense of the warrant for *might*-claims as it is for *can*-claims.

A third and final issue with von Fintel and Gillies’ solution to the problem of warrant is that it implies a level of intentionality on the part of speakers that is implausible. To see this, it will help to first consider the case of quantificational claims, which von Fintel and Gillies suggest may also be made against a cloud of admissible evaluation contexts (2011, 123–24).

With quantificational claims, there does not seem to be the systematic mismatch between hearer-assigned interpretations and speaker-warranted ones that there

allegedly is for *might*-claims.<sup>54</sup> Yet sometimes there can be such a mismatch, and it can be exploited as a way of gaining information (cf. von Fintel and Gillies 2011, 123–24). For example, consider a professor who knows (i) that every student in good standing was at last night’s department party and (ii) that Alex and Billy, who have been in bad standing, were not. What this professor does *not* know is (iii) whether Alex and Billy have continued on in the program despite their bad standing. Believing it would be impolite to ask directly, this professor utters (197) to the administrative assistant, knowing he has all the relevant records and will interpret (197) as (198), whereas the professor only has warrant for (199).

(197) Every student was at last night’s department party.

(198) Every registered student was at last night’s department party.

(199) Every student in good standing was at last night’s department party.

This mismatch between the hearer-assigned interpretation and the speaker-warranted one is just what the professor wants. For she knows that, given this mismatch, she will be able to discretely gain the desired information about Alex and Billy, however the department assistant responds.<sup>55</sup> And she knows too that, if the assistant disagrees with her claim, she can defend it by appealing to (199) as her intended interpretation for (197).

If von Fintel and Gillies are correct, then the sort of intentionality on display in the above example is on display with almost any *might*-claim. For there is a *systematic* mismatch, on their view, between the hearer-assigned interpretation of such claims and the speaker-warranted ones. Speakers further seem to actively exploit this mismatch on their view: they make *might*-claims “based just on [their] own evidence but ... [use them] as a probe or test or trial balloon into the hearer’s evidence” (von Fintel and Gillies 2011, 123). And this, it seems to me, is just implausible. Speakers do not display the same level of intentionality in making *might*-claims that the professor does in uttering (197).

Even if they did, this level of intentionality would be misplaced. Say my ultimate concern is with whether the concierge is a spy. On von Fintel and Gillies’ view, uttering (200), and exploiting the fact that it has no determinate interpretation in the discourse context, helps me attain my end. For I both express my own information state *and* probe other states on which I have no warrant to pronounce, thereby gaining useful information that moves me closer to what I am ultimately concerned to know.

(200) That concierge might be a spy.

---

<sup>54</sup> This by itself reinforces the criticism that von Fintel and Gillies have made the wrong choice in accounting for the warrant for *might*-claims. There is nothing problematic with the view that a claim may sometimes be warranted on the basis of an interpretation other than the one a hearer permissibly and appropriately assigns it. Yet to have to posit this mismatch on a systematic basis suggests that something in the analysis has gone awry.

<sup>55</sup> If the assistant agrees with her, then she will know that Alex and Billy have both left the program. And if he disagrees, stating that either Alex or Billy, or both, were not there, then she will know that either none or only one of them have left.

Yet in uttering either (201) or (202), I will also manage in many contexts to do both these things. I will make the determinate claim about my own information state that I do in (201) as a way of asking the determinate question about your information state that I ask with (202), and vice versa.

- (201) I don't know whether that concierge is a spy.  
 (202) Do you know whether that concierge is a spy?

Nothing at all seems to be gained by the alleged indeterminacy of (200).<sup>56</sup>

To summarize: an obvious objection to von Fintel and Gillies' solution to the problem of warrant is that it assumes that speakers' *might*-claims may permissibly be interpreted as being exclusively about their own information states. And this assumption, I have argued, is undermined by the dismissal data. But even if one does not accept this argument, von Fintel and Gillies' solution faces serious other problems which suggest a lack of fit with the data. A different solution to the problem of warrant is required.

#### 4.3.2.2 Relativism

An alternative solution to the problem of warrant—and the final epistemic solution to the puzzle about disagreement with which we began—is the relativist one (Egan et al. 2005; Egan 2007; Stephenson 2007; MacFarlane 2011; 2014). For relativists, you will recall, a speaker's *might*-claim does not include an information state as part of its content but instead has a truth value that is determined relative to *assessment contexts* individuated by information states. Since information states differ, and a prejacent will be compatible with the content of some but not others, a single claim may thus have different truth values at different assessment contexts for relativists.

Removing information states from content in the relativist way is supposed to explain how a speaker could ever be warranted in making a claim to which potentially anyone's information state could be relevant. Since the claim's truth value is relativized to assessment contexts that are individuated by information states, anyone's information state will naturally be relevant to it. And if a claim is false at someone's context, then she should reject that claim and express her disagreement with it.<sup>57</sup> However, as long as the claim is true at the speaker's context, and he justifiably believes it to be so, he would seem to be warranted in making it (cf. Egan et al. 2005, 153; MacFarlane 2011, 160–61; 2014, 255). For what you justifiably believe to be true, you are generally warranted in asserting. And this, in sum, is the relativist solution to the problem of warrant.

On the whole, I think the relativist is closer to the truth than any of the other theorists we have considered so far. As relativists suggest, we should not take information states to be part of the content of *might*-claims. Yet we should not take such claims to have truth values determined relative to such states either, contrary to what relativists suggest. It is one thing to hold that a speaker *assesses* a claim to be

<sup>56</sup> There are exceptions, as when a prejacent is not the potential answer to a polar question but instead a *wh*-one. These exceptions are what von Fintel and Gillies focus on in developing their theory.

<sup>57</sup> For further discussion of relativist disagreement, and why it is important that it center on the truth-value of the claim rather than of its content, cf. Richard (2015).

true, given his information state, and another altogether to hold that that claim *is* true at a context of assessment individuated by that information state.<sup>58</sup> Nothing pushes us to the latter in accounting for a speaker's warrant for a *might*-claim except the unchallenged presupposition of its epistemicity, as is perhaps expressed most clearly in MacFarlane (2011)'s statement of the "paradox" facing contextualism:

[A]lthough the truth of a claim made using epistemic modals must depend somehow on what is known—that is what makes it "epistemic"—it does not seem to depend on any *particular* body of knowledge (2011, 155).<sup>59</sup>

Given this unchallenged presupposition, the otherwise promising relativist solution to the problem of warrant turns out to be no solution at all. Let me explain.

In removing some alleged truth-conditional contributor from content to the evaluation matrix, relativists are not proposing that speakers are blind to the truth-conditional dependence of their claims on that contributor. In fact, quite the opposite (cf., a.o., Egan et al. 2005, 158; MacFarlane 2014, 255). For relativists, speakers recognize that their *might*-claims are true at their assessment contexts *in virtue of* their information states being a certain way; that their taste-claims are true at their assessment context *in virtue of* their standards of taste being a certain way; etc. But speakers also recognize that these claims could be false at their interlocutors' assessment contexts in virtue of their interlocutors having different information states, standards of taste, etc. (cf. Stephenson 2007, 508–11).

This recognition need not by itself affect the claims that speakers are willing to make. In claiming that tempeh is tasty, even while recognizing that this claim may be false at your assessment context, I could hope—if it turns out to be so—to register the difference between our assessment contexts, to make your context more like my own, etc. What *should* affect a speaker's willingness to make a claim, however, is the recognition that their own assessment contexts could very easily change, such that a claim true at their earlier context would be false at a later one. Given such a change, that claim would have to be retracted, given that the operative assessment context for evaluating one's own claims is always the current one for a speaker (MacFarlane 2014, chaps. 5.4, 5.7).

The mere possibility of having to later retract a claim should not by itself prevent a speaker from making it. However, as MacFarlane suggests, a high likelihood that one will immediately have to do so should (2011, 161, fn. 11).<sup>60</sup> And while one's tastes are not likely to change in an instant, one's information state is (especially in conversation, where information exchange is often the goal). Speakers should thus

---

<sup>58</sup> Information states play the former role, but they do not seem to play the latter. I can terminate interminable debate about whether it is plausible that *p* by saying, *well look, it's plausible for me*. I cannot terminate interminable debate about whether it is possible that *p* by saying, *well look, it's possible for me*. Information affects whether I assess a possibility claim to be true. It does not make that assessment so.

<sup>59</sup> Similar unquestioned presuppositions of epistemicity are to be found in other relativists, including Egan et al. (2005, 133) and Stephenson (2007, 487).

<sup>60</sup> It is worth noting that MacFarlane makes this suggestion, after having claimed that relativism avoids the problem of warrant, in an apparent concession to the fact that it might not so straightforwardly do so.

recognize a high likelihood that immediate retraction will be required for almost any *might*-claim they may make, and so should avoid making such claims, unless they presume that their claim is also true at their interlocutors' assessment contexts. Yet given such a presumption, speakers could just as well have made a claim about their and their interlocutors' joint information state. And it was precisely a problem of how speakers could typically have warrant for such a claim that relativism was trying—but seems equally unable—to avoid.

This last claim will seem counterintuitive. But I am not claiming that relativism about every domain will face the same problems of warrant that contextualism about that domain may. Rather, I am making this claim only for allegedly epistemic modal claims. In this case, in my opinion, the problem of warrant arises from the assumption of epistemicity rather than from contextualism itself. The problem with relativism, as with cloudy contextualism, is not necessarily the general framework, but instead its fit with the specific data.

Say, however, that you are unconvinced that the relativist solution to the problem of warrant is no solution at all. There would still be serious difficulties with it.

Here is a first difficulty. Some *might*-claims seem *unwarranted*: e.g., the one made with (203) (the style of example and general line of criticism are due to Dietz 2008).

- (203) Mathematicians all agree that Fermat's Last Theorem has been proved. But I don't believe them. I can't follow their proofs at all. That theorem might very well be false!

Yet relative to the assessment context of its speaker, a so-called "ignorant assessor", it would seem the *might*-claim made with (203) could both be true and justifiably believed to be so. Hence, the claim *would* be warranted. And yet it clearly does not seem to be so.

MacFarlane's response to this difficulty is to suggest that the relevant information state in a given assessment context need not always be the assessor's (2011, 175–76). If, for example, the relevant state for the ignorant assessor's context for (203) was the derided mathematicians', then the claim made with (203) would not be warranted. This response may seem unprincipled, but MacFarlane contends that it is not. For what determines the relevant information state for an assessment context in any given case, he claims, are the features of *that* context (2011, 176–77). In particular, MacFarlane seems to suggest that the assessor's interests in that context—whether simply "to guide her own inquiry" or instead to "determine whether [someone else] might be a trustworthy source of information"—are what determines the relevant information state (2011, 176–77). The former sort of interest, MacFarlane suggests, may make only the assessor's information state relevant, while the latter may make others' so.<sup>61</sup>

---

<sup>61</sup> MacFarlane's suggestion has recently been formalized as sensitivity to the question under discussion, or QUD (cf. Beddor and Egan 2018). For these authors, as well as for MacFarlane and Dietz, the primary focus with ignorant assessors is their assessment of others' modal claims, rather than their production of their own. This leads to slightly different issues from the issues of warrant I am considering here, but the general problems I have raised, and will raise, carry over. (They also carry over from MacFarlane to Beddor and Egan too.)

There are two problems with this suggestion. The first is that (203)'s speaker may have the sort of interests for which MacFarlane suggests only his own information state need be relevant in the assessment context. For example, he may continue on with (204).

(204) I guess it's time to crack open my algebra books and construct a counter-proof!

Perhaps MacFarlane would say that others' information states could also be relevant for this inquiry-guiding interest too, but this brings us to the second, more general problem. For MacFarlane has not said how it is that an assessor's interests determine the relevant information state for her assessment context. Is this determination subjective? Or is it objective, and something about which a speaker can be mistaken? Neither option is very promising. With the former, it is not clear why an ignorant assessor could not always maintain that it is her own information state that is relevant in her assessment context. With the latter, it is not clear how we would be justified in thinking that we know whose information state our interests determined to be relevant in our context. The former option leaves us with too many warranted claims, the latter with too few.<sup>62</sup>

Here is a second difficulty. MacFarlane explains the retraction of a warranted claim in terms of its falsity at the speaker's current state, not his prior one (cf. MacFarlane 2014, 256). Yet he also allows that retraction can sometimes be resisted despite a change in information state (2014, 258–60). MacFarlane does not give an example of what such resistance would look like, but from the context it seems that a response like (205) would be a good candidate.

(205) I stand by my claim. That concierge might have been a spy.

Now it is not clear to me that (205) should be understood as a refusal to retract, but it is also not clear that MacFarlane has any other option. For he does not want to interpret the claim made with the second sentence as an implicit epistemic counterfactual (2014, 253). And it is not clear for MacFarlane how a non-epistemic counterfactual could be relevant to the speaker's standing by her claim. Hence, MacFarlane is left to interpret (205) as being used to make a claim about the speaker's past information state (2014, 258–60, 271–72), and as implying that the original *might*-claim was a pragmatically enriched claim about that same state too (2014, 258–60, 271–75). Retraction is resisted because that pragmatically-enriched claim was an assessment-invariant one.

Such pragmatic enrichment is common, MacFarlane holds, and should not give us pause. I am less sure. There seems a fundamental difference between enrichment that merely changes the content of a claim and enrichment that changes whether a claim has an assessment-variant truth value or instead an assessment-invariant one. However, the point can be waived. For what *should* give us pause is that when a

---

<sup>62</sup> With this latter option, we might also just as well make an unmediated appeal to an objective relation and eliminate the relativization of truth to information-state-individuated assessment contexts altogether.

speaker can apparently resist retraction with (205), it seems that she could just as well have chosen to retract with a sentence like (206), something for which MacFarlane argues assessment-variant claims are required (2014, 260).

(206) Ok, I guess I was wrong. That concierge can't be a spy.

If this is correct, then it will often be indeterminate whether a speaker has made an assessment-variant claim or rather an assessment-invariant one. But it is unsatisfying, to say the least, to have to posit such an indeterminacy in order to explain the mutual availability of both responses. At this point, the original explanation of warrant and retraction in terms of assessment-sensitivity begins to seem otiose and baroque. Better simply to hold that, with the first response, the speaker stands by the warrant for her claim but not its truth. But for the claim to have been warranted, without having been true, we require it to have had a content with non-epistemic truth conditions.<sup>63</sup>

The source of these and other difficulties for relativists is their unchallenged presupposition of epistemicity, and the assumption it generates that the warrant for a claim that comes *from* an information state is to be explained in terms of the truth of the claim *at* that state. Without linking warrant to truth in this way, we can simply hold that a warranted claim is one that a speaker could reasonably (even if inaccurately) have believed to be true, while an unwarranted one she could only have unreasonably (even if accurately) believed to be so. We do not need to adopt a special theory of warrant that first relativizes the truth of *might*-claims to information-state-individuated assessment contexts only to then allow variability not only in the information states that individuate an assessment context (MacFarlane 2014, 260–61) but also in the relativization of claims to assessment contexts at all (cf. MacFarlane 2014, 258–60).

#### 4.4 Dissolving the puzzle on a non-epistemic analysis

On a non-epistemic analysis of *might*, disagreement about what might be the case is in principle much the same as disagreement about what can happen. Both are disagreement about the possibilities and potentials there are, given the circumstances.

What is required for disagreement about circumstantial modalities to be genuine? One thing that does *not* seem to be required is for those who disagree to have complete and accurate beliefs about the circumstances. In much the same way, complete and accurate beliefs about laws are not required for there to be genuine disagreement about what is legally permissible or obligatory, and complete and accurate beliefs about the composition of some domain are not required for there to be genuine disagreement about what is true of all or some or none of its members.

---

<sup>63</sup> Wright (2007, 272–75; 2008, 177–82), who is also skeptical of the explanation of warrant in terms of truth, wonders how you could ever move from a claim's being warranted, given its speaker's information state, to its being true at an assessment context individuated by that information state. What does truth at this original context buy us, he wonders, given that warranted claims often turn out untrue? The answer, if we accept truth-conditional dependence on information states, is that this assessment-sensitivity buys us the very warrant for those claims in the first place! The question is still on point, however, and suggests that we should simply reject such truth-conditional dependence. Truth at the original assessment context is an idle wheel.

Incomplete or inaccurate beliefs about any of these things will often be accompanied by correct *de re* beliefs—about what the propositions *believed* to be the circumstances or the laws leave open or close off, for example. But while these *de re* beliefs certainly help to explain why speakers make the claims they do, I see no reason to think that they provide the content or the intended truth conditions of those claims.<sup>64</sup> Instead, when speakers recognize the inaccuracy or incompleteness of the relevant beliefs which formed the basis for their claim, it is perfectly natural for them to respond by retracting that claim, as in (207) or (208).

- (207) a. [S1] John can win this morning's race.  
b. [S2] Actually, he got stoned last night.  
c. [S1] Oh, I didn't realize that. I guess I was wrong then. He *can't* win.
- (208) a. [S1] All the registered runners are here.  
b. [S2] Actually, we had two more runners register just last night.  
c. [S1] Oh, I didn't realize that. I guess I was wrong then. Not all the registered runners *are* here.

Genuine disagreement is thus not threatened by speakers' different information, I conclude. Yet it is threatened by speakers' different intentions. Imagine, for example, that a young David Lewis utters (209) to his friends as he pores over an anatomy book, and that his parents, overhearing, interject with (210) (cf. Lewis 1976, 150).

- (209) I can speak Finnish, you know.  
(210) Don't listen to him. He can't. He quit taking lessons less than half a month in.

Here there is no genuine disagreement. For Lewis had a specific type of circumstances in mind in making his modal claim. And he could make his intended meaning clear, and resolve the apparent disagreement, by uttering (211).

- (211) Given my *anatomy*, I can speak Finnish, or any human language for that matter.

Genuine disagreement about a circumstantial modal claim requires agreement about the intended modal background (whether specific or generic). It does not require agreement about the modal base.

Intentions are not always clear, whether from the context or to the speaker herself. Sometimes it may thus be difficult to tell whether disagreement is genuine or instead only apparent. Further difficulty may arise if intentions that are clear to a speaker appear to be in conflict with context (cf. Dowell 2011).<sup>65</sup> There may then be meta-

---

<sup>64</sup> Indeed, the results of thinking as much are distinctly bizarre. For example, in the case of the laws, the result would be that the only genuine disagreement is precisely the sort of interpretive disagreement for which we might think there may be no single correct answer.

<sup>65</sup> While discourse contexts do not determine unique sets of circumstances for contents of modal claims to be evaluated against (cf. Ch. 3.3.1), it is not implausible to think they constrain admissible ones (cf. von Stechow and Gillies 2011).

level disagreement that masquerades as object-level disagreement. That said, I see no reason to think that the opacity of intentions threatens the possibility of genuine disagreement in every single case. Instead, in most cases where speakers have not referenced a specific type of circumstances in making a modal claim, it is safe to assume that they are talking about the possibilities and potentials left open by circumstances *in general*. This seems to be the locus of genuine disagreement.

How could such disagreement ever be warranted, it might be wondered? For speakers could never be warranted in thinking that they know what all the circumstances are. Yet nor could speakers ever typically be warranted in thinking that they know what all the state's laws are. But this does not prevent the possibility of warranted (and genuine) disagreement about what the state's laws leave open. What is required for warrant in both cases is a reasonable, defeasible assumption that one knows the relevant laws and/or circumstances.<sup>66</sup> This assumption will not always be reasonable, and speakers will not always have warrant for the relevant claims, as (212) and (213) show.

- (212) I can't say whether Phelps can win. I know he's been making a lot of visits to his doctor recently, but I don't know what for.
- (213) I can't say whether you can shoot pigeons during hunting season. I know it's been legalized in Michigan, but with restrictions on who can obtain a license.

Yet while a speaker should always concede that there may be relevant laws and circumstances that she is unaware of, generally for her claim to lack warrant she will need positive reason to think there are such. (Or she will need the possibility she implicitly recognizes to be explicitly raised; cf. Lewis 1976, 354–55; 1996, 559–67.) Sometimes there will be such reason, but not always, and so a speaker may often assume warrant by default.

This account of warrant, it might be protested, is available to epistemic theorists just as well as to myself. For we could say that all that is required for warrant is a reasonable, defeasible assumption that the *known* laws and circumstances leave open a possibility, and that, while such an assumption may not always be reasonable, in general it will be. For generally a person will not have any reason to think she is in a worse position than everyone else, so that the known laws and circumstances are somehow *unknown* to her.

This protestation rests on an equivocation (cf. also the previous discussions of Khoo 2015; Yanovich 2014). If by *known*, all that we mean is *commonly known* or *publicly available*, then, yes, there is no special problem of warrant. But there will also be many cases where disagreement is not genuine either, as with the responses in (214) and (215), where one of the interlocutors has privileged access to uncommon knowledge.

- (214) You can't have known this, but I'm that concierge's father. And I can assure you, he simply can't be a spy. He's far too trusting and naïve for that.

---

<sup>66</sup> Most laws and circumstances will be relevant for relatively few modal claims.

- (215) You can't have known this, but I'm Phelps' doctor. And I can assure you, he simply can't win a medal. His recent injury's far too serious for that.

If such privileged access does not seem to threaten genuine disagreement, as I think it does not, then we have to interpret *known* as meaning something like *known by anyone* (cf. MacFarlane 2011, 150–52) And here more is required for warrant than is required in the non-epistemic case. Instead of needing positive reason to think that my information is relevantly incomplete, I now need positive reason to think it complete. By incorporating reference to knowledge into my modal claim, I thus incur a higher justificatory burden, just as I do in choosing to utter (217) rather than (216).

- (216) That concierge is a spy.  
(217) I know that that concierge is a spy.

I conclude, then, that while an epistemic analysis of *might* does not readily make sense of genuine and warranted disagreement about what might be the case, a non-epistemic one can do so, and as readily as it makes sense of such disagreement about what can happen.<sup>67</sup>

## 4.5 Truth-value judgments on a non-epistemic analysis

### 4.5.1 Truth-value judgments and the evaluation matrix

I have focused in this chapter on discourse-internal disagreement. However, there is also a sort of disagreement that is discourse-external, and that has created problems for standard epistemic analyses of *might* too. This is disagreement in truth-value judgments made outside of—but elicited by—a given discourse context (cf. Egan et al. 2005; Egan 2007; MacFarlane 2011; 2014). The basic problem is that such judgments can seem to conflict, while also all seeming appropriate. Following Knobe and Yalcin (2014), I will call this the problem of *extra-contextual assessments*.

The problem can be illustrated with a case from DeRose (1991). John has symptoms indicative of cancer. Tests have been run which rule out that possibility. John's wife, Jane, knows that the test results are available to his doctor, but she does not know yet what they are. In their separate discourse contexts, Jane and the doctor respectively utter (218) and (219), and we as theorists want to maintain that each has said something true.

---

<sup>67</sup> The two cases are not completely parallel, of course, for the possibilities that we express with *might* are not completely parallel to the potentials that we express with *can* either (cf. Chapters 6 and 7). In particular, the progression of time and the accumulation of circumstances will affect what potentials there are, but it will rarely affect the status as a non-trivial possibility of a previously settled proposition. Additionally, relative to the total current circumstances, a settled proposition can only be a trivial possibility. Yet the total current circumstances are evidence of past potentials, not necessarily current ones, and so there is not the same threat of triviality here, especially insofar as potentials can exist unrealized. The atemporality and potential triviality of the possibilities expressed by *might* can make them seem less objective, and so can make disagreement about them seem less objective too. Yet this is no reason to think that genuine disagreement cannot occur. What is key is to recognize that speakers' claims are generally to be interpreted as claims that a prejacent is a non-trivial possibility, given circumstances in general, and that disagreement centers on whether this is so.

- (218) John might have cancer.  
 (219) John can't have cancer.

From this starting point, imagine that a nurse reports to the doctor what Jane has said. The doctor responds by uttering (220), seemingly appropriately, we as theorists also want to maintain.

- (220) I know that Jane has different evidence than we do. And I know that her evidence is compatible with John's being contagious. Still, what she has said is false. John cannot be contagious.

Yet if Jane and the doctor have both said something true, then how could the doctor also appropriately say that what Jane has said is false?

The answer to this question is simple. When a proposition is asserted, we judge the truth of the *assertion* by looking to the context in which it was made. We look to see whether the asserted proposition is true at the evaluation matrix determined either in or by that context. But not all truth-value judgments are about the assertion. Instead, some are judgments about the asserted *proposition* itself, and abstract away from the context in which it was asserted. This seems to be the explanation for the apparently (but not genuinely) conflicting truth-value judgments in (221) and (222).

- (221) [*Context: King Leonidas addresses his compatriots at Thermopylae*]  
 The Athenian poets used to complain that no Spartan had ever contributed to the Panhellenic defense. Those poets said something true: no Spartan *had* done so. But today, thanks to your efforts, what they've said is now false.
- (222) Our suffragette forebears used to remonstrate that no woman had ever been elected to national office. Our forebears said something true. No woman ever *had* been elected. But today, in this historic year of women, what they've said is now false. Over fifty women have been elected!

If I am correct, this is also the explanation for the apparently conflicting truth-value judgments in the modal case. Some are judgments of the assertion itself, while others are judgments of the asserted proposition that abstract away from the context in which it was asserted and evaluate it at the current context.

Note that the crucial distinction in this explanation is between the assertion itself and the asserted proposition. It is *not* a distinction between different propositions that may be asserted in the use of a single sentence: for example, a sentence like (223).

- (223) I'm hungry.

If John's utterance of (223) is reported to Mary, then it is simply incoherent for her to respond with (224).

- (224) John said something true. But what he has said is false. For I'm not hungry.

Thus, the appearance of conflicting truth-value judgments seen in the cases above cannot be attributed to different propositions being asserted in the different uses of a sentence (cf. MacFarlane 2014, 8–11; MacFarlane cites Kölbel 2002, 39). Instead, it must be attributed to a difference between judgements about an *assertion* and about the asserted *proposition*.

In the modal case, this explanation reinforces the conclusion from the previous chapter that the modal base should be removed, from the asserted proposition, to the evaluation matrix against which that proposition is evaluated. It is this removal that makes it possible for judgments about the truth of an *assertion* of a modal proposition at a given context to not automatically determine judgments about the truth of the asserted *proposition* at another. For the asserted proposition has no absolute truth value at the evaluation world.

If we adopted an epistemic analysis of *might*, I do not see how we could avoid this conclusion about the modal base, given examples like the one above. For Jane's claim clearly has a limited scope, and the doctor seems not to have misunderstood this scope either. Yet if both these things are true, and if Jane's claim also included her evidence, or information state, as part of its content, then that content would have a uniform truth value at evaluation matrices within the actual world. Hence, either the claim made by Jane with (218), or the claim made by the doctor with (220), would have to be false. And neither seems to be so.

For perspective, contrast the modal case with another case of context-sensitivity. Imagine that Jane and the doctor live in different towns, and that each makes a true claim about their own town by uttering (225) and (226) in their respective discourse contexts.

(225) The local bars serve caviar.

(226) The local bars do not serve caviar.

If a nurse subsequently reports to the doctor what Jane has said, (227) does not seem at all appropriate for that doctor to utter, but instead evinces confusion or intentional obtuseness.

(227) I know that Jane lives in a different town than I do. And I know that the local bars in her town serve caviar. Still, what she has said is false. The local bars do not serve caviar.

There is a real problem for epistemic theorists in the contrast between the two cases, then, and the obvious solution is the relativist one of removing information states from the asserted proposition.

The solution is obvious, I say, and yet it has also been strenuously resisted (cf. Wright 2007; von Fintel and Gillies 2008; 2011; Dowell 2011; Braun 2012; Yanovich 2014). The reason is that relativists have paired the proposal to remove information states from the asserted proposition with a proposal that *should* be resisted. This is a proposal to relativize not just the truth of the asserted proposition, but also of the very assertion itself, to information-state-individuated evaluation matrices: i.e., to so-called assessment contexts (cf. Egan et al. 2005; Egan 2007; MacFarlane 2011; 2014;

Beddor and Egan 2018). The relativist proposal may seem counter-intuitive, given how I have presented and responded to the problem of extra-contextual assessments. And yet relativists have persistently (and maddeningly) conflated judgments about the truth of an asserted proposition with judgments about the truth of the assertion itself, as Beddor and Egan show in asking:

How should assessors evaluate this sort of centered *content* for truth or falsity? Here is a natural first thought: when assessing an *utterance*  $u$  that expresses a set of centered worlds  $p$ , ... (2018, 5; emphasis mine).<sup>68</sup>

Given this conflation, and given the relativization to assessment contexts, relativists predict truth-value judgments of false whenever extra-contextual assessors' information rules out the prejacent of a *might*-claim.<sup>69</sup> Yet this prediction has been challenged, both from the armchair and by empirical studies (cf. Yalcin 2011; Knobe and Yalcin 2014; Khoo 2015). To make sense of unpredicted variability in extra-contextual assessments, relativists have thus had to claim that the relevant information state for an assessment context is not necessarily the assessor's own, but instead determined by something like his interests (cf. MacFarlane 2014). In the most recent version of this proposal, also claimed to be the most predictive, this determination is sensitive to the QUD in the assessment context, which can include whether the original speaker was warranted (or competent) in making her claim (cf. Beddor and Egan 2018, secs. 4.2-4.3). The result is that the literal truth of a speaker's assertion at any given assessment context can depend on the assessor's determining that assertion to have been warranted!

Surely this is a *reductio* of the relativization of the truth values of assertions—rather than just of asserted propositions—to assessment contexts. Better to say simply that the relevant truth value judgments track different questions than to add that the answers to those questions literally determine the relative truth value of a speaker's assertion at an assessor's context. And the former *can* be done, without the latter, if we distinguish an assertion's truth at an assessment context from the truth of the asserted proposition at that context. Asserted propositions may have relative truth values; assertions of them need not. Extra-contextual assessments push us to remove the modal base from the asserted proposition. They do not push us to relativize the truth of that assertion itself to so-called assessment contexts.

#### 4.5.2 Binding and the evaluation matrix

A circumstantial analysis of *might*, I have argued, should be accompanied by a removal of the modal base from content to the evaluation matrix. Yet Schaffer (2011) argues that the modal base, which he presumes to be epistemic, should be taken to be part of content (2011, 203–6). I respond to that argument here.

---

<sup>68</sup> Beddor and Egan are not alone. Cf. also Egan et al. (2005, 154) and Weatherson and Egan (2011, 12–13). MacFarlane (2014) also tends toward conflation with his talk of propositions being true as used at one context and assessed from another. For is the assessment of the *proposition* or of its *use*?

<sup>69</sup> Cf. MacFarlane's claim that "only a relativist semantics can explain why earlier epistemic modal claims are always evaluated in light of what we know now (at the time of assessment), even when we know much more than was known at the time the claim was made" (2011, 173).

Schaffer's argument is based on what is sometimes called the *binding criterion* of Stanley (2000) (cf. Recanati 2002). The binding criterion is deployed to make sense of examples like (228), which are most naturally interpreted as being used to make claims about variable locations, as in (229), rather than as being used to make a claim about a constant location, as in (230).

- (228) No matter where I vacation, it rains.
- (229) No matter where I vacation, it rains there.
- (230) No matter where I vacation, it rains {in my hometown / in Chicago / etc.}.

The co-variation of the vacationing and raining events in the claim made with (228), Stanley argues, is evidence for the existence of an implicit location variable in the logical form of the embedded sentence, (231), that is being bound in (228) by the higher-scoping quantifier.

- (231) It rains.

And in general, the possibility of co-varying readings is supposed to provide evidence for the element that co-varies with the domain of quantification being the value of a variable, whether implicit or explicit, and so also being part of content rather than of the evaluation matrix (Stanley 2000, 410–14). This is the so-called binding criterion.

Against this backdrop, Schaffer provides a number of examples in which a *might*-sentence is embedded beneath a quantifier phrase, and in which it seems that there is co-variation of information states with the domain of quantification. These examples are given in (232)-(235).

- (232) Every boy has a father who might be a genius. [Schaffer's (72)]
- (233) Anytime you are going for a walk, if it might be raining, you should bring an umbrella. [≈ Schaffer's (70)]
- (234) Sometimes it is not the case that Billy [might] be guilty. [Schaffer's (73)]
- (235) Usually if Billy might be the one who stole the cookies, then Suzy might be too. [Schaffer's (74)]

In (232), co-variation is supposed to be of different individuals and their information states. Every boy has a father who, *as far as that boy knows*, is a genius. In (233)-(235), it is supposed to be of different times and an individual's information states at those times. For example, anytime you are going for a walk, you should bring an umbrella if, *as far as you know at that time*, it is raining. Both cases of co-variation are supposed to provide evidence for the view that the logical form of a sentence like (236) has an implicit variable ranging over information states. It is the binding of this variable that is supposed to explain the apparent co-variation in the cases above.

- (236) John might be a genius.

Schaffer's argument leads to the conclusion that the value of an information state variable is part of the content of the claim made with (236). This runs counter to both

my claim about the relevant modal base (that it is circumstantial, not epistemic) and to my claim about its role (that it is part of the evaluation matrix, not of content). Yet I do not think this is problematic. For Schaffer relies on Stanley's binding criterion in a surprisingly uncritical way (cf. Schaffer 2011, nn. 22, 41 for a somewhat cavalier dismissal of the very points of contention about this criterion).

Stanley's binding criterion is anything but uncontroversial. It has been criticized both for over-generating implicit variables in making sense of apparent co-variation (cf. Partee 1989; Cappelen and Lepore 2002; Recanati 2002; Sennet 2008), and for failing to account for other ways of making sense of this phenomenon too (in addition to the authors already cited, cf. Bach 2000; Neale 2000; Elbourne 2008). Both sorts of criticism are relevant here.

Focus on the first criticism. We can indeed use (232) to communicate (237), as Schaffer suggests.

(237) Every boy has a father who *that boy thinks* might be a genius.

However, we can also use (238) to communicate (239) and (240) to communicate (241).

(238) Every boy has a father who could be a star MLB pitcher.

(239) Every boy has a father who *that boy thinks* could be a star MLB pitcher.

(240) Every boy has a father who is the greatest person in the world.

(241) Every boy has a father who *that boy thinks* is the greatest person in the world.

In the case of *could* and *greatest*, however, it is clear we do not want to posit implicit variables ranging over information states to account for this fact. And the simplest hypothesis is that we should not do so for *might* either. The phenomena seem to be uniform, due to the *every boy has a father who* construction rather than to any of the specific terms occurring within its relative clause. How to account for these *epistemic intrusions* into the content of the claims made with (232), (238), and (240) need not concern us here.<sup>70</sup> Presumably, some sort of pragmatic process is involved, but the main point is that nothing special need be said about *might*. Apparent co-variation of individuals and information states provides no evidence in this case for *might* having an implicit information state variable.

The same holds for the *might*-sentence embedded in the conditional antecedent in (233). Conditionals are used to convey that some sort of connection holds between the antecedent and the consequent, one often implicitly mediated by knowledge. Thus what a conditional like (242)—with a *non*-epistemic modal in its antecedent—is used to communicate is likely to be (243).

(242) If you can shoot pigeons, you should do so.

(243) If you can shoot pigeons, *and you know this*, you should do so.

The same could hold for (244) and (245).

<sup>70</sup> The term *epistemic intrusion* is due to Alexander Williams.

- (244) If it might be raining, you should bring an umbrella.  
 (245) If it might be raining, *and you know this*, you should bring an umbrella.

A conditional construction could itself result in an epistemic intrusion into meaning. In this case, at least, we need not posit an implicit variable ranging over information states for *might* any more than we need to do so for *can*.

Turn now to the second criticism, and to examples like (234) and (235). Here the co-variation seems to be of circumstances, for circumstances incriminate individuals, not information states. With (234) we claim that sometimes the circumstances do not incriminate Billy, while with (235) we claim that when the circumstances incriminate Billy, they typically incriminate Suzy too. That said, while there is thus co-variation that could be attributed on my circumstantial view of *might* directly to the modal itself, it would be premature to think that the co-varying material had to be a part of content even when the *might* sentence was not embedded beneath a quantifier phrase. For, as Recanati (2004) has argued, we can view quantifier phrases as introducing not only the quantifier itself but also the variable that it binds (2004, 107–11). In that case, there would be no modal base variable in logical form when the *might*-sentence was unembedded, and so the claim made with such a sentence would not have the modal base as part of its content. And there are other possibilities too (for explanation of co-variation in terms of bound event variables, cf. Bach 2000; and for a general survey of the space of possibilities, cf. Elbourne 2008).

Indeed, co-variation could even simply be attributed to intensional operators that manipulated the values of the evaluation parameters. Stanley himself considers this possibility of semantic binding “without the mediation of variables” (2000, 413, fn. 24), but argues against it, at least in some cases, on the basis of the intuition that the relevant co-varying materials were “parts of the contents of propositions” rather than of “the entities relative to which the truth of propositions is evaluated” (2000, 413, fn. 24). Now this intuition is more or less orthogonal to Stanley’s central concern, which is the division of labor between semantic and pragmatic processes in generating the intuitive truth conditions of a claim. However, it is not at all orthogonal to Schaffer’s, which is the division of labor between content and the evaluation matrix. Hence, the possibility of appeal to intensional operators is a real threat to Schaffer’s argument.<sup>71</sup>

Are there cases of apparent co-variation that *do* require variables? Indeed, but here the facts tell against Schaffer. For, in the modal case, there is most definitely not the co-variation that we would expect if modals introduced modal base variables.

The relevant cases involve so-called *donkey anaphora*, as in (246).

---

<sup>71</sup> For co-variation in the case of modals, whether with *or* without intensional operators, we would want to adopt a situation semantics (Barwise 1981; Barwise and Perry 1983; Kratzer 2019). On the view being considered now, we could then treat *sometimes* and *usually* as intensional operators that manipulated the value of the situation parameter, taking as their input a situation-variant content and yielding as their output a situation invariant-one. The claim made with (234) would then be true if sometimes the relevant situations did not incriminate Billy. And the claim made with (235) would be true if, usually when Billy was incriminated in a relevant situation, so was Suzy. In this way, we could account for covariation without having to posit the covarying material itself as antecedently being part of content.

(246) No vegan who owns a donkey beats it.

These are cases of anaphora where neither term involved in the anaphoric dependency has scope with regard to the other (cf. Geach 1962), a fact represented in (247), with the brackets indicating scope and the subscripts anaphoric dependency.

(247) [No [vegan [who [owns a donkey<sub>i</sub>]]][beats it<sub>i</sub>]]

How to account for this sort of anaphoric dependency has been a matter of intense research and debate since the advent in the early eighties of Discourse Representation Theory (Kamp 1981) and File Change Semantics (Heim 1983). I will have nothing to say about this debate. (For recent summaries, with references, cf. Geurts 2012; King and Lewis 2018.) The main point of concern is that, in anaphoric dependency without relative scope, it seems impossible for co-variation to be attributed to a quantifier phrase introducing both a quantifier and a variable, à la Recanati, or to an intensional operator manipulating the value of an evaluation parameter, à la Stanley. For the anaphoric dependency is not a scope-based one, as these mechanisms require. Instead, it seems that the anaphoric dependency must be attributed to variable(-like) elements that are somehow antecedently associated with the second term involved.

Cases of so-called donkey anaphora are widely attested for any number of context-sensitive terms, suggesting they are associated with implicit variable(-like) elements. Relevant classes of terms include (i) relational adjectives like *local*, *foreign*, and *different*, as in (248)-(250); (ii) temporal and locative adverbs, as in (251)-(252); and (iii) verbs like *win*, *apply*, and *notice*, as in (253)-(255), that are interpreted on their intransitive uses as having definite but null (or unpronounced) complements (cf., a.o., Fillmore 1986; Partee 1989; Condoravdi and Gawron 1996; Williams 2012). And there are other classes of context-sensitive terms that fall into this category too.

(248) No man who visited his relatives for the holidays went to the local church.

(249) Everyone who enters an established community at first has foreign customs.

(250) Every beginning general who loses a battle switches to a different strategy in his second. [Partee's (22d)]

(251) Every man who stole a car abandoned it two hours later. [Partee's (13a)]

(252) Every man who stole a car abandoned it 50 miles away. [Partee's (13b)]

(253) Every man who put chips on 17 won. [Williams' (8)]

(254) Every boy who saw an ad for dish washers applied. [Williams' (12b)]

(255) Every man who shaves off his beard expects his wife to notice. [Partee's (10c)]

For all of the examples above, the most prominent readings are the ones that involve co-variation. For example, no man went to the church that was local *to his relatives*, not the church that was local to the speaker. Every man who stole a car abandoned it 50 miles away *from where he stole it*, not 50 miles away from where the speaker is at.

And every boy who saw an ad for dish washers applied *for that advertised job*, not for the job for lawn-mowers that the speaker had advertised.

Modals like *might*, however, are part of a class of context-sensitive terms that does not give rise to donkey anaphora. This is transparent if the relevant context-sensitivity is to information states. The claims made with (256) and (257) are clearly false, for example, even though they could be true in a community of self-deceived dog-owners if the possibilities co-varied with their information states (Alexander Williams, p.c.).

(256) No dog with a doting owner might be ill-behaved.

(257) Every dog with a doting owner might be the best behaved.

Similarly, (258) and (260) do not have the co-varying readings indicated in (259) and (261). Each reflects the judgment of the speaker, not the contestants.

(258) No entry submitted by an overconfident student is a possible finalist.  $\neq$

(259) No entry submitted by an overconfident student<sub>*i*</sub> is a possible<sub>*i*</sub> finalist.

(260) Every entry submitted by an underconfident student is plausible.  $\neq$

(261) Every entry submitted by an underconfident student<sub>*i*</sub> is plausible<sub>*i*</sub>.

Co-varying readings also seem to be unavailable if we assume that modals like *might* are sensitive to sets of circumstances, not information states, as I have argued that we should. For example, (262) does not have the obviously false reading that, *relative to their flu-like symptoms*, only some individuals might have the flu.

(262) Only some individuals with flu-like symptoms might have the flu.

And this lack of co-variation seems to hold for modals of all flavors, as can be seen from examples like (263) and (264).

(263) No person who violates an idiotic set of rules does what he should.

(264) No motivated athlete who's received a poor training performs as well as he can.

(263) does not have the likely tautologous reading that the rule-violators fail to do what they should *according to the set of rules they violate*. And (264) does not have the presumably false reading that motivated athletes fail to perform as well as they can *given the limitations of the training they have received*.

The generalization thus seems to be that modals comprise a class of context-sensitive terms that does not give rise to donkey anaphora, signifying that *might* is not an outlier in this regard. This in turn suggests that modals are not associated with implicit modal base variables, at least not of a familiar sort. As a result, the argument from binding is turned on its head. The modal base should be taken to be part of the evaluation matrix, not part of content, reinforcing the conclusion from Section 4.5.1 and the proposal from Chapter 3, Section 3.3.1.

#### 4.6 An objection to a non-epistemic analysis

I have argued that what is required to make sense of genuine and warranted disagreement about what might be the case is much the same as what is required to make sense of such disagreement about what can happen. Once we have adopted a circumstantial analysis of *might*, the former turns out to be no more puzzling or unexpected than the latter. And this tells in favor of such an analysis.

A key part of my argument in this chapter, however, has been that the so-called entrenchment data are not indicative of whether genuine disagreement has occurred. Instead, just as entrenchment in debates about what can happen suggests the exact opposite, so it could do the same in debates about what might be the case, provided that the seemingly parallel examples are interpreted in parallel ways. Yet it might be objected that this proviso cannot be met. For in (178), repeated here as (265), the speaker's last claim is a counterfactual one.

(265) Well look, I wasn't saying that he *would* win. I was just saying that he *could*. And he could have won.

Thus, if the examples are parallel, then the speaker's last claim in (176), repeated here as (266), should also be interpreted counterfactually.

(266) Well look, I wasn't saying that he *was* a spy. I was just saying that he *might* be. And he might have been one.

And a counterfactual interpretation of this claim, it might be objected, is implausible. For whatever a speaker does in uttering (266), she does *not* seem to be claiming that, if things had gone differently for the concierge, he might have become a spy. Her claim seems to be about her past evidence rather than about the concierge's once-open futures. And, if this is true, then perhaps the entrenchment data show exactly what they have been alleged to in the case of *might*, and so something like the cloudy contextualism of von Fintel and Gillies turns out to be required to make sense of its compatibility with the dismissal data.

As a first response to this objection, I would simply deny that there could not be *some* sort of counterfactual interpretation of (266). For example, just as (179), repeated here as (267), makes explicit the counterfactual interpretation of (265), and is likewise acceptable as a continuation of the dialogue in (174), so the same holds for (268) with regard to (266) and (173), *mutatis mutandis*.

(267) Ok, so Phelps can't win. But he could have won if he hadn't had plantar fasciitis. And I didn't know that he did.

(268) Ok, so that concierge can't be a spy. But he might have been one if the vetting process here weren't so rigorous. And I didn't know that it was.

This first response may seem insufficient, however, given examples like (269) in which the claim that the concierge might have been a spy also plays a role in action explanations (cf. von Fintel and Gillies 2008, 87–88).

- (269) Ok, so that concierge can't be a spy. But he might have been one. And because he might have been one, I hid out of sight when he came by.

The worry here is that the action explained seems to have been a sensible one, and yet it clearly would not be sensible for the speaker to have hidden because of a counterfactual possibility. The worry is legitimate but applies equally well to the action explanation in (270) (cf. also MacFarlane 2011, 169–73; 2014, 271–75).

- (270) Ok, so Phelps can't win a medal. But he could have won one. And because he could have won one, I placed a massive bet on his doing so.

What explains the speaker's actions in both cases is not the counterfactual possibility itself, but instead the speaker's belief that it was an actual one. How the speaker manages to successfully convey these explanations in uttering (269) or (270) is an interesting question,<sup>72</sup> but not one that need make us think that either *That concierge might have been a spy* or *Phelps could have won a medal* is used to make a past epistemic possibility claim rather than a counterfactual one.

Indeed, we can even turn the tables and move from defense to offense. For (271) is used to make a claim about a past epistemic state, and (272) provides a good explanation for why that state obtained.

- (271) It was plausible that that concierge was a spy.  
(272) But only because I didn't know at the time about the hotel's vetting process.

Yet (272) does not provide a good explanation for the truth of the claim made with (273), even though it would do so if that claim were about a past epistemic possibility. Instead, the truth of the claim made with (273) seems to be best explained in much the same way that the truth of the counterfactual claim made with (274) is.

- (273) That concierge might have been a spy.  
(274) Phelps could have won.

That is, the truth of these claims seem to be best explained by appeal to facts on the ground that, in the absence of acknowledged contravening facts, could have opened up the relevant possibilities. Such appeals are made, respectively, in (275) and (276).

- (275) But only because he moves so stealthily.  
(276) But only because he's so damn talented.

Thus, if I am correct, there is no obstacle to a counterfactual interpretation of both *might have* and *could have* in the entrenchment data, and so also no obstacle to interpreting the data in a parallel way either. Indeed, there may even be a requirement

---

<sup>72</sup> Stephenson (2007) proposes that *because* shifts the value of an information state parameter to the information state of the individual whose "conscious reasoning or rationale" explains the action undertaken (2007, 505–8). Yet this proposal would not work for action explanations involving modals like *can* that are not sensitive to information state parameters.

for both these things, thus making the case for a circumstantial analysis of *might* even stronger.<sup>73</sup>

#### 4.7 Conclusion

Legitimate disagreement about what might be the case seems possible, even between strangers. Yet on an epistemic analysis of *might* it is not easy to see how this could be so. For if we assume that the strangers are making claims about information states, either we have to hold that they are talking past each other, or that they are making claims for which they have no warrant. There is no evidence for the former, and indeed there is strong evidence to the contrary. It is also unintuitive to think that speakers' *might*-claims are typically unwarranted. So proponents of epistemic analyses have had to adopt fairly radical revisions to the semantics and pragmatics for *might*-claims so as to account for warrant. We can avoid these revisions—which are fraught with difficulties and may not solve the problem of warrant anyway—if we simply adopt a circumstantial analysis of *might*. On a circumstantial analysis, it is no more difficult to explain how people with different evidence could have legitimate disagreement about what might be the case than it is to explain how people with different evidence could have legitimate disagreement about what can happen. In both cases, the disagreement is about the possibilities and potentials left open by the circumstances. Speakers naturally do not have all the same information when it comes to the circumstances, just as they do not have all the same information when it comes to the laws. But just as such differences of information do not threaten the possibility of legitimate disagreement about what the laws leave open, so they do not threaten the possibility of legitimate disagreement about what the circumstances do. The best explanation for how there can be legitimate disagreement about what might be the case is that *might* expresses possibility relative to the circumstances, not to an information state.

---

<sup>73</sup> For more on the relevant counterfactual interpretation of *might have*, cf. Chapter 7, Section 7.6.3.

## Chapter Five: Applying the Analysis to Modal Belief

To believe something might be the case ... is not really to embrace any positive thesis about how the world is. Rather, it is a way of lacking information.

~ Yalcin (2011, 310)

### 5.1 A puzzle about belief reports

Modals embedded beneath belief verbs seem to behave differently depending on the flavor of modality that they express. Take the *deontic* claim made with (277), for example, and the report made with (278) when (277) is embedded beneath a belief verb.

(277) The president can declare war on his own.

(278) Mary {believes/thinks} that the president can declare war on his own.

For the claim made with (277) to be true, there needs to be a certain deontic possibility: the president's declaring war on his own needs to be compatible with the relevant norms. Similarly, for the belief report made with (278) to be true, Mary needs to have a belief about a deontic possibility. She needs to believe that certain relevant norms are compatible with the president's declaring war on his own. Now contrast this with apparently epistemic uses of modals. For the claim made with (279) to be true, it is alleged, there needs to be an epistemic possibility.

(279) There might be landmines.

However, for the belief report made with (280) to be true, Mary does not need to have a belief about such a possibility (Stephenson 2007).

(280) Mary {believes/thinks} that there might be landmines.

Mary does not need to believe that anyone's information, whether her own or someone else's, is compatible with there being landmines (cf. Section 5.2 below). This is puzzling.

The most promising epistemic response to this puzzle has been to hold that embedded "epistemics" are anaphoric on, and so interpreted directly against, the information state determined by the attitude verb itself (Yalcin 2007; Hacquard 2010). Just as the claim made with (279) is supposed to be true if and only if there being landmines is compatible with the information determined by the discourse context, so the claim made with (280) is supposed to be true if and only if there being landmines is compatible with the information determined by (and corresponding to) Mary's beliefs. The attitude verb thus plays the same role for the embedded modal in (280) that the discourse context does for the unembedded modal in (279). As a result of this unique role, embedded "epistemics" turn out to be unlike other embedded modals: they modify the force of an attitude rather than contributing to its content. (280) indicates not that Mary believes some modal content to be true but instead that

she has some less-than-full-belief level of credence that the content of the modal's complement, or its *prejacent*, is true. The attitude reported with (280) is thus something like suspicion, or *weak belief*, that there are landmines.

In this chapter I argue against this anaphoric response to the apparent puzzle about belief reports and in favor of a response that dissolves the puzzle by adopting a circumstantial analysis of *might*. Given such an analysis, what we expect to report with (280) is not a belief that someone's information leaves open the possibility that there are landmines, but instead a belief that circumstances do. And it is not at all implausible to think that this is exactly what we do report. Modals on my analysis thus behave in the uniform way we expect them to: when embedded beneath belief verbs, they are used to report beliefs about the same sort of modalities that they are used to make claims about when unembedded.

This circumstantial response to the puzzle about belief reports is preferable, I argue, for both theoretical and empirical reasons. From a theoretical standpoint, it is preferable insofar as it does not require embedded "epistemics" to be interpreted idiosyncratically (and so also does not incur the incidental expenses that come in trying to make an idiosyncratic interpretation feasible). And from an empirical standpoint, it is preferable insofar as it readily makes sense of cases in which the report made with a sentence like (280) can be true even when the prejacent is *not* compatible with the content of the belief-holder's beliefs (cf. Section 5.5.2 below). In contrast, anaphoric analyses must make exceptions that ultimately undermine their explanatory force in order to make sense of such cases.

Here is the structure my argument will take. In Section 5.2, I show that what is reported with (280) need not be a belief about information, or what I will call a *second-order sort of belief*. This is in contrast to what most epistemic analyses predict, as I show in Section 5.3. In Section 5.4, I present the epistemic analysis that avoids this prediction, the *anaphoric analysis*. In Section 5.5, I then detail the theoretical and empirical challenges facing this analysis. In Section 5.6, I show how a non-epistemic analysis can be applied to belief reports while avoiding such problems. I conclude in Section 5.7 by responding to an objection that a non-epistemic analysis cannot account for data that an anaphoric analysis can (cf. Yalcin 2007; 2011).

## 5.2 Belief reports without second-order commitments

The belief reported with (280) need not be a second-order sort of belief. It need not be about information.<sup>74</sup> First and foremost, the reported belief need not be about the information relevant in the speech situation, or *global context* (Egan et al. 2005; Stephenson 2007). Say that the cadets in officer training camp have been tasked independently with determining whether there are fake landmines planted in the practice fields. One cadet, reporting back, has the following conversation with the drill sergeant.

---

<sup>74</sup> What matters, of course, is information *qua* information. For on a *de re* reading of *information*, beliefs about information are just beliefs about propositions, and all beliefs are trivially beliefs about information. This reading cannot be the intended one for epistemic theorists, however, else modals like *can* would just as well be analyzed in informational terms too, which they are not.

- (281) a. [*Cadet*] There might be landmines.  
b. [*Sgt.*] Interesting. Mary {believes/thinks} there might be landmines too.

In this speech context, it is clear that the cadet's information state is the relevant one. (The sergeant presumably knows where the fake landmines are.) Yet Mary, another cadet, need not have any beliefs about this other cadet's information for the report made with ((281)b) to be true. Indeed, Mary need not have any beliefs about this other cadet at all. Hence, ((281)b) does not have to be used to report a belief about the information relevant in the global context.

Sometimes it is claimed that *might* is sensitive, when embedded, to an information state provided by the *local context*: in this case, the information state of the subject of the belief verb (Egan et al. 2005; Stephenson 2007). Perhaps so. But, as observed in Yalcin (2007) and Hacquard (2010), ((281)b) need not be used to report a reflexive belief about the belief-holder's information either. Say that Mary is prone to jump to quick conclusions, easily misled by planted evidence, and seemingly constitutionally incapable of following up on a lead, unless prompted by others. It would be entirely possible and perfectly coherent for Mary to sincerely utter ((281)a) and then to subsequently disavow having reflected on her own information.<sup>75</sup> On the basis of her sincere utterance of ((281)a), we would normally take the report made with ((281)b) to be true. Yet that report could not be true in this scenario were ((281)b) used to report a second-order belief, for Mary would not have such a belief. Insofar as that report *does* seem true, even in this scenario, then ((281)b) need not be used to report a second-order belief.

Here is a further piece of evidence. With the report made with (282), we incur commitments about the second-order capacities of non-human animals.

- (282) Fido {believes/thinks} that his information is compatible with there being a bone under the table.

In contrast, as Yalcin (2007; 2011) notes, we do not incur such commitments with the report made with (283).

- (283) Fido {believes/thinks} there might be a bone under the table.

So clearly, with belief report sentences with *might* in the complement, we are not required to report second-order beliefs.

### 5.3 Epistemic analyses with second-order commitments

#### 5.3.1 Contextualists and relativists

The predictions of most epistemic analyses of *might* contrast with the observations of the previous section. This is obvious for so-called contextualist theorists. Standardly, contextualists have taken the content of a bare *might*-sentence, or BMS, to include a

---

<sup>75</sup> Alexander Williams, p.c. Mary may be particularly unreflective, but she need not have misunderstood the meaning of *might*, or have used it inappropriately. She would instead be in the position many of us find ourselves in from time to time: namely, that of making a claim without having thought about whether our information supports it.

contextually-determined information state, regardless of whether that BMS occurs in a matrix clause or in an embedded one (cf. DeRose 1991; Papafragou 2006; Dowell 2011; Schaffer 2011; Crabill 2013; Yanovich 2014).<sup>76</sup> Contextualists thereby incur second-order commitments for belief report sentences embedding BMSs, unless more is said. Dowell makes this commitment explicit:

[W]hich proposition the complement *might*  $\Phi$  has as its content [in belief reports] is often determined at least in part by the [reported belief-holder's] information .... [I]n the default case, one's own modal beliefs are *beliefs about what's compatible with what one or one's group knows* (2011, 19; emphasis mine).

Recently, more has been said. In particular, a non-standard contextualist account of *might* has been developed to avoid second-order commitments (Silk 2017). It does this by taking truth-conditional reference to information states to help *determine* content without itself being a part of content (Silk 2017, 1783). On Silk's account, you will recall, reference to information states for epistemic modals, like reference to speakers for first-person pronouns, plays the role of a Kaplanian *character*: that is, a function from speech contexts to speech contents (Kaplan 1989). In the modal case, the content determined by the character is a set of propositions: the propositions that some particular information state comprises. What we report with a sentence like (284), on this account, is not a belief about information at all but instead "the belief that a certain set of propositions is compatible with the proposition *b* that the butler is the killer" (Silk 2017, 1783).

(284) Alice thinks that the butler might be the killer. [Silk's (11)]

Silk's account thus avoids the second-order commitments that standard contextualist accounts incur. However, there is also good reason for his account being non-standard. As Kratzer (1991) makes explicit, and as is already implicit in Kratzer (1977; 1981), logical relations between sets of propositions are not contingent while the contents of our modal beliefs and claims are (1991, 641–42). The conjunction of (284) and (285) describes a perfectly coherent set of beliefs for Alice to have, but it would not if Alice's reported belief in (284) was about a non-contingent relation.

(285) But she also thinks that, if only the butler had learned to control his temper, there wouldn't even have to be this possibility.

---

<sup>76</sup> Recall that a BMS is a sentence, like (279), in which *might* occurs without being restricted by an epistemic modifier (e.g., a phrase like *given what the cadets know*). Some theorists hold it is not the content of a BMS, but instead of a claim made with a BMS, that includes a contextually-determined information state (cf. Bach 2011; Braun 2012; these authors also hold that the information state is determined in, but not by, context). Other theorists hold that discourse contexts need not determine unique evaluation contexts, and so hold that BMSs need not have unique contents in discourse contexts either (cf. von Fintel and Gillies 2011). These differences are unimportant for our purposes here. What matters is only that these theorists take the relevant implicit parts of contents to be information states.

To secure the contingency of modal beliefs and claims, we need the descriptions that contingently characterize sets of propositions—descriptions like *the norms*, *the facts*, *our information*—to not play the role of a Kaplanian character but instead be part of content. Yet if we take descriptions like *our information* to be part of content, we incur second-order commitments. Epistemic contextualists seem to have no good options.<sup>77</sup>

A common alternative to contextualism is relativism. Relativists, you will recall, move information states from a contextually given aspect of the proposition expressed by a BMS to an aspect of the evaluation matrix relative to which that proposition is true or false (Egan et al. 2005; MacFarlane 2011; 2014). In moving information states from content, relativists correctly predict that BMSs embedded in belief reports will not be globally context-sensitive. (The relativist content of the reported belief does not need an information state to be supplied by context, after all.) That said, they still incur second-order commitments, as some relativists make explicit:

*a* believes the [relativist] proposition *b might be F* ... iff *a* believes it is consistent with what they know that *b* is *F* (Egan et al. 2005, 158).

Other relativists are less explicit but hold that the content of a factive attitude report can be pragmatically enriched, when necessary to avoid presupposition failure, to include the attitude-holder's information state, all without changing the sort of attitude reported (cf. MacFarlane 2014, 276–77; MacFarlane is responding to von Fintel and Gillies 2008). So relativists in general seem committed to the position that adopting an attitude toward the purportedly relativist content of a BMS is equivalent to having a second-order attitude.

### 5.3.2 Expressivists

In contrast to truth-conditional theories like contextualism and relativism, recall that expressivist theories do not take *might* to make the content of a speech act depend truth-conditionally on an information state (cf. Halliday 1970; Palmer 2001; Schnieder 2010). Instead, they take its function to be to modify the force of that act. With *might*, the prejacent is put forth with weak, non-assertoric commitment. This enables a speaker to express, without having to assert, a less-than-full-belief level of credence in the prejacent.

Expressivists correctly predict that belief report sentences embedding BMSs need not be used to report beliefs about information states. For they do not take *might* to contribute any such state to content to begin with. However, it is not clear that they

---

<sup>77</sup> Some contextualists appeal to descriptions that can sometimes make it unclear whether they are *epistemic* contextualists. Yanovich (2014), for example, talks of believing, for the purposes of an investigation, that something might be so. Elsewhere, Yanovich makes clear, however, that to think that *p* might be so, “for the purposes of an investigation”, is to think that *p* is consistent with the facts relevant to the investigation that are either known or knowable (cf. Yanovich 2014, 76–77). And to have a belief about what is knowable is, undoubtedly, to have a second-order sort of belief.

can make sense of the attitude that a sentence like (280), repeated here as (286), *is* used to report, nor of how this differs from the attitude that a sentence like (287) is.<sup>78</sup>

(286) Mary {believes/thinks} that there might be landmines.

(287) Mary {believes/thinks} that there are landmines.

One possibility is to hold that the embedded *might* modifies the force of the attitude, so that (9) is used to report a less-than-full belief while (287) is used to report a full one. In response, I would argue that the ability of a purely expressive term to modify attitudinal force seems, understandably, to be parasitic on its ability to modify speech act force. Take *gee* in (288), for example.

(288) Mary {believes/thinks} that, gee, there are landmines.

With the embedded *gee*, the speaker indicates an additional attitude of surprise on the part of Mary that she does not with (287). However, (288) also seems to involve some sort of coerced direct quotation. Apart from such coerced quotation, *gee* is not able to modify attitudinal force, as (289) shows.

(289) a. Tell me. Do you sincerely and fully believe that, gee, there are landmines?

b. You know, I've thought long and hard about the matter. And yes, I do sincerely and fully believe that, gee, there are landmines.

In (289), the first *gee* involves a sort of projected speaker comment, and the second a sort of parroting. Neither changes the attitude that is the object of inquiry. The second speaker could have left out *gee* and still answered the first. This is what we should expect of purely expressive terms. To the extent that they can be used to modify attitudinal force, it is only because the attitude being reported had been expressed in a speech act, the force of which they had also been used to modify.

*Might*'s ability to modify attitudes, however, does not seem to be parasitic in this way. In particular, (286), unlike (288), does not seem to involve coerced quotation. And *might* in (290), unlike *gee* in (289), does change the nature of the attitude that is the object of inquiry.

(290) a. Tell me. Do you sincerely and fully believe that there might be landmines?

b. You know, I've thought long and hard about the matter. And yes, I do sincerely and fully believe that there might be landmines.

The second speaker in (290) could not have replaced *might be* with *are* and still have directly answered the first.

---

<sup>78</sup> For arguments from embedding against expressivist views, compare, among others, Papafragou (2006), Swanson (2011), Hacquard and Wellwood (2012), and Anand and Hacquard (2009; 2013).

*Might*, in summary, does not seem to act as an expressivist term when embedded.<sup>79</sup> To account for its effect in embedded contexts, we thus need *might* not to modify force but instead to help determine truth conditions.

#### 5.4 Avoiding second-order commitments via anaphora

For epistemic contextualists and relativists, the relevant information state for a BMS embedded in a belief report sentence is determined by (and corresponds to) the beliefs or knowledge that the belief-holder *believes* she (or someone else) possesses. Some authors have recently proposed that we can avoid second-order commitments, while maintaining an epistemic analysis, if we instead take the relevant information state to be determined by (and correspond to) the belief-holder's *actual* beliefs (Yalcin 2007; 2011; Hacquard 2010; Anand and Hacquard 2009; 2013).

To accomplish this, these authors take advantage of a standard formalization of both belief verbs and modals as quantifiers over possible worlds (cf. Hintikka 1962; Kratzer 1977). On a Hintikkan analysis of belief verbs, the claim made with (291) is true iff there are landmines in every world compatible with the content of Mary's beliefs. And on a Kratzerian analysis of modals, the claim made with (292) is true iff there are landmines in some world compatible the content of the contextually relevant information state.

- (291) Mary believes that there are landmines.  
 (292) There might be landmines.

What authors like Yalcin and Hacquard do, given this common formalization, is to make an embedded epistemic modal anaphoric on the embedding belief verb, so that the set of worlds each quantifies over is the same, as in (293). (Let  $\mathcal{B}$  stand for worlds compatible with the content of a belief-holder's beliefs.)

- (293)  $\forall w \in \mathcal{B}: \exists w' \in \mathcal{B}$  such that there are landmines in  $w'$ .

For Yalcin, this anaphoric dependency is accomplished by having the belief verb reset the value of a special parameter that *might* is supposed to be relativized to (Yalcin 2007). For Hacquard, it is accomplished by having the verb introduce an event variable whose value provides, via a binding relation, the value of an event variable that modals now also introduce (Hacquard 2010).<sup>80</sup> Either way, the result is the same. The modal quantifies over worlds compatible with the content of the belief-holder's actual beliefs, the worlds the belief verb also quantifies over. The result is to make the belief verb's quantification vacuous, and so essentially to modify the force of the attitude, for (293) is just equivalent to (294).

- (294)  $\exists w' \in \mathcal{B}$  such that there are there are landmines in  $w'$ .

<sup>79</sup> To foreshadow somewhat, it also does not seem to modify attitudinal force. There is nothing odd with (290), despite adverbs like *fully* modifying force in ways incompatible with *might*'s alleged modification.

<sup>80</sup> Hacquard, it should be noted, has independent reasons for having modals introduce event variables and takes all modals to do so, not just modals that are allegedly interpreted epistemically.

At the end of the day, for these authors, the claim we are thus supposed to be making with a belief report sentence embedding a BMS is that the prejacent is compatible with the content of the belief-holder's beliefs. Alternatively, we are supposed to be claiming that, relative to the information state determined by the content of those beliefs, the prejacent is a possibility. And this would seem to be a way of saying that the prejacent is an epistemic possibility for the belief-holder.

It is worth noting the essential role that anaphoric dependency plays on this analysis in allowing second-order commitments to be avoided (Yalcin 2007, 997; Hacquard 2010, 105). Stephenson (2007) had proposed that an embedding belief verb and an embedded epistemic modal could also quantify over the same set of worlds, without anaphora, if we assumed that believing that  $p$  required believing that one knew that  $p$  (2007, 502–3). On that assumption, the worlds compatible with the content of one's beliefs would be the same as the worlds compatible with the content of one's presumed knowledge, that is, of the content that one believed one knew. Stephenson's interpretation of (286), namely, (295), would thus entail (296), Yalcin's and Hacquard's.

- (295) In all worlds  $w'$  compatible with the content of Mary's beliefs in  $w$ , it is the case that, in some world  $w''$  compatible with the content of Mary's knowledge in  $w'$ , there are landmines in  $w''$ .
- (296) In all worlds  $w'$  compatible with the content of Mary's beliefs in  $w$ , it is the case that, in some world  $w''$  compatible with the content of Mary's beliefs in  $w$ , there are landmines in  $w''$ .

As a result of the alleged entailment, Stephenson, like Yalcin and Hacquard, would take the report made with a belief sentence embedding a BMS to be true iff the prejacent was compatible with the content of the reported belief-holder's beliefs. However, what would be reported would also be a second-order sort of belief. The representation of Mary's belief in (295) requires her to have a belief about information, as would any representation referencing her or anyone else's information in the worlds  $w'$  compatible with her beliefs. This is similar to the way that the belief represented in (297) would be a belief about laws, even if what Mary believed to be true was somehow coextensive with what she believed to be required.

- (297) In all worlds  $w'$  compatible with the content of Mary's beliefs in  $w$ , it is the case that, in some world  $w''$  compatible with the content of the laws in  $w'$ , there are landmines in  $w''$ .

To avoid second-order commitments, we need to avoid reference to any information state other than the one determined by the attitude verb. Thus, if we assume an embedded modal is interpreted relative to an information state, we require anaphoric dependency of the modal on the attitude verb itself.

## 5.5 Challenges for an anaphoric analysis

### 5.5.1 Theoretical challenges

An anaphoric analysis makes epistemic interpretations of modal predicates idiosyncratic, unlike other interpretations of modals beneath belief verbs. As Yalcin writes of epistemic possibility modals, the beliefs they are used to report “do not correspond to a distinctive class of believed contents; rather, they correspond to a distinctive way of being doxastically related to a proposition” (2011, 309). On their non-epistemic interpretations, in contrast, modal predicates would contribute to the content of a belief. Indeed, even some epistemic modal predicates seem to do this: for example, terms like *plausible* and *plausibly*. While we do not incur any commitments about the second-order capacities of non-human animals with (298), on its allegedly epistemic interpretation, we do with (299).

(298) Fido thinks there’s possibly a pill mixed in with his food.

(299) Fido thinks there’s plausibly a pill mixed in with his food.

Only some modal predicates interpreted epistemically, then, would fail to contribute to the content of a reported attitude and instead modify its force. This makes the apparently non-uniform behavior of embedded modal predicates more difficult to explain. It makes this behavior less explicable on independent conceptual grounds, as opposed to arbitrary lexical ones, than it otherwise would have been.

Speaking of arbitrary lexical properties, belief verbs will also have to display them vis-à-vis other attitude and speech act verbs on anaphoric analyses. For, in some cases, a BMS occurring as a complement of an attitude or speech act verb seems to contribute to the content of the reported attitude or speech act, rather than to modify its force (Crabill 2013; Yanovich 2014). Take (300), for example.

(300) John’s doctors have {reported/showed} that he might be contagious.

On a natural interpretation of (300), what John’s doctors have done is to report, or to show, the world to be a certain way. That is, they have reported or shown a certain content to be true. Yet if this is correct, anaphoric analyses must explain why belief report sentences embedding BMSs are not used to report beliefs with these modal contents. (I presume that whatever can be the content of a report or demonstration can also be the content of a belief.) And here, it seems, the explanation could only be in terms of an arbitrary lexical property of attitude verbs like *believe*. We would thus have more or less the equivalent of a perfect storm. Attitude reports embedding BMSs could be used to report attitudes about modal contents. And belief reports embedding other epistemic modal sentences could be used to report beliefs about modal contents too. However, belief reports embedding BMSs would not be used to report beliefs about modal contents, presumably due to the combined arbitrary lexical properties of both *believe* and *might*.

One possibility for anaphoric analyses is to offer an alternative interpretation of (300), and of other attitude and speech act reports for which we would naturally

appeal to modal contents. Yanovich, however, argues that this would not work in the case of (300).<sup>81</sup> And it is not clear that it would work in other cases either.

Take indirect speech reports, for example. On an anaphoric analysis, an epistemic possibility claim functions to express the compatibility of the speaker's information state with the prejacent, and to coordinate on that same state of mind with one's interlocutors (cf. Yalcin 2007; 2011). We could thus interpret an indirect speech report like (301) in light of either of these functions, as in (302) and (303) (cf. Anand and Hacquard 2009).

- (301) John's doctors said that he might be contagious.
- (302) John's doctors expressed the compatibility of their information state with John's being contagious.
- (303) John's doctors aimed to coordinate the compatibility of their and their interlocutors' information states with John's being contagious.

Yet these interpretations yield the wrong results in a number of cases: for example, in cases in which possibility claims serve an evasive function, as in (304).

- (304) a. [S1] Do you know whether John is contagious?
- b. [S2] I do.
- c. [S1] Is he?
- d. [S2] He might be.

In such cases of evident evasion, the report made with (301) would still be true—to see this, witness the coherence of (305)—but the reports made with (302) and (303) would not.

- (305) John's doctors just said that he might be contagious. But they clearly know whether he is and just don't want to tell me.

The function of an epistemic possibility claim is sometimes formulated in terms of the common ground, however, instead of in terms of information states. A speaker, for example, is alleged to be “mak[ing] explicit that [the prejacent is] compatible with the common ground” (Yalcin 2007, 1010). On this formulation, it might seem that anaphoric analyses could predict the report made with (301) to be true, even in cases of evasion. In such cases, after all, the speaker is making explicit what is *not* in the common ground. Yet even this formulation will not suffice. For in some cases of evasion, the prejacent (or its negation) may be in the common ground, the speaker's

---

<sup>81</sup> One possible anaphoric interpretation of (300) is (iii):

- (iii) The content of the doctors' {report/demonstration} is compatible with the proposition that John is contagious.

Yet if the doctors report simply that John *is* contagious, then the conditions in (iii) would be met for the report made with (300) to be true. However, the report made with (300) would be false in this case. More sophisticated interpretations Yanovich shows to be equally problematic (cf. Yanovich 2014, 103–6).

purpose being not to withhold information but instead, as in (306), to avoid public commitment.

- (306) While it's common knowledge that John is contagious, his doctors have steadfastly refused to say that he is. Instead, they say only that he might be.

Hence, for indirect speech reports and reports of demonstrations, at least, it seems that an anaphoric analysis will not do. Instead, we require *might* to contribute to the content of the reported attitude or speech act in these cases.

Indeed, even for belief reports the bare form of the anaphoric analysis will not do (cf. Yalcin 2011). This form does not require the subject of the belief verb to have entertained the prejacent: it only requires the prejacent be compatible with the content of her beliefs. But many propositions are thus compatible without ever having been entertained. Call such propositions *bare epistemic possibilities*, in contrast to *live* ones, propositions that have been entertained. Bare epistemic possibilities can explain actions, as in (307).

- (307) [*Context: John is at the office party, despite being contagious. Colleagues avoid him, except for Mary, who is ignorant of his medical history.*]
- a. Why is Mary talking to John?
  - b. She doesn't know that he's contagious.

In (307), the bare epistemic possibility that John is not contagious is what explains Mary's actions. Without emphasis on *know*, Mary is not intentionally taking even a minor risk in talking to John. Contrast this with (308), where Mary *is* intentionally taking a health risk, and what explains her action is the live epistemic possibility that John is not contagious.

- (308) a. Why is Mary talking to John?  
b. She thinks that he might not be contagious.

((308)b) cannot be used to explain Mary's actions in terms of a bare epistemic possibility. Mary must have entertained the prejacent in order for the claim made with ((308)b) to be true. And to account for the fact that belief report sentences embedding BMSs cannot be used to report bare epistemic possibilities, Yalcin (2011) refines the anaphoric analysis, claiming that “[t]o count as believing  $\diamond\phi$ ,  $\phi$  should be compatible with one's beliefs; but in addition, one's state of belief should also be sensitive to a question for which  $\phi$  is an answer” (2011, 315–16).

This requirement for question sensitivity makes sense of what would otherwise be an anomalous fact on the anaphoric analysis: namely, that we are often hesitant to make the report we do with a sentence like ((308)b) apart from its subject having uttered (309).

- (309) John might not be contagious.

Our best evidence that Mary is sensitive to the question whether John is contagious, after all, will often just be that she has made such an utterance. However, a question-sensitivity requirement also complicates what appears to be an otherwise standard relation between utterances and belief reports. Normally we take the former to justify the latter in virtue of a sincere utterance requiring the speaker to believe the asserted content of that utterance. With the introduction of question sensitivity, however, we are forced to develop a parallel account of this relation for the case of BMSs.<sup>82</sup>

One final challenge that arises from denying distinctive modal contents of belief for belief reports embedding BMSs is to explain relations of anaphora, reference, and inference without them (on inference, cf. Braun 2013; MacFarlane 2014, 277–79). Take (310), for example, in which belief contents would play all three roles.

- (310) a. Sue believes that John might be contagious. Mary believes that too.  
 b. So the two believe the same thing.

It is certainly possible to account for the validity of the inference and the success of the anaphoric reference without contents of belief. (We could simply appeal to states of mind in their place.) However, to account for these things in a compositional framework will be a challenge.

One might have thought we could give an account in terms of bound variables ranging over information states, similar to the account we would give to explain the anaphora and inference in (311) when Sue and Mary have different fathers and so believe different propositions.<sup>83</sup>

- (311) a. Sue believes that her father loves her. Mary believes that too.  
 b. So the two believe the same thing.

Yet given an appeal to bound variables, we should also be able to report Sue and Mary as not believing the same thing, or not being in the same state of mind, given the differences in those variables' values. We do this coherently with (312) as a continuation of (311). Yet (313), in contrast, is not a coherent continuation to (310).

- (312) Yet they also do not believe the same thing, for the two have different fathers.  
 (313) Yet they also do not believe the same thing, for the two have different information.

---

<sup>82</sup> We must also explain why not all evidence of question sensitivity *cum* compatibility justifies a belief report. Mary's saying that she does not know whether John is contagious, for example, provides just such evidence. And yet Mary could in some cases reasonably object to that utterance being the basis for the belief report made with ((308)b) (cf. Yanovich (2014, 96–97) for a similar example used to make a similar point). To respond to this sort of problem, Yalcin defines question-sensitivity as "being *appropriately sensitive* to information which speaks to a question" and holds that what counts as appropriate sensitivity may vary with the context (2011, 317; emphasis mine). Mary's objection would thus reflect that she did not count as being appropriately sensitive in that context. This works, I guess, but further complicates a seemingly standard relation between utterances and belief reports

<sup>83</sup> Alexander Williams, p.c.

If we cannot appeal to bound variables, there seem to be two options. Either we take the noun-phrase complements of *believe* to be elliptical for the sentential complement (cf. Elbourne 2013). Or, instead, we take their semantic values in context to be some sort of entity that can combine with the semantic value of *believe* to determine the required state of mind. Yet this entity could not be a propositional content, for the prejacent certainly would not suffice and, by hypothesis, there is no appropriate modal content. Hence, we would have to complicate the semantic entry for *believe* or posit an ambiguity. None of these options seem promising. A simpler option is to appeal to distinctive modal contents of belief.

To conclude: an anaphoric analysis faces several theoretical challenges. It requires *might* not to behave like other epistemic modals when embedded beneath belief verbs, and belief verbs not to behave like other attitude verbs when embedding *might*. In doing away with distinctive contents for modal beliefs, it also complicates an otherwise standard relation between utterances and belief reports, and struggles to make sense of anaphora and inference. None of these challenges may be decisive. But they are serious enough, in my opinion, that the anaphoric analysis should only be adopted as a measure of last resort. Only if it provides greater empirical coverage than all the other alternatives should it be adopted, that is. And yet, as we will see in the next section, its empirical coverage is also limited.

### 5.5.2 Empirical challenges

From an empirical standpoint, anaphoric analyses are obviously an improvement on other epistemic analyses. In most seemingly true utterances of belief reports embedding BMSs, the prejacent will be compatible with the content of the subject of the belief verb's beliefs, even though in many cases she will not have second-order beliefs about its being so. Still, even anaphoric analyses are not entirely empirically adequate. For, in some cases, the relevant reports can be true even when the prejacent is *not* compatible with the content of the belief-holder's beliefs.

The first sort of case trades on the fact that not all of our beliefs have the same status. Some we may take to be foundational, and be unwilling to give up even if convinced that there is strong evidence that they could be false. Others, less foundational, we would give up as soon as we were so convinced. Presented with evidence of evil, for example, I may continue to believe that God is good. Presented with evidence that I am hallucinating, in contrast, I am unlikely to continue to believe that it is raining.

Say I do believe that God is good, then, but take myself not to have completely firm grounding for this belief. It seems I could accurately be reported as believing that God is good, while also recognizing that he might not be so, or even, simply, as believing that God might not be good.<sup>84</sup> Yet the proposition that God is not good is

---

<sup>84</sup> This especially so if I am attending to the fact that the grounding for my belief is not completely firm. Indeed, to claim to believe something, as opposed to claiming to know it, is often just to leave open the possibility that one could be wrong (and so to admit that the negation of what is believed is an epistemic possibility for oneself). This phenomenon is actually not infrequent. After game three of the 2017 NBA Finals, for example, I say to my friends: "I think the Cavs are going to lose this series. Still, they might win" (Drew Harr, p.c.). My second assertion is sincere, and so a reflection of what I believe. I thus believe that the Cavs are going to lose, while also believing that they might win.

*not* compatible with the content of my belief that God *is* good, and so is not compatible with the content of my belief state as a whole. These reports should thus be false on an anaphoric analysis, when in fact they seem to be true.

This first sort of case highlights that the subject matter of our beliefs matters. For more mundane matters, believing some proposition may—as a matter of practical rationality—rule out taking its negation to be a live possibility. Yet not all matters are mundane, and so believing a proposition may not always have this effect. Anaphoric analyses do not account for this fact. Instead, they encode into the semantics what seems to be a limited matter of practical rationality.

In contrast to the first case, which involves a sort of uncertainty on the part of the belief-holder, the second sort of case involves no uncertainty at all. These are cases of what have been called *exocentric* belief reports (cf. Lasersohn 2005; Stephenson 2007; Dowell 2011; Braun 2013; Yanovich 2014; Silk 2017). In “exocentric” reports, admittedly less common, it seems to be the information of someone other than the belief-holder that is relevant. (Reports in which the belief-holder’s information seems relevant have been called *autocentric* reports.) Here is an example of such a report.

(314) [*Context: Officer training camp. Fake landmines have been planted in certain fields. The sergeant, who knows which fields these are, is teaching her cadets to identify tell-tale signs of possible mining activity.*]

- a. [*Sgt.*] Are you sure you want to charge? There might be landmines.
- b. [*Cadet 1*] Do you really think so?
- c. [*Sgt.*] Sure I do. Look at the terrain of that field. [Etc.]

[*A second cadet approaches*]

- d. [*Cadet 2*] What’s going on? Why aren’t we charging?
- e. [*Cadet 1*] I’d ignored the possibility till now, but Sergeant thinks there might be landmines.
- f. [*Cadet 2*] Hmm, I think she’s right. We’d better proceed with caution.

In the example above, the belief report seems fine even though it is clear to the cadets that the sergeant has not forgotten in which fields the fake landmines have been planted. It is thus clear that she knows whether the prejacent is true and so, very possibly, that it is not. But if the sergeant knows that the prejacent is *not* true, then the prejacent is not compatible with the content of her beliefs. Hence, on anaphoric analyses, the truth of the cadet’s report should be contingent on the truth of the prejacent, which it does not seem to be.

Can anaphoric analyses account for these two sorts of cases? Indeed, they can. But the cure may be worse than what it cures. To account for a similar case involving a different attitude verb, for example, Yalcin appeals to a “tacit shift in the information parameter under the scope of [the verb]” (2007, 1013). The parameter shifts from the information state determined by (and corresponding to) the attitude verb itself to a salient information state about which the subject could have an attitude of that sort. That same appeal could be made for our cases. Its effect would be to make what is

reported a second-order belief. I do not think an appeal to tacit shifts should be made, however, here or elsewhere. Let me explain.

The first difficulty with the appeal to a tacit shift is its commitment to the beliefs and other attitudes being reported being second-order ones. Here are three cases that suggest otherwise.

*Case One:* Belief verbs. It is common knowledge among the sergeant and her cadets that the cadets have obtained a copy of her legend indicating which fields have fake landmines. It is also common knowledge that the field about to be charged does not. It is thus common knowledge that the sergeant does not believe that the beliefs of her cadets are compatible with there being landmines. Yet the belief report in (314) still seems true.

*Case Two:* Verbs of reporting and demonstrating. If we accept Yanovich's argument that there is no good anaphoric interpretation of (300), repeated here as (315), then a natural thought is to appeal to a tacit shift.

(315) John's doctors have {reported/showed} that he might be contagious.

With a tacit shift, John's doctors would be claimed to report or show a certain content to be true. And yet this content does not seem to involve an information state, as can be seen from the contrast between the alternative continuations in (316).

[*Context: John's doctors have been trying to convince the hospital director to lift his quarantine. John's partner reports back to his parents about their efforts.*]

(316) The director may lift the quarantine! John's doctors have reported that,  
*finally,*  
a. It's plausible that he's not contagious.  
a'. He might not be contagious.

With ((316)a), the facts about John's case do not need to have changed for the doctors' report to be true. For what the doctors report is a change in an information state. And while the latter sort of change may often be precipitated by the former, it is also possible without it. (The director, for example, may have a new assessment or new awareness of old facts.) With ((316)a'), however, the facts about John's case do need to have changed for the report to be true. This suggests that the doctors have reported a content involving not an information state, but instead involving the facts themselves.

*Case Three:* Verbs of uncertainty. In most true utterances of (317), Mary's beliefs (and so also her knowledge) will be compatible with John's being contagious (DeRose 1991). Yet anaphoric analyses make the opposite prediction, interpreting (317) as (318).

(317) Mary doesn't know whether John might be contagious.  
(318) Mary's knowledge state is *not* compatible with John's being contagious.

To avoid this result, Yalcin appeals to a tacit shift to a target information state: e.g., the state possessed by an epistemic authority (2007, 1013). Yet with this shift, we would incur commitments with (319) about the representational capacities of non-human animals that in fact we do not. Contrary to what the anaphoric analysis leads us to predict, (320) commits us to more than (319) does.

- (319) Fido doesn't know whether there might be an intruder downstairs.  
 (320) Fido doesn't know whether an intruder being downstairs is compatible with the target information state.

A second difficulty with an appeal to a tacit shift is that that shift would have to be only selectively available. To see this, consider an example used to motivate the anaphoric analysis (Yalcin 2007).

- (321) Imagine that John is contagious and that he might not be.

Yalcin claims that sentences like (321) are anomalous. The explanation of this anomaly, on an anaphoric analysis, is that we are making an incoherent request of our interlocutor. We are asking her to imagine that John is contagious while also making the content of her imaginings compatible with the proposition that he is not. Our request would not be incoherent, however, if a tacit shift were available. For we then would be requesting our interlocutor to imagine both that John is contagious and that, relative to someone or other's information, it is possible that he is not. We could make such a request with (322), for example, in the context of the office party above.

- (322) Imagine that John is contagious and that Mary doesn't know it.

To explain the anomaly of (321), it is thus crucial that a tacit shift *not* be available. Yet the principles governing the selective availability of this shift are not obvious. They cannot be principles of charitable interpretation, for the charitable interpretation of (321) would have us favor a coherent request over an incoherent one. Nor can they just be principles of salience and relevance. For, out of context, the most salient and relevant information states are those of the (intended) subjects of the attitude verbs. And yet we do not interpret (317) as (323), nor (321) as (324).

- (323) Mary doesn't know whether John's being contagious is compatible with *her* information state.  
 (324) Imagine that John is contagious and that *you* don't know it.

No immediately obvious principle governs the availability of the tacit shift, then. Yet some principle needs to be given, lest anaphoric analyses lose all explanatory power, making the shift selectively available for their apparent counterexamples but, crucially, not for their central explananda.

A third and final difficulty relates to the fundamentally different ways belief reports embedding BMSs would function with and without tacit shifts. In the former case, we would report a belief with a modal content. In the latter, more standard case,

we would report something like suspicion, or weak belief, that the modal's prejacent is true. Yet it would be simpler and more uniform to say that, in all cases, we reported beliefs with modal contents, and that such beliefs were often, though not always, accompanied by weak belief in the prejacent. It would be simpler and more uniform, that is, to take the belief-holder's weak belief to be a standard (but non-universal) implication drawn from belief reports embedding BMSs rather than to be the content of those reports in most (but not all) cases. The anaphoric analysis seems no more principled than if an expressivist who argued that we reported a belief-holder's pro-attitude toward an action with a report like (325) were then to claim, if presented with apparent counter-examples, that we could also report beliefs with moral contents whenever the required pro-attitude was lacking.

(325) John thinks that it is good to care even for the contagious.

The more principled response would seem to be that we always reported beliefs with moral contents with reports like (325), but that, in the vast majority of cases, those beliefs were also accompanied by certain related pro-attitudes.

To conclude: without a tacit shift, an anaphoric analysis incorrectly predicts the claim made with a belief report sentence embedding a BMS to be false whenever the prejacent is not compatible with the content of the belief-holder's beliefs. This tacit shift, optional in the case of belief verbs, would have to be obligatory for other attitude and speech act verbs: for example, for any for which *might* could only be analyzed as contributing to the content of the reported attitude or speech act. There would also have to be a third class of verbs, however, for which a tacit shift was unavailable: verbs like *suppose* and *imagine*. There is thus no general account to be given of attitude and speech act verbs, and there may be no principled one either. With a tacit shift, anaphoric analyses also predict us to be reporting second-order attitudes. Yet second-order attitudes are no more required with an alleged tacit shift than without one. Second-order attitudes *are* required for other epistemic modal predicates, however, regardless of whether the embedding attitude verb is one for which a tacit shift is supposed to be optional, obligatory, or unavailable.<sup>85</sup>

An anaphoric analysis misses the forest for the trees. It explains the fact that *might* does not seem to contribute an information state to the content of a reported belief by taking it not to contribute to the content of that belief at all. All the while, it ignores the fact that *might* does seem to contribute to the content of other attitudes and speech acts, despite not seeming to contribute an information state in those cases either. And it also ignores the fact that there are epistemic modals that do contribute information states in all these cases. It is time to look for a more general analysis of the fact that we need not report second-order beliefs with belief report sentences embedding BMSs.

---

<sup>85</sup> Valentine Hacquard (p.c.) has suggested to me that certain attitudes may naturally tend toward being autocentric, whereas other attitudes may naturally tend toward being exocentric, and that this could provide the basis for a principled account of the availability and/or obligatoriness of the tacit shift for the different attitude verbs. This suggestion dovetails nicely with Anand and Hacquard (2009; 2013)'s claim that the non-availability of "epistemic" readings for (semi-)modals under attitude verbs like *want* is due to the type of attitude reported. However, this suggestion does not make sense of the differences between *might* and *plausible* with regard to the alleged tacit shift.

## 5.6 Belief reports on a non-epistemic analysis of *might*

### 5.6.1 Avoiding second-order commitments

Say that we accept a non-epistemic analysis of *might* and take (327) to provide the content of the claim made with (326).

(326) There might be landmines.

(327) Circumstances leave open the non-trivial possibility that there are landmines.

We would then have a straightforward explanation for why a belief report sentence embedding (326) was not used to report a second-order belief. For to believe that circumstances leave open some possibility is *not* to have a belief about information. In much the same way, to believe that circumstances make it possible for something to happen is not to have a belief about information either. This is plausibly the sort of belief we express in a sincere utterance of (328).

(328) John can jump ten feet.

And just as we do not expect to report a second-order belief with a belief report sentence embedding (328), so on a non-epistemic analysis of *might* we should not expect to do so with a belief report sentence embedding (326) either.

Does an individual who believes that John can jump ten feet need to *explicitly* believe that the circumstances make it possible for him to do so? No. Nor does an individual who believes that there might be landmines need to explicitly believe that the circumstances leave open the non-trivial possibility that there are. These are merely ways of glossing that the modal relation an individual has a belief about is a circumstantial one: it holds, or would hold, in virtue of circumstances being the way they are, rather than in virtue of the content of an information state being what it is. This is a distinction that individuals who hold modal beliefs can be implicitly aware of, even if they never explicitly formulate it as such.<sup>86</sup>

Indeed, this distinction would seem to be one that we are aware of. Take (329), for example, in which causal explanations like those from Chapter 2 are included within the scope of a belief verb.

(329) Norma believes that she might be the murderer. She also believes the only reason she might be ...

---

<sup>86</sup> It would seem to matter, however, that we not gloss the modal belief as a *de re* one, whether about the circumstances themselves or about the propositions assumed to constitute the circumstances. For the beliefs one has about circumstances affect the beliefs one has about circumstantial modalities. But the former can include false beliefs, and so we could not take the latter to be *de re* beliefs about circumstances. But nor could we take them to be *de re* beliefs about propositions. If they were, learning that one had a false belief about the circumstances should not cause one to revise one's modal beliefs. And yet it can have this effect. Learning that John's training was as a rower and not, as I had mistakenly believed, as track and field athlete, I will not continue in my belief that John can jump ten feet but instead maintain only that he could have done so, had he trained as I'd believed.

- a. ?is because her memory of that awful night is so spotty.
- a'. is because her temper on that awful night was so uncontrollable.

With ((329)a), unlike with ((329)a'), we report what seems to be an unreasonable belief for Norma to have. This makes sense if her belief that she might be the murderer is a belief about a circumstantial modal relation, not an epistemic one. If her belief was about an epistemic relation, we would expect ((329)a) to sound as reasonable as (330) does.

- (330) Norma believes that it's plausible that she's the murderer. She also believes the only reason it's plausible is because her memory of that awful night is so spotty.

The same distinction holds when it comes to modal beliefs for which the information of someone other than the belief-holder is supposed to be relevant. These "exocentric" beliefs seem not to be about epistemic relations, as ((331)a) would require, but instead about circumstantial ones, as ((331)a') suggests.

- (331) I just spoke with the sergeant about our group's progress. She agrees that there might be landmines. But she also thinks that's only because ...
  - a. ?we haven't heard back from our scouting party yet.
  - a'. the terrain is so pocked.

If these "exocentric" beliefs were about epistemic relations, in contrast, we would expect ((331)a) to sound as reasonable as (332) does.

- (332) I just spoke with the sergeant about our group's progress. She agrees that it's plausible that there are landmines. But she also thinks that's only because we haven't heard back from our scouting party yet.

What holds for belief verbs holds for other attitude verbs too, which is nice given the struggles of the anaphoric analysis to generalize across attitude verbs. Take (333).

- (333) John's wife realized that he might be contagious.
  - a. ?But she also realized it's only because she didn't ask him his diagnosis.
  - a'. But she also realized it's only because he works in the worst hospital ward.

The sort of content John's wife adopts an attitude toward is apparently one whose truth she can realize is underwritten by non-epistemic facts, but not one whose truth she can realize is underwritten by epistemic ones. She apparently adopts an attitude toward a non-epistemic content, that is. This follows straightforwardly from a non-epistemic analysis.

### 5.6.2 Making sense of uncertainty implications

On an anaphoric analysis, believing  $\diamond p$  was identified with a state of mind similar to a specific sort of uncertainty about  $p$ .<sup>87</sup> Yet recall from Section 5.5.2 that there were at least three different states of mind a belief-holder could be in with regard to the prejacent  $p$  when a true report was made with a belief report sentence embedding a BMS.

*State of Mind 1*: Uncertain whether  $p$ . Believes neither  $p$  nor  $\sim p$ .

*State of Mind 2*: Uncertain whether  $p$ . Believes  $\sim p$ .

*State of Mind 3*: Certain whether  $p$ . Believes  $\sim p$ .

For the anaphoric analysis, the second and third states of mind had to be accommodated by appeal to tacit shifts. In this section, I want to show how we can make sense of these different states of mind that accompany true belief reports if we appeal to beliefs with contents like (327). I also want to show how this appeal can make sense of the relative frequency with which these beliefs are accompanied by this first state of mind, with its specific sort of uncertainty. To do so, let me introduce a distinction between two types of modal beliefs.

Our beliefs about circumstantial modalities are formed on the basis of our beliefs about the circumstances. Often we take, or intend to take, all of the latter into account in forming the former. In doing so, we form what we might call *all-things-considered* modal beliefs. Sometimes, however, we only take some of our beliefs about the circumstances into account, thus forming what we might call *relative* modal beliefs.

For the most part, our relative beliefs will be expressed in explicitly relativized ways, as in (334) and (335) (cf. Lewis 1976, 150; Kratzer 1981, 303–6).

(334) Given his *symptoms*, John might be contagious. Given his *test results*, he can't.

(335) Given his *training*, John can swim. Given his current state of *inebriation*, he can't.

Yet in some contexts we can drop the modifier restricting a modal sentence and express a relative belief by uttering just that bare modal sentence itself. We can do so in contexts in which some of our beliefs about the circumstances have been tabled, the strength of our evidence for these beliefs having been called into question. We can also do so in contexts in which it is clear that only a very specific type of circumstance is relevant, oftentimes instructional contexts. These are the two types of contexts that Section 5.5.2 made relevant. In what we might call the *default* context, however, the belief expressed in uttering a bare modal sentence will be an all-things-considered one. This is what we expect given basic principles of informativeness. Cooperative speakers can ignore considerations they take to be relevant to some matter—for example, to the truth of the prejacent—but they are generally expected to

---

<sup>87</sup> While the compatibility of the content of one's beliefs with  $p$  does not by itself require uncertainty, the addition of question-sensitivity, at least as defined by Yalcin (2011), seems to.

make this explicit.<sup>88</sup> As such, we can expect beliefs expressed in sincere utterances of bare modal sentences to generally be all-things-considered ones.

Given this distinction between all-things-considered and relative beliefs, let us return to our three states of mind and see how a person who believed that  $p$  was a non-trivial circumstantial possibility could be in any one of them. If this person's belief was an all-things-considered one, she would be in our first state of mind. If she was in our second or third state of mind, her belief would be a relative one. A person who believed  $\sim p$  could not form an all-things-considered belief that  $p$  was a non-trivial circumstantial possibility. Instead, the belief she formed on the basis of all of her beliefs about the circumstances would have to be that  $p$  was trivially impossible. Hence, to be in our second or third state of mind requires that the belief that  $p$  is a non-trivial circumstantial possibility be a relative one. If someone held this relative belief as a result of questioning the strength of her evidence for her belief that  $\sim p$ , she would be in our second state of mind. If she held it without such questioning, but instead as a result of attending to only a very specific type of circumstance, she would be in our third.<sup>89</sup>

Turn now to the relative frequency with which a reported belief-holder is in our first state of mind. There seem to be two reasons for this. The first deals with the fact that modal beliefs, to borrow a phrase from Yanovich, are always “something positive [on] our theory” (2014, 104). They require entertaining and endorsing a specific content. Yet someone who believed  $\sim p$  is unlikely to entertain whether  $p$  is a non-trivial possibility. Or rather, she is unlikely to do so unless forced to attend either (i) to the strength of her evidence for believing  $\sim p$  or (ii) to only a very specific type of circumstance. Someone attending to the question of whether  $p$  obtains, however, but who remained agnostic on the matter, would *ipso facto* be likely to entertain (and endorse) that there was this possibility. Hence, the belief that  $p$  is a non-trivial possibility is a belief that someone in our first state of mind is more likely to form.<sup>90</sup>

An all-things-considered belief is also more likely to be expressed in an utterance of a BMS, and so also more likely to be (faithfully) reported with a belief report sentence embedding a BMS. Indeed, it may be even more difficult in the default context to report a relative belief with an embedded BMS than it is to express one with an unembedded one. For we need not just a context in which a relative belief can be transparently expressed, but also positive reason to believe that the reported belief-holder has that sort of belief. Imagine, for example, that David Lewis utters (336) to make a point about relative modalities and context sensitivity (Lewis 1976, 150).

(336) I can speak Finnish. My anatomy, unlike an ape's, is of the right sort.

---

<sup>88</sup> Cf. Mandelkern (2017, 120–21) for a related point that speakers are generally expected to present their strongest evidence on a given matter.

<sup>89</sup> Yet I am not suggesting that these three states of mind are the *only* possible ones a belief-holder could be in when a true report was made with a belief report sentence embedding a BMS. In particular, a person could also believe  $p$ , rather than  $\sim p$ . What we would then report would either be an all-things-considered belief that  $p$  was a trivial possibility or a relative belief that  $p$  was a non-trivial one.

<sup>90</sup> Likely, but not necessarily guaranteed. For, as DeRose (1991) pointed out, and as Yalcin struggled to explain, someone non-committal on  $p$  could also remain non-committal on whether  $p$  was a possibility. They would do so, on my account, if they did not take themselves to be a competent judge of whether circumstances left open the non-trivial possibility that  $p$ .

Lewis would thereby have expressed that he had the right anatomy to speak Finnish, even if not the right training. Still, if he subsequently uttered (337), it would be very difficult to hear him as reporting anything other than the default belief that he had the right training to do so.

(337) My parents think that I can speak Finnish too.

To hear him as reporting a relative belief of the sort he had just expressed, we would need reason to believe that his parents attended to the sorts of questions that interested him. Hence, it seems to be difficult to report relative modal beliefs of *any* sort with bare modal sentences embedded in belief reports, a second reason for the relative frequency with which a reported belief-holder is more likely to be in our first state of mind.

### 5.6.3 Comparison to epistemic analyses

To conclude my discussion of how a non-epistemic analysis of *might* can be applied to belief report sentences embedding BMSs, let me draw two comparisons. The first is to epistemic analyses in general, the second to anaphoric analyses in particular.

On epistemic analyses, there is a fundamental divide between types of belief reports. For so-called autocentric reports, a belief-holder's information is supposed to be relevant for the interpretation of an embedded BMS. For so-called exocentric reports, like those from Section 5.5.2, someone else's is. I deny that a BMS is ever interpreted relative to an information state, however. Instead, for both reports, I hold that the beliefs reported are beliefs about what the circumstances leave open. The difference between these types of reports, on my analysis, is a difference in whether they report all-things-considered beliefs or relative ones. And this sort of difference, it turns out, is a general one, applying to other sorts of beliefs about circumstantial modalities to which an autocentric/exocentric divide could not apply.

On anaphoric analyses, there is a further fundamental divide between types of belief reports. Autocentric ones report an attitude toward the prejacent. Exocentric ones report a belief with an epistemic modal content. This divide, as we have seen, is problematic in any number of ways. But these problems can be avoided if we take both sorts of reports to report beliefs with modal contents, and to differ only in whether those beliefs are all-things-considered or relative ones. With non-epistemic modal contents of belief we can make sense, without the troublesome tacit shift, of how *might* contributes to the contents of other reported attitudes and speech acts. We can correctly predict that it will not make an epistemic contribution in these cases any more than in the others, unlike *plausible*, which makes an epistemic contribution in all. We can avoid making unique appeal in the semantics for epistemic modals to question sensitivity. We can have attitudinal contents available for anaphora and inference. And we can make sense of the different states of mind of the subjects of "autocentric" and "exocentric" reports, while still taking both to have modal beliefs. I thus conclude that a non-epistemic analysis is preferable as a whole to an anaphoric one in explaining the lack of second-order commitments for belief report sentences embedding BMSs.

### 5.7 An objection to a non-epistemic analysis

On the non-epistemic analysis that I have developed, it is supposed to be possible to have the relative belief that  $p$  is a non-trivial possibility while also believing that  $\sim p$ . A similar possibility is likely to hold for most other attitudes too. And yet, it might be objected, if there is this sort of possibility, then (338) and (339) should report and command possible (and possibly coherent) states of mind.

(338) I believe that God is not good and that he might be good.

(339) Imagine that John is not contagious and that he might be contagious.

However, the objection continues, (338) and (339) seem anomalous, which is not predicted on my non-epistemic analysis. To explain the anomaly of (338) and (339), an anaphoric analysis is supposed to be necessary. Since that analysis takes the modal to be evaluated directly relative to the content of the information state determined by the attitude verb, we end up reporting or commanding an incoherent state of mind when embedding an “epistemic contradiction” like (340).

(340)  $\sim p$  and  $\diamond p$

For the prejacent’s negation is supposed to be added to the content of the information state, but the modal requires the prejacent to be compatible with that content too, and these are incompatible demands (Yalcin 2007; Anand and Hacquard 2013; for alternative accounts, cf. Schnieder 2010; Crabill 2013; Yanovich 2014).

In response to this objection, I would remind readers that there are constraints on the expression of relative modality whenever the circumstances relative to which a prejacent is to be evaluated—i.e., its *modal base* (Kratzer 1981; 1991)—have not been explicitly specified. One plausible constraint is that an assertion of  $p$  results in  $p$  being added to the presumed circumstantial modal base for subsequent occurrences of matrix-clause modals, unless that assertion has somehow been waived or bracketed. Hence, if you wanted to express a coherent relative belief that  $\sim p$  was a non-trivial possibility, you could not do so with a BMS once  $p$  had been asserted. Instead, you would be expressing the incoherent belief that, relative to circumstances that included  $p$ ,  $\sim p$  was a possibility.

What holds for matrix clauses plausibly holds for embedded ones too. The occurrence of a sentence within an embedded clause results in the content of that sentence being added to the presumed circumstantial modal base for subsequent occurrences of modals within the clause. As such, embedded epistemic contradictions cannot be used to report (or command) coherent relative attitudes, but instead report (or command) incoherent ones. It is coherent to believe or imagine that some circumstances leave open the possibility that  $\sim p$ , and also to believe or imagine that  $p$ . It is not coherent, however, to imagine that circumstances that include  $p$  leave open this possibility.

What I am proposing as a solution to the apparent anomaly of embedded epistemic contradictions is thus in some ways very similar to what anaphoric analyses propose. The anomaly results from the content of the one conjunct being added to the modal

base for the other conjunct. However, I take this modal base to be circumstantial, not epistemic. I thus do not take there to be anything incoherent *per se* in believing that  $\sim p$  is a possibility while also believing that  $p$ . I take the incoherence to arise only from believing that  $\sim p$  is a possibility *relative to*  $p$ .

## 5.8 Conclusion

The observation that we do not report second-order beliefs with belief report sentences embedding BMSs has motivated the view that we instead report some other attitude like suspicion, or weak belief. This view, which proceeds from the assumption that *might* is to be interpreted epistemically, requires far-reaching complications to theories of attitude verbs and their interactions with modals. It also has not insignificant empirical limitations, the only potential solution to which would undermine the explanatory force of the view. These complications and limitations can be avoided. Both the illusion of an explanandum and the problems with the explanans are generated by the assumption of epistemicity. If we instead interpret *might* circumstantially—something we independently have reason to do—we would never expect the relevant belief report sentences to be used to report second-order beliefs at all. With this circumstantial interpretation of *might*, we maintain a standard view of the attitudes, and of attitude reports, and also provide greater (and more principled) empirical coverage of the relevant reports. The best explanation of the lack of second-order commitments for belief report sentences embedding BMSs is that *might* is a circumstantial modal, not an epistemic one.

## **Part Three: Objections**

## Chapter Six: The Meaning Difference Between *Might* and *Can*

Epistemic modality [expressed by *might*] and circumstantial modality [expressed by *can*] involve a different categorization of the facts.

~ Kratzer (1981, 302)

### 6.1 Two competing analyses

In formal semantics, the standard analysis of the meaning difference between *might* and (non-deontic uses of) *can* is the Kratzerian one. On this analysis, both (341) and (342) express the compatibility of a common prejacent with a set of facts but differ in how those sets are characterized.

(341) John might easily be a military man.

(342) John can easily be a military man.

For (341), the relevant facts are characterized in epistemic terms: as our evidence, or what we know. For (342), they are instead characterized simply as the circumstances (cf. Kratzer 1981, 302; 1991, 646; 2012, 33; 2013, 188–89). Now the entire argument of this dissertation is *against* an epistemic characterization of the relevant facts for *might*. In this final section of the dissertation, I thus want to do two things.

My first aim is to call into question the adequacy of the Kratzerian analysis of the difference between (341) and (342) as a difference in epistemic versus circumstantial modality. I do this in the current chapter in two ways, having first demonstrated the basic continuity of Kratzer's analysis from her earlier to her more recent work in Section 6.2.<sup>91</sup>

First, in Section 6.3, I focus on a contrast between *might* and *can* that provides the *intuitive* motivation for the Kratzerian analysis. This contrast, I argue, turns out to not be well explained by that analysis at all. For the contrast extends to cases where the modality expressed by *might* is standardly assumed to be metaphysical, not epistemic. Further, there is a related contrast between *might* and *can* that points to the inverse of the Kratzerian analysis if given the same sort of explanation as the original one. Thus, even if a Kratzerian analysis is possible for the original contrast, it would not seem to be explanatory. Rather, there seems to be some other, underlying difference between the modalities expressed by *might* and *can* for which an epistemic/circumstantial distinction could provide only a partial, and redundant, explanation.

Second, in Section 6.4, I focus on the Kratzerian non-ambiguity project that provides the *theoretical* motivation for her analysis of the difference between (341) and (342). This is the project of attempting to account for *modal flexibility* (or the ability of a single modal to be used to express multiple, distinct flavors of modality) without having to posit *modal ambiguity* (or multiple semantic entries for that modal). This project, I argue, is also ill-served by Kratzer's assumption that the difference

---

<sup>91</sup> In essence, I show that while Kratzer's analysis of *how* to generate the meaning difference has changed in well-known ways from her earlier to her more recent work—from contextually-provided conversational backgrounds to syntactically-provided modal anchors—her analysis of what this difference is seems to have remained the same.

between the modalities expressed by (341) and (342) is a difference in the intensional characterization of facts, or what I will sometimes call the *intensional assumption*.

In Section 6.5, I conclude the chapter by revisiting the arguments for Kratzer's non-ambiguity project. Abandoning the intensional assumption does not necessarily require abandoning this project; yet I will suggest that these arguments provide no compelling reason for adopting it either. This fact will be significant in the following chapter when, in response to the different temporal properties of *might* and *can*, I propose different semantic entries for the two: a proposition-level semantic entry for *might* and an event-level semantic entry for *can* (cf. Brennan 1993; Palmer 2001).

My second aim in this final section of the dissertation is in fact just this: to provide an alternative to the Kratzerian analysis of the contrast between *might* and *can*. I do this formally in the following chapter, in the context of the argument that the different temporal properties of the two cannot be attributed to the epistemic/circumstantial distinction. Yet throughout this chapter I also informally develop this alternative too. The basic idea is that with *might* and *can* we attribute relative modal properties to different sorts of entities. With *might*, we attribute the relative property of being possibly true to a *proposition*. With *can*, we attribute the relative property of being able to develop in a certain way to an *individual, location, situation*, etc. (This property is often called a *potential*.) In attributing relative properties to these different sorts of entities, we express different sorts of relations, not just the same compatibility relation to different sorts of facts. And this relational difference makes sense not only of the Kratzerian contrast between *might* and *can* that provides the intuitive motivation for her analysis but also of the additional contrasts from Section 6.3 that Kratzer's proposed fact-based difference does not.

## 6.2 Kratzer's analysis

Kratzer's analysis of modals has changed in important ways from her earlier work (cf. Kratzer 1981; 1991) to her more recent work (cf. Kratzer 2012; 2013). Yet I will argue in Section 6.2.1 that it has not changed in regard to the fundamental concern of this chapter: namely, the way in which *might* and *can* differ in meaning. In Section 6.2.2, I then review the ways in which the analysis *has* changed and show that they affect only how that meaning difference is generated, not what it is. Yet Kratzerian exegesis is not my main aim in this chapter. Although I think that my interpretation of Kratzer's recent work makes the best sense of various comments and commitments, I am also open to interpretations that see in it further support for my own contention that *might* is not epistemic, and that *might* and *can* thus do not differ along an epistemic/non-epistemic dimension either. In Section 6.2.3, in fact, I even argue that there is implicit in Kratzer a relational alternative to what I take to be her explicit fact-based analysis. Unfortunately, if I am correct, this alternative has so far remained implicit.

### 6.2.1 The continuity of Kratzer's analysis

From her earlier to her more recent work, Kratzer has displayed remarkable continuity on two points in her analysis of the difference between the modalities expressed by *might* and *can*, which she calls *epistemic* and *circumstantial*.<sup>92</sup>

First, she continues to maintain that this is a difference in the categorization, or kind, of facts that these modalities involve. Reiterating the epigraph of this chapter, for example, Kratzer claims that, “semantically, [the two] modals differ with respect to the kinds of facts they depend on” (2012, 24). And this claim is reiterated again and again (cf. Kratzer 2012, 50, 51; for her older work, cf. Kratzer 1981, 302, 306; 1991, 646).

Second, she continues to use the same terms to describe this fact-based difference. In her early work, she often uses terms like *evidence* and *circumstances* to describe the different kinds of facts relevant for the two types of modalities (cf. Kratzer 1981, 302; 1991, 644, 646). And in her more recent work she also uses these exact same terms. She writes, for example, that “the kind of facts that are targeted by the two types of modals are different in kind ...: external or internal *circumstances* of people, things, or places ... contrast with *evidence* of things (2012, 54; emphasis mine). And these terms are likewise reiterated again and again in her characterization of the alleged fact-based difference (cf. Kratzer 2012, 33; 2013, 188–89).

There is thus fundamental continuity in Kratzer's work in that part of her analysis which is the concern of this chapter. Yet it should be noted that while Kratzer talks interchangeably in developing this analysis of different *kinds* of facts, and of a different *categorization* of the facts, the latter is strictly speaking the more accurate description. For there is no kind of fact that is a piece of evidence *per se*, as Kratzer herself is aware. Indeed, Kratzer writes that “whatever exists in a world ... should in principle qualify as potential evidence of things in that world” (2012, 53). As such, the relevant internal and external circumstances for a *can*-claim could very well turn out in some cases to be the same as the relevant evidence for a *might*-claim too. This is the lesson of Nauze (2008), and is demonstrated with cases like (343)-(344) and (345)-(346).

(343) Given only his psychiatric profile, John might easily be a military man.

(344) Given only his psychiatric profile, John can easily be a military man.

(345) Given only his digital dexterity, John might easily be a professional pianist.

(346) Given only his digital dexterity, John can easily be a professional pianist.

---

<sup>92</sup> Kratzer in fact defines circumstantial modality broadly, to encompass all non-epistemic modalities, including deontic ones. For Kratzer, these modalities differ amongst themselves, and many of them still further from epistemic modalities, in the value of what she calls their *ordering source*. However, for the non-deontic uses of *can* that concern us here—what Kratzer calls *pure circumstantials*—there is not necessarily this second source of variation (cf. Kratzer 1991, 646). Instead, pure circumstantials have what Kratzer calls an empty ordering source, and epistemics may too. Unless noted otherwise, I will generally restrict my attention in the following two chapters to such non-deontic uses of *can* and the “pure” circumstantial modalities expressed by them. Hence, it is appropriate to talk in terms of *the* difference between them and epistemic modalities on the Kratzerian analysis, even if this would not be appropriate for other sub-classes of circumstantial modalities.

Hence, it is not as though there are different *kinds* of facts which could be relevant for *might*- and *can*-claims but instead, on Kratzer's analysis, different *characterizations* of what could sometimes turn out to be the same facts.

## 6.2.2 The discontinuities in Kratzer's analysis

Despite the continuity of her fact-based analysis, there are important discontinuities in Kratzer's work, and these have tended to obscure what has remained the same (cf., e.g., the discussion in Matthewson 2016). These relate to what Kratzer calls epistemic and circumstantial *conversational backgrounds*, which are supposed to be contextually-provided functions from possible worlds to propositions, respectively, that are either known in, or circumstances, of those worlds. But I will argue that these discontinuities do not affect the argument of the previous section.

### 6.2.2.1 Conversational backgrounds

In her early work, Kratzer had used epistemic and circumstantial conversational backgrounds to generate what she sometimes calls the *modal base*. This is the set of worlds that the modal quantifies over, testing to see whether its prejacent is true in all or some or none of them. And it is the same for Kratzer as the set of worlds verifying the value of the epistemic or circumstantial background (Kratzer 1981, 297; 1991, 644).<sup>93</sup> She also appealed to additional types of conversational backgrounds to function as what she calls *ordering sources*: that is, to induce an ordering on the worlds in the modal base based on their compliance with a given set (and type) of norms. These additional types of backgrounds Kratzer refers to as *non-realistic*. Unlike epistemic and circumstantial backgrounds, the propositions they assign to a world do not necessarily obtain in it. For the propositions they assign represent the content of a set (and type) of norms in that world, and what is normatively required is rarely the same as what is actually the case.

Together with non-realistic ones, epistemic and circumstantial backgrounds helped to determine the modal flavor of a claim in Kratzer's early work. However, mere differences in the set of worlds quantified over, or in how these worlds are ordered, do not themselves suffice to generate flavor differences. Indeed, it is flavor-internal differences of precisely this sort that help to explain how, for example, a deontic claim that is true in one context may be false in another. This may be due to changes in the relevant circumstances from one context to the other (in which case, there will be a change in the modal base) or instead to changes in the applicable norms (in which case, there will be a change in the ordering of the worlds in the modal base).

---

<sup>93</sup> von Fintel and Heim (2011) suggest that Kratzer uses the term *modal base* to refer to the conversational backgrounds themselves, not to the sets of accessible worlds determined in this way by their values (2011, 42). This seems right in some contexts. Yet Kratzer continues to use the term even when, as we will see, she no longer espouses conversational backgrounds. So for ease of exegesis, I will generally use the term *modal base* in this chapter to refer to the set of worlds that a modal quantifies over. However, the use suggested by von Fintel and Heim should also be kept in mind, especially when Kratzer is being quoted. Finally, over and above these two uses, there is the use of the term *modal base* that I have adopted throughout the rest of this work: namely, to refer to the *value* of a conversational background. In this chapter, I will avoid this third use, unless I explicitly note otherwise.

What determines the modal flavor of a claim is thus not the set of worlds quantified over, nor the ordering of the worlds, but instead the accessibility and ordering relations themselves. Modal flavor depends on whether the accessible worlds are epistemically or circumstantially accessible ones, on whether the ordering of those worlds is deontically or teleologically determined, etc. Modal flavor, in other words, depends in part in Kratzer's early work on the *type* of conversational background used to determine the modal base: namely, whether it is epistemic or circumstantial.

Two subsequent developments in the literature, however, led Kratzer to reconsider her use of epistemic and circumstantial conversational backgrounds. One was Nauze (2008)'s criticism of the attempt to derive the differences between modal flavors from differences in context. The other was Hacquard (2006; 2010)'s argument that an observed correlation between flavor and syntactic height could be accommodated within the Kratzerian framework. In response to these two developments, Kratzer claims that she had earlier made an "erroneous assumption that the two types of modals semantically select modal bases with distinctive semantic properties: circumstantial backgrounds ... and epistemic backgrounds" (2012, 24). Instead of there being such semantically distinctive modal bases provided by context, Kratzer suggests that there are, at different heights in the syntactic tree, "different kinds of semantic objects from which modal bases can be systematically projected" (2012, 55). Here she is following the lead of Hacquard, who calls these objects *modal anchors*.

This, then, is the first major discontinuity in Kratzer's work. Yet despite rejecting epistemic and circumstantial conversational backgrounds, everything that Kratzer has said so far is consistent with it still being the intensional characterization of facts that determines the *type* of the modal base, even if that modal base is now syntactically projected rather than semantically selected. Indeed, this would even seem to be Kratzer's position. For she holds both (i) that such modal bases "target relevant bodies of facts in the evaluation world and track them via counterpart relations" and (ii) that the relevant facts differ in kind from *might* to *can* (2012, 54). Since the relevant facts are further said to differ in the exact same way as in in her early work (cf. Section 6.2.1), Kratzer would thus seem to still be committed to characterizing modal bases as epistemic or circumstantial, depending on how they are determined, even if she rejects the characterization of conversational backgrounds in these terms.

This commitment, however, is easily obscured by Kratzer's subsequent comments in response to Nauze' criticism. "It now seems to me," she continues,

a hopeless enterprise to try to characterize formal objects like conversational backgrounds as "circumstantial" versus "epistemic." Both types of backgrounds are functions that map possible worlds to sets of factual premises. What is it that would allow us to single out some of those functions as epistemic, but not circumstantial, or the other way round? (2012, 24).

In fact, the enterprise is, in some ways, not nearly as hopeless as Kratzer makes it out to be. In giving a compositional semantics for phrases like *given what we know* and *given the circumstances*, for example, we will use functions that we characterize as epistemic but not circumstantial, and vice versa. Of course, the functions themselves

are not epistemic or circumstantial, but we can characterize them as such insofar as we use them to represent the meaning of a given phrase or lexical item. What is hopeless is thus not the attempt to characterize formal objects as epistemic or circumstantial, but instead the attempt to have *context* provide such characterizations. Context provides objects, not meanings (cf. Section 6.4.2, as well as Chapter 3, Section 3.3.3).

While Kratzer's rejection of epistemically or circumstantially characterized conversational backgrounds may thus seem to imply a similar rejection for modal bases, in fact the opposite is true. Not only do Kratzer's claims continue to commit her to characterizing modal bases as either epistemic or circumstantial, depending on how the modal anchors they are projected from are intensionally characterized, but so must *any* attempt to use syntactically provided modal anchors to generate the flavor difference between *might* and *can* maintain some such commitment. The facts provided by modal anchors may help to determine the modal base but, by themselves, they do not help to determine its type. For what if anchors for the two modals ever happened to provide the same set of facts, from which their modal bases were projected? And, more importantly, for the vast majority of modal anchors that provide different sets of facts, how do we distinguish flavor-internal differences in the modal bases projected from these sets from flavor-external ones?

The answers to these questions come from the intensional characterization of the facts provided by the modal anchors. If the relevant internal circumstances of John are also our relevant evidence, then it is the intensional characterization of those facts *as* evidence that allows us to differentiate (343) from (344) on the Kratzerian analysis. This is also what allows us to distinguish flavor-internal differences (where the facts differ, but the characterizations remain relevantly the same) from flavor-external ones (where the facts may be the same, but the characterizations differ).

Now it may not be necessary to characterize the facts provided by the modal anchors in the way that Kratzer does: namely, in evidential versus non-evidential terms. Yet whatever characterization of facts is chosen, it must be general enough to ensure that the differences between (347) and (349), and between (348) and (350), count as flavor-internal, while those between (347) and (348), and between (349) and (350), count as flavor-external.

(347) John might run four-minute miles.

(348) John can run four-minute miles.

(349) Mary might bench 125 pounds.

(350) Mary can bench 125 pounds.

And for this my guess is that Kratzer's characterization is the only possible one.

To conclude: a mere difference in the set of accessible worlds (or modal base) does not by itself generate a difference in modal flavor. Instead, we also need there to be a difference in the accessibility relation itself. Yet Kratzer does not change her view on what this difference is, though she does change her view on how the facts that determine it are provided, moving from context and conversational backgrounds in her early work to syntax and modal anchors in her more recent work. Despite this

change, it is still in both cases the intensional characterization of facts in evidential versus non-evidential terms that allows us to distinguish the flavor of modality expressed by *might* from the flavor expressed by *can*.

### 6.2.2.2 Modes of domain projection

Epistemic and circumstantial conversational backgrounds vary in whether their values are construed in contentful terms: for the former, they are; for the latter, they are not. In her recent work, Kratzer similarly claims that there are contentful and non-contentful (or factual) modes of domain projection (cf. Kratzer 2013, 193–97). However, Kratzer claims that the distinction between these two modes of domain projection cross-cuts her earlier distinction between epistemic and circumstantial conversational backgrounds. This is the second apparent major discontinuity in Kratzer’s work, and it may seem, to those who are aware of it, to undermine the entire argument of the previous section. Obviously, I do not agree. But first let me explain what motivates the new distinction, and then let me relate it to the old one.

To characterize a fact as a piece of evidence is to characterize it in a contentful way: namely, as the content of an evidential state. However, some facts are also themselves content-bearing, while others are not. The fact that I am feeling fit is a content-bearing piece of evidence that I am fit. The fact that I can run a four-minute mile is also a piece of evidence that I am fit, but not a content-bearing one. Focusing on content-bearing pieces of evidence used to support modal claims, Kratzer noted an interesting contrast. Whether a speaker can distance herself from the modal claim varies with the modifier used to specify that piece of evidence. For example, while (351) has no reading on which it is bad, with (352) there is such a reading (cf. Kratzer 2012, 24–25, 33–35; cf. also Matthewson 2016, 538–41).

- (351) According to the rumor, Roger must have been elected chief. But he actually doesn’t have to have been. [≈ Matthewson’s (24)]
- (352) Given the rumor, Roger must have been elected chief. ?But he actually doesn’t have to have been. [≈ Matthewson’s (23)]

To account for this contrast, Kratzer suggests that in the first case we are claiming, more or less, that the *content* of the rumor entails that Roger has been elected chief. Since we recognize that this content may well be false, we can subsequently distance ourselves from its alleged entailment. In the second case, however, we are supposed to be claiming that that the *fact* that the rumor exists entails that Roger has been elected. (Maybe a rumor with this content would only exist in a world where Roger *had* been elected chief.) And since we are claiming that one fact entails another, then we cannot subsequently distance ourselves from that second fact.<sup>94</sup> This is Kratzer’s account of the contrast between (351) and (352), and it suggests that with content-bearing pieces of evidence, like rumors, we have two distinct modes of domain projection: contentful and factual. Yet on Kratzer’s original analysis, evidence used to project a domain would always be used as the value of an epistemic conversational background, which would seem to entail an exclusively contentful mode of domain

<sup>94</sup> The reading on which (352) would be fine on this analysis, brought out by emphasis on *the rumor*, would be one in which we were making a claim about what the content of the rumor entailed.

projection. So the distinction between contentful and factual modes of domain projection seems to cross-cut the distinction between epistemic and circumstantial conversational backgrounds, a second major discontinuity in Kratzer's work.

I am not convinced that Kratzer's account of her contrast is on the right track, as I will argue in Section 6.3.2.<sup>95</sup> Assuming that it is, however, all that it really shows is that we may sometimes project a modal domain from the *matrix* content of an evidential state and sometimes from its *embedded* content. In this case, the matrix content of the evidential state could be *that there is a rumor that Roger was elected chief*, and the embedded content of that state would be *that Roger was elected chief*. Yet projection from either sort of content will be a contentful mode of projection.

For no fact is a piece of evidence *per se*, just as no fact is a known fact *per se*. To characterize a fact in either of these ways is to characterize in a contentful way: as the content of an evidential or knowledge state. Hence, if we treat these characterizations as essential, as what differentiates one flavor of modality from another, for example, then we are already committed to a contentful mode of projection for any facts used to project (or determine) the modal base. The alleged distinction between factual and contentful modes of projection from a body of evidence is really a distinction between projection from its matrix and embedded contents. There is no discontinuity, then, in the move from epistemic conversational backgrounds to the allegedly factual mode of projection from bodies of evidence, or at least not when it comes to the meaning difference between *might* and *can*.<sup>96</sup>

### 6.2.3 An implicit alternative to Kratzer's explicit analysis

My interpretation of the commitments and implications of Kratzer's recent work is not standard. Indeed, many have seen in Kratzer's work support for my own contention that *might* is not epistemic, and that *might* and *can* thus do not differ along an epistemic/non-epistemic dimension either. Of course, I would be glad for such support. But I do not think it is there. The intensional characterization as evidence of the facts used to determine the modal base plays the same critical, flavor-determining role in her early work (where those facts are provided as the value of a conversational background) as in her later work (where those facts are provided by a modal anchor). That said, I do think there is implicit in Kratzer a relational alternative to (what I take to be) her explicit fact-based analysis. Here are two examples of what I mean.

First, in her early work, Kratzer writes:

---

<sup>95</sup> In fact, it should be noted that the content of content-bearing pieces of evidence like rumors and feelings, etc., would not be used to project domains on Kratzer's original analysis but instead, functioning as the value of the ordering source, to restrict and order already projected domains. Why Kratzer does not attempt to account for the contrast between (351) and (352) in terms of her contrast between modal bases and ordering sources I am not completely certain.

<sup>96</sup> Perhaps it will seem there is an important discontinuity in the move from the *known* facts (the value of an epistemic conversational background) to the facts that constitute our *evidence* (the factual mode of projection from a body of evidence). Yet, for better or for worse, Kratzer seems to treat a fact as a piece of evidence if and only if it is a known fact. She writes, for example, that "the epistemic conversational background (*in view of the available evidence*) determines for every world the set of worlds which are epistemically accessible from it" (Kratzer 1991, 644; emphasis mine). Kratzer may well be mistaken in her view of the relation between evidence and knowledge. But, given that view, her view of the type of modal base involved in both cases remains the same.

If we use an epistemic modal [like *might*], we are interested in what else *may or must be the case* in our world, given everything we know already. And if we use a circumstantial modal [like *can*], we are interested in what *can or must happen*, given circumstances of a certain kind. ... Epistemic modality and circumstantial modality involve a different categorization of the facts (1981, 302; emphasis mine).

Now what may or must be the case are *propositions*, and what can or must happen are *events*. Kratzer is thus implicitly appealing to the distinction between propositional and event modality that is explicit in Palmer (1979; 1986; 2001) and that can also be seen as the precursor for my proposal in this final section of the dissertation.<sup>97</sup> For Kratzer is really concerned in the first case with the modal properties of propositions: that is, with whether, relatively speaking, they are possibly or necessarily true. And Kratzer is really concerned in the second with the modal properties of individuals, situations, locations, etc.: that is, with the potentials that determine what eventualities are possible or necessary for the future. However, Kratzer attempts to reduce the differences between the *entities* to which these properties are attributed to a difference in the characterization of the *facts* responsible for those properties.

Kratzer does something very similar in her recent work, here even building into her characterization of the facts a description of this relational difference itself.

The kind of facts that are targeted by the two types of modals are different in kind, though: external or internal circumstances of people, things, or places *that determine their possible futures* contrast with evidence of things *implying or suggesting the presence of other facts* in the past, present, or future (2012, 54; emphasis mine).

Here Kratzer comes her closest to explicitly formulating the distinction between event and propositional modality: for possible futures comprise events and facts are just true propositions. But since the internal and external circumstances of things may sometimes be the same as our evidence, Kratzer's attempt to build this distinction into her characterization of the facts does not suffice in this regard.

Kratzer claims, in what has so far been this chapter's theme, that "there is a subtle semantic difference between the two kinds of modals I grouped under the two headings. It is a difference in the kind of facts relied on" (2012, 51). Yet she adds that "the facts relied on seem to be different in a way that has proven difficult to

---

<sup>97</sup> In Palmer's formulation, propositional modality concerns "[a] speaker's attitude to the truth-value or factual status of the proposition," while event modality refers to "events that are not actualized, events that have not taken place but are merely potential" (2001, 8). These definitions can be generalized, both to bring them in line with my own proposal and to make them more empirically adequate. For example, propositional modality could involve the attribution of *any* modal property to a proposition. And with event modality, we could take the basic fact to be that the modal properties of individuals, situations, locations, etc., are specified with reference to events that need not be actualized, even though they sometimes are. This accommodates modal claims that entail that the relevant properties *have* been actualized on some occasions, or so-called *actuality entailments* (cf. Hacquard 2009).

characterize in formal terms” (2012, 50). Now I think that Kratzer’s difficulties are due to it being not the *facts* themselves that differentiate the modals but instead the *relations* they express to those facts. Hence, Kratzer ends up having to try to build into her characterization of facts a description of this relational difference itself. But what motivates Kratzer to adopt a fact-based analysis in the first place? In the next two sections of this chapter, I address both the intuitive and the theoretical motivation for the fact-based analysis and show that neither is ultimately very well-served by it.

### 6.3 On the intuitive motivation for Kratzer’s analysis

#### 6.3.1 Epistemic contradictions

Say that we are admiring a neighbor’s garden, and that one of the relevant facts in the context is that there are no hydrangeas growing there. Then it seems that I will say something false if I utter (353), but not necessarily if I utter (354) (cf. Kratzer 1991, 646).

(353) Hydrangeas might be growing here. [ $\approx$  Kratzer’s (21b)]

(354) Hydrangeas can grow here. [Kratzer’s (21a)]

A similar contrast arises for *possible* with regard to finite and non-finite complements (cf. DeRose 1991, 601–5; Vetter 2013, 4–5). If one of the facts that is relevant in the context is that Frank does *not* run four-minute miles, then the claim made with (355) seems necessarily to be false, but the claim made with (356) seems as though it could possibly be true.

(355) It is possible that Frank runs four-minute miles. [Vetter’s (1a)]

(356) It is possible for Frank to run four-minute miles. [Vetter’s (1c)]

This contrast in truth-value judgments carries over to cases where the proposition that we have assumed to be relevant in the context is instead asserted in conjunction with the modal proposition. While the conjunction in (357) seems anomalous—what Yalcin (2007) calls an *epistemic contradiction*—the conjunction in (358) does not.

(357) #Frank doesn’t run four-minute miles anymore, but it’s possible that he does so.

(358) Frank doesn’t run four-minute miles anymore, but it’s possible for him to do so.

Similarly, while (359) seems anomalous, (360) does not.

(359) #Hydrangeas aren’t growing here, but they might be growing here.

(360) Hydrangeas aren’t growing here, but they can grow here.

In sum, knowing or asserting that the content of *might*’s complement is false will generally have the effect of making a *might*-claim bad. Yet the same does not hold for

a *can*-claim, *mutatis mutandis*.<sup>98</sup> With *might*, but not *can*, it seems possible for there to be a certain sort of contradictoriness involving the prejacent.

The Kratzerian explanation for this contrast is that with *might* there is an epistemic characterization of the facts that determine the modal base, but with *can* there is not. Hence, we cannot ignore any of the known facts, or the available evidence, in making a *might*-claim, but we can, and do, in making a *can*-claim. Since it is a known fact that hydrangeas are not growing here, and since this known fact is incompatible with the proposition that they are, the *might*-claims above turn out to be false or contradictory. Yet since this known fact can be ignored for circumstantial modals, the *can*-claims are neither, at least not on this account (cf. Kratzer 1991, 646).

This contrast between *might* and *can* when it comes to so-called epistemic contradictions provides the intuitive motivation for the Kratzerian analysis of the meaning difference between the two as an epistemic/non-epistemic one. Yet I will now show that we should really be talking more generally of *modal* contradictions, and that the best analysis of the meaning difference between *might* and *can* here cannot be in terms of an epistemic/non-epistemic divide.

### 6.3.2 Modal contradictions

The contrast between *might* and *can* from Section 6.3.1 is not limited to cases in which *might* is assumed to have an epistemic interpretation. Instead, it carries over to future-oriented cases, like (361), in which *might*'s interpretation is standardly assumed to be metaphysical, perhaps necessarily so (cf. Klecha 2016).<sup>99</sup>

- (361) #John might win tomorrow's race, but he won't.  
 (362) John can win tomorrow's race, but he won't.

In addition, there are also cases where it is not a *might*-claim that is anomalous in a conjunction, but instead a *can*-claim, as in (363)-(366).

- (363) ?John never hits the bullseye, but he can.  
 (364) John never hits the bullseye, but he might.  
 (365) ?These crops have never grown above six feet, but they can.  
 (366) These crops have never grown above six feet, but they might.

---

<sup>98</sup> Assuming, that is, that the content of *can*'s complement can even be an object of knowledge or assertion in the first place. If the contrast between *might* and *can* has the same sort of explanation as the contrast with *possible* between its finite and non-finite complements, then this is not at all obvious.

<sup>99</sup> DeRose (1998) has argued that any apparently metaphysical reading of a future-oriented *might* is derived from its actual epistemic one. But the derivation is implausible. It requires the facts that determine the epistemic modal base to include the total laws and history, and whatever follows from them, in most cases not as *known* but instead merely as *knowable*, in some contextually relevant way. Further the motivation for the derivation is a desire to avoid positing an ambiguity between future- and present-oriented *might*, in conjunction with the belief that the badness of the examples in Section 6.3.1 can only be attributed to *might*'s having an epistemic meaning. But this belief I have already shown to be false (cf. Chapter 3, Section 3.4 and Chapter 5, Section 5.7). Hence, the desire to avoid an ambiguity can be satisfied in some other way: for example, by positing a single, circumstantial reading for both cases.

Just as our evidence does not support the claim that hydrangeas might be growing here if we know that they are not, so it intuitively does not support the claim that John can hit the bullseye if we know that he never does (cf. also Thalberg 1972 for what constitutes evidence for an ability attribution). In both cases, however, our evidence *may* support the alternative modal claim.

If this contrast in what our evidence supports has the same explanation as the contrast from Section 6.3.1, then we have to hold that with *might* we can ignore certain known facts that we cannot ignore with *can*: for example, the known fact that the crops have never grown above six feet, or that John never hits the bulls-eye. To explain the two contrasts, we would thus require both epistemic *and* non-epistemic characterizations of the relevant facts: both for *might* and for *can*. And this would suggest that the real difference between the two is not how the relevant facts are characterized. Instead, a further, underlying difference can make an epistemic characterization seem apt in some cases but not others.

Much the same holds for the extension of Kratzer's original contrast from *might*'s apparently epistemic uses to its apparently metaphysical ones. If both uses contrast with *can* in the same way, this suggests that something more fundamental, and common to both, is responsible for their contrast with *can*.

Say, for example, that both uses of *might* attribute to a proposition the relative property of being possibly true. A proposition cannot have this property relative to its negation, with which it is trivially incompatible. Hence, adding the prejaçant's negation to the modal base—here understood as the set of facts relative to which the content of the modal claim is evaluated for truth—should result in contradictoriness *independently* of how the facts that the modal base comprises are characterized. Say, in contrast, that with *can* we attribute modal properties to individuals, etc. These modal properties, also known as potentials, are specified with reference to eventualities that need not be actualized. Hence, no contradictoriness should result from adding to the modal base the proposition that the relevant eventuality is *not* currently being actualized. And again, this is independently of any characterization of the facts that the modal base comprises.

This sort of difference between the entities to which we attribute modal properties would also seem to explain the contrasts in (363)-(366). While the non-realization of a potential—a modal property of an individual, etc.—is no indication of its non-existence, we also have no good basis for attributing one if the hypothesized potential has never been realized on any occasion on which it could have been. This explains why (363) and (365) are anomalous, even while (358) and (360) are not. Likewise, while a proposition  $p$  cannot possibly be true relative to a set of facts that includes its negation, a proposition  $p_1$  can possibly be true relative to a set of facts that includes  $\sim p_2, \sim p_3$ , etc. That is, while the proposition that hydrangeas are growing here at  $t_1$  is incompatible with the proposition that they are not growing here at  $t_1$ , the proposition that John has not hit the bulls-eye at  $t_1, t_2 \dots t_n$  need not be incompatible with the proposition that he hits the bulls-eye at  $t_{n+1}$ . And this explains why (364) and (366) are fine, even while (357) and (359) are not.

Of course, more must be said to precisify this characterization of the difference between *might* and *can* in terms of the different sort of entities to which we attribute

modal properties. But the more that must be said cannot be that *might* and *can* differ in whether the relevant facts are characterized epistemically, at least not if we are to provide the same sort of explanation for the contrasts in this section as in the previous one. We could, perhaps, attempt to explain only the first contrast in terms of how the facts are characterized. But there is something unsatisfactory about this attempt, especially given that the first contrast does not have to be explained in this way. If we could attribute the two contrasts to the same underlying principle, it would be better to do so. And we can, if we hold that with *might* we are concerned with the relative modal properties of propositions, while with *can* we are concerned with potentials, which are the relative modal properties of individuals, locations, situations, etc.

A different argument against the attempt to account for the meaning difference between *might* and *can* in epistemic and non-epistemic terms comes from a case where other allegedly epistemic modals show *similarities* to acknowledged non-epistemics when it comes to modal contradictions. Recall Kratzer's contrast between modifiers from Section 6.2.2.2, intended to motivate the distinction between contentful and factual modes of domain projection for allegedly epistemic modals. It turns out that this same sort of contrast arises with what she would take to be circumstantial modals, including deontic and teleological ones. While the normative (367) has no reading on which it is bad, for example, the normative (368) does.

- (367) According to the rumor, Roger must go to confession. But he actually doesn't have to.
- (368) Given the rumor, Roger must go to confession. ?But he actually doesn't have to.

The simplest account of this contrast would hold that with (367) we are simply reporting the modal content of the rumor's claim. And, as with any speech report, we are free to distance ourselves from the reported content, though not from the report itself, without fear of backtracking or contradiction. With (368) on its non-reportative reading, in contrast, we are saying that the fact that the rumor exists necessitates, whether deontically or more likely teleologically, that Roger go to confession. Maybe we believe the content of the rumor to be false. But given that the rumor even exists, and given the current political climate, if Roger wants to save his career, confession is the only way to go.<sup>100</sup> And of course, once we ourselves have made this modal claim, contradiction or backtracking are the only way to distance ourselves from it. The contrast between (368) and (367) is simply the contrast between making a modal claim and reporting one.

This account of the contrast between (367) and (368) makes no reference to the rumor as a piece of evidence. And the simplest account of the similarity between this contrast and the contrast between (351) and (352) is that no such reference is needed there either. With (351), for example, we could simply be reporting the modal content of the rumor's claim. And with (352), on its non-reportative reading, we could be saying that the fact that the rumor exists entails that Roger has been elected chief.

---

<sup>100</sup> Alexander Williams has pointed out to me an alternative non-reportative reading: namely, that Roger must go to confession, given that the content of the rumor is *true*. This sort of reading is also available in the allegedly epistemic case from Section 6.2.2.2.

Maybe the content of the rumor has nothing to do with Roger directly. But given that a rumor with its content would only arise if certain policies had been enacted—and given that Roger was the only candidate who would enact those policies—then Roger must be the chief if that rumor exists. In the first case, we are reporting a modal claim; in the second, we are making one. And the distinction between reporting and claiming suffices to explain why we can distance ourselves in the one case from the modal content without contradiction or backtracking but cannot do so in the other. We do not need to reference content-bearing pieces of evidence from which either “factual” or “contentful” modes of domain projection are alleged to be possible. And, since it would be simpler and more uniform *not* to do so, I would also suggest that we should not.<sup>101</sup>

When it comes to modal contradictions, then, both the similarities and the contrasts between allegedly epistemic modals and acknowledged non-epistemic ones suggest that the meaning difference between the two does not depend on an epistemic versus non-epistemic characterization of the relevant facts used to determine the modal base.

#### 6.4 On the theoretical motivation for Kratzer’s analysis

The contrast from Section 6.3.1 provides the intuitive motivation for the Kratzerian analysis of the meaning difference between *might* and *can*. Yet on closer inspection, there is also something remarkably unintuitive about that analysis. Kratzer claims, for example, that we sometimes “have to neglect certain facts [in making a *can*-claim], although we might be aware of them” (1981, 303). But there is no intuitive sense in which we have to ignore or neglect the known fact that hydrangeas are not growing here in order to claim that they can. For this claim seems to involve the attribution of a potential, and potentials exist independently of their realization. Hence, it is not at all intuitive that (or why) we should have to ignore the known fact that a potential is not being realized in order to say that it exists, contrary to what Kratzer implies.

Why then does Kratzer claim that we sometimes must ignore known facts when making *can*-claims? In Section 6.4.1, I am going to argue that Kratzer’s claim is about what is required from a formal standpoint, not an intuitive one. Yet the formal requirement is not absolute. Instead, it depends on formal choices that I argue in Section 6.4.2 are motivated by, though ultimately untenable with, the Kratzerian non-ambiguity project. In Section 6.4.3, I conclude by showing how that project fares without the intensional assumption that it motivates.

##### 6.4.1 The formal framework

The possible worlds framework provides a useful way to label the different flavors of modality expressed by *might* and *can*. We can say that the former quantifies over *alethically* accessible worlds, for example, and the latter over *dynamically* accessible ones.<sup>102</sup> Yet the framework itself does not commit us to any particular analysis of

---

<sup>101</sup> My guess is that Kratzer does not consider the alternative presented here because she is working under the assumption, argued against in the previous chapter, that “epistemics” do not contribute to the embedded content of attitude and indirect speech reports.

<sup>102</sup> In discussing the general possible worlds framework, not just Kratzer’s specific implementation of it, I will adopt standard practice and use terms like *accessible worlds* and *accessibility relation*

what makes one world alethically but not dynamically accessible, or vice versa. Instead, accessibility relations are black boxes, placeholders for analyses. Now it is possible that the various types of accessibility relations will all have the same sort of analysis: that there will be a single template, with variables to be filled in, that will suffice to generate these types and, correspondingly, their modal flavors. It is equally possible, however, that various accessibility relations and modal flavors will require their own separate analyses. The possible worlds framework is compatible with either view.

Kratzer, it should be clear, adopts the template view of accessibility relations. She holds that what makes a world accessible in one but not another way is if it instantiates a set of propositions characterized in one but not the other way. She also holds, for all the different types of accessibility relations, that they place the same sort of constraint on accessible worlds: namely, that a proposition be true in all or some of those worlds, depending on the force of the modal.

This template view has ramifications for how Kratzer models claims about potentials, or what she calls pure circumstantials. The potential for hydrangeas to grow here, for example, will be modeled by having it be true of at least one of the accessible worlds that hydrangeas *are* growing here. And these accessible worlds will just be ones that instantiate some relevant set of facts. Given these choices, a problem arises if one of the relevant facts is that hydrangeas are *not* growing here. For all the accessible worlds will instantiate that fact, and so at none of those worlds will it be true that hydrangeas are growing here, which is what the attribution of the potential requires. Hence, the relevant facts need to exclude that hydrangeas are not growing here. From a formal standpoint, we need somehow to be able to “ignore” this known fact when making the *can*-claims above.<sup>103</sup>

For Kratzer, the way to formally ignore the known fact that hydrangeas are not growing here is to look for some way of characterizing the facts that ensures that it does not count as relevant. Yet while she rules out the epistemic characterization that would *not* have this effect, she struggles to provide a positive characterization that would. In her early work, she speaks vaguely of “circumstances of a certain kind” and of “certain sorts of facts” (1981, 302; 1991, 646). In recent work, she is sometimes more specific, speaking, for example, of “internal and external circumstances of people, things, [and] places that determine their possible futures” (2012, 54). Yet it is not clear that the second half of this description does not beg the question. And it is not clear that the first half, by itself, would allow us to formally ignore precisely those circumstances that we need to. (If I am sitting, for example, but can also stand, do my relevant internal and external circumstances include that I am sitting?)

---

slightly differently from how I have so far. For example, I will speak of deontically accessible worlds rather than deontically-ordered, circumstantially-accessible ones. My reasons for doing so will soon become clear.

<sup>103</sup> A alternative response to modeling potentials in this way is to hold that attributions of potentials are claims about what is true at some accessible world at some *future* point (cf. Thomas 2014; Rullmann and Matthewson 2018). Then the fact that hydrangeas are not *currently* growing here does not have to be ignored but can instead be held constant across all accessible worlds. I address this view, which supposedly helps explain the temporal interpretation of modals, in the following chapter.

My point here, however, is not to catalog any difficulties that Kratzer may or may not have in formally ignoring known facts. Instead, it is to show there are alternatives to the template view, and so alternative ways to formally ignore known facts too.

To help see this, consider the difference between the modal base as defined in the rest of the dissertation and as Kratzer has defined it. The first, *ModalBase*<sub>1</sub>, is the set of facts relative to which the content of the modal claim is evaluated for truth. The second, *ModalBase*<sub>2</sub>, is the set of worlds that the modal quantifies over. Now Kratzer has assumed there is a straightforward relation between *ModalBase*<sub>1</sub> and *ModalBase*<sub>2</sub>. *ModalBase*<sub>1</sub> determines a set of worlds—namely, the set of worlds which verify all of its members—and this is identical to *ModalBase*<sub>2</sub>. Yet maybe we should not adopt a single mechanism for moving from *ModalBase*<sub>1</sub> (the set of facts relative to which we evaluate the content of a modal claim for truth) to *ModalBase*<sub>2</sub> (the set of worlds that we formally use to model the truth conditions for that claim). After all, insofar as the non-realization of a potential is no indication of its non-existence, it seems it could be true that hydrangeas can grow here, even relative to the fact that they are not. If so, then maybe instead of trying to formally ignore known facts by utilizing a single mechanism for moving from *ModalBase*<sub>1</sub> to *ModalBase*<sub>2</sub>, while appealing to different characterizations of the facts, we could instead attempt to do so by utilizing different mechanisms while appealing to the same intensional characterizations (or eschewing such characterizations altogether).

A different possibility is to allow modals to place different sorts of constraints on different sorts of accessible worlds. With alethically accessible worlds, for example, we could assume the traditional constraint: that a specific proposition be true at some or all of those worlds. With dynamically accessible worlds, however, we could perhaps instead require that a specific sort of *event* occur at some or all of them. Now obviously, if an event occurs at a possible world, there will be a corresponding proposition that is true of that world. Yet the event constraint is crucially more general than the propositional one. While the latter involves the temporal structure of the world, the former need not. And this could be one possible way to formally ignore known facts. If it is true in all accessible worlds that hydrangeas are not growing here, there could still be hydrangea-growing events in those worlds. An event constraint on possible worlds, as opposed to a propositional one, thus provides another way to formally ignore known facts: i.e., to ensure that they do not have unintended formal consequences.

There are likely to still be other ways to formally ignore known facts as necessary. What is most important, however, is to recognize that this necessity is *merely* formal. It is not conceptual. And even within the possible worlds framework there are formal choices available to us. There is no necessity, and no *a priori* reason either, to adopt the template view of accessibility relations. It is thus part of a substantive analysis within that framework to do so, and to hold, as Kratzer does, that the intensional characterization of propositions and facts is what differentiates one modal flavor from another.<sup>104</sup> In this following section, I will thus explore why I think it is that Kratzer

---

<sup>104</sup> Or rather, it is part of a substantive analysis provided that the intensional characterizations are not question-begging ones. For example, if we hold that *can* requires compatibility with the relevant facts that bear on the existence of a potential, and *might* with those that bear on the truth of a proposition, then we will have begged the question whether they express (i) the same relationship to

adopts the template view. And I will argue that her reason for doing so is ultimately ill-served by her intensional assumption.

#### 6.4.2 The intensional assumption

Kratzer's reason for making the intensional assumption relates to her attempt to account for the phenomenon of *modal flexibility*, or the ability of a single modal to be used to express multiple flavors of modality (Kratzer 1977). This phenomenon turns out to be fairly widespread, both across languages *and* language groups (cf. Hacquard 2010). Yet Kratzer notes that with the template view of accessibility relations we can avoid encoding a modal's flavor into its semantic entry, and so can avoid having to posit multiple semantic entries for the multiple flavors of modality that flexible modals can be used to express. To avoid such ambiguity, all that is required is that the variable positions in the template for accessibility relations be able to be filled in either by context or by syntax (cf., respectively, Kratzer 1981; 1991; and Kratzer 2012; 2013). And a natural way for context and syntax to do this is to make available intensionally characterized sets of propositions, the differences in characterization generating the different types of accessibility relations.

The intensional assumption is thus natural given Kratzer's non-ambiguity project. Yet whether it is also required for that project is another matter. If so, the project should be abandoned, independently of any argument in this dissertation (cf. Nauze 2008; Viebahn and Vetter 2016). And if not, then the intensional assumption should be, so that Kratzer's project itself does not have to. Let me explain.

Kratzer, as already seen from Section 6.2.2.1, proposes two models for generating different types of accessibility relations and modal flavors.<sup>105</sup> The first relies on context alone (Kratzer 1981; 1991). Different types of accessibility relations are generated by combinations of different types of conversational backgrounds (i.e., by contextually-provided functions from possible worlds to sets of propositions). One of these conversational backgrounds will determine the Kratzerian modal base, and the other will function as an ordering source for the worlds in that base. The second model supplements context with syntax (cf. Hacquard 2006; Kratzer 2012; 2013). Adapting the work of Hacquard, Kratzer no longer uses conversational backgrounds to generate modal bases. Instead, she relies on an alleged difference in the syntactic height of modals, and in the semantic objects available at those heights, to do so. This difference had always been correlated with her different types of modal bases, and so presumably could replace the different types of backgrounds previously used to generate them. In combination with the different types of conversational backgrounds functioning as ordering sources, this height difference is now supposed to generate the different types of accessibility relations and modal flavors (2012, 49–55).

---

facts characterized in different ways or instead (ii) different relationships to potentially the same set of facts. Kratzer's characterizations do not generally seem to be question-begging in this way. However, they are also incomplete, as pointed out earlier, and it is not immediately clear how to complete them without begging the question either.

<sup>105</sup> In this section, I will again follow standard practice and use terms like *accessibility relation* in a broader sense than the Kratzerian one used in Section 6.2.2.1. A change in how the latter sort of relation is generated will thus imply a change in how the former is, but not necessarily vice versa.

Now the first of these two models Kratzer rejects because she ultimately rejects her earlier distinction between the relevant conversational backgrounds used to generate modal bases: namely, circumstantial and epistemic backgrounds (cf. Nauze 2008). These are supposed to differ in whether their values are propositions construed as the circumstances or instead as the content of a knowledge state. However, no proposition construed as a circumstance could not also, in principle (even if not in practice), be construed as the content of a knowledge state, and vice versa. Kratzer concludes from this fact that the distinction between types of conversational backgrounds must be rejected, nothing allowing us “to single out some . . . functions as epistemic, but not circumstantial, or the other way around” (2012, 24). The functions themselves are not antecedently of either type, and so by themselves they do not allow us to differentiate modal flavors either. If we have already done the latter, then we can label such a function and use it to represent that flavor. But we cannot hope to go in the opposite direction (for further discussion, cf. Chapter 3, Section 3.3.3).

Kratzer does not reject her distinction between types of conversational backgrounds functioning as ordering sources, but she should by this same reasoning (Nauze 2008).<sup>106</sup> *Deontic* and *bouletic* conversational backgrounds, for example, are supposed to differ in whether their values are propositions construed as the content of a body of laws or instead of a body of desires. However, no proposition construed as the content of the law—e.g., *no murder occurs*—could not also, in principle (even if not in practice), be construed as the content of a desire. And vice versa. Hence, if we cannot appeal to different types of conversational backgrounds to determine different types of modal bases, neither can we appeal to different types of conversational backgrounds to induce different types of orderings on the worlds in those bases. Conversational backgrounds can play no role in determining the type of accessibility relation (and so in determining modal flavor), whether they are supposed to have generated the modal base or to have induced an ordering on it. Kratzer’s second model should thus also be rejected alongside her first.

Problematically for Kratzer, however, only for modal flavors differentiated in terms of her modal bases is there evidence of different syntactic heights (cf. Hacquard 2010). For modal flavors differentiated exclusively in terms of her ordering sources—deontic and bouletic flavors, for example—there is none. Some authors have even proposed instead that for specific modal flavors there may be flavor-internal height differences (cf. Brennan 1993 on ought-to-be vs. ought-to-do deontics), while others have proposed more general flavor-internal height differences correlated with modal force (cf. Cinque 1999 on possibility vs. necessity modals).

If these authors are correct, then there are two problems facing us if we attempt to rely on syntax alone to generate different types of accessibility relations. First, we will require relevantly different types of semantic objects at the same syntactic height for any modal flavors not differentiated height-wise. Second, we will sometimes

---

<sup>106</sup> Nauze (2008)’s observations about the conversational backgrounds used to determine modal bases are often cited in support of rejecting the distinction between epistemic and circumstantial backgrounds (cf. Kratzer 2012; Matthewson 2016). Oddly enough, his equivalent observations about conversational backgrounds functioning as ordering sources are never addressed, as far as I can tell. But the observations point to the same conclusion. If we accept this conclusion in the one case, we need some argument for why we should not accept it in the other. I have seen none given.

require relevantly the same types of semantic objects at different heights to accommodate any flavor-internal height differences.<sup>107</sup> Kratzer's second model thus cannot be easily revised to generate all the different types of accessibility relations, and hence modal flavors, that we require. This suggests that something has gone wrong.

What has gone wrong is that while context may provide sets of propositions, or the values of conversational backgrounds, it does not intensionally characterize them (Nauze 2008). Hence, if this is what differentiates one type of accessibility relation and modal flavor from another, it will have to be provided by a modal's semantic entry, multiple entries being posited for modals used to express multiple flavors of modality. That said, once we have hardwired a modal's flavor into its semantic entry, there is no longer any reason to adopt the template view of accessibility relations or to think that all flavor differences must be differences in intensional characterization of the sort that Kratzer assumed she required for her non-ambiguity project to succeed.

The intensional assumption may be motivated by the Kratzerian non-ambiguity project, then, but it turns out to be incompatible with it. Whether we give up on the project of accounting for modal flexibility without positing modal ambiguity, or instead attempt to make it work without the intensional assumption, either way the result is the same. We should simply give up on the intensional assumption itself. Instead of holding that *might* and *can* differ with regard to the facts that they express modal relations to, we should consider the distinct possibility that they differ in the relations they express to those facts.

### 6.4.3 Non-ambiguity without the intensional assumption

Despite his criticism of its reliance on intensionally characterized conversational backgrounds, Nauze seems to suggest that Kratzer's non-ambiguity project may still be viable. In advocating an ambiguity framework, for example, Nauze writes that "obviously we need within such a framework a way to resolve the ambiguity," and he continues by claiming that "this ambiguity is often resolved by context" (2008, 153–54). Yet this is not a theoretical disadvantage of his framework compared to the non-ambiguity one, Nauze avers. Instead, he takes his argument to have shown precisely that, even within the Kratzerian framework, "a new contextual parameter [over and above conversational backgrounds] is needed to determine the [flavor] of modality involved in a particular utterance" (2008, 157–58). And Nauze asserts that this "same parameter could be used in [an ambiguity] framework ... for the same purpose" (2008, 158). Nauze thus seems to be conceding that the non-ambiguity project can be made to work, just not in its standardly assumed form.

Unfortunately, Nauze obscures what is at stake in choosing between ambiguity and non-ambiguity frameworks with these comments. Context simply does not play the same role in these two theories, and there is no contextual parameter that is needed for the ambiguity framework. To suggest that there is is equivalent to suggesting that there is a contextual parameter for the numeric subscripts that we use to represent the different senses of a word like *bank*. When an ambiguous word like *bank* is used, various contextual pressures may well lead us to prefer one sense over another. But

---

<sup>107</sup> An additional concern is whether we should hope to rely on a single or instead multiple syntactically-provided objects to generate different types of accessibility relations.

this is not a matter of accessing an underspecified semantic entry, for which context somehow supplies the relevant missing numeric subscript. Rather, it is a matter of accessing one fully specified semantic entry, which we represent with a subscript, as opposed to another. This is in contrast to what happens when we interpret a third-person pronoun like *he* in context. Here there is a single, underspecified semantic entry for which context *does* need to supply the referent. And the choice between ambiguity and non-ambiguity frameworks for flexible modals is the choice between interpreting such modals as analogous to *bank* or instead as analogous to a pronoun like *he*. Contrary to what Nauze suggests, there is not some contextual parameter to be used for the same flavor-specifying purposes in the two frameworks.

I am, on the whole, thus somewhat less sanguine than Nauze about the prospects of the Kratzerian non-ambiguity project without the intensional assumption. In fact, as far as I can tell, the best hope for that project is to abandon the template view of accessibility relations, and the search for appropriate objects to fill into that template to generate the flavor differences. Instead, we should hold that different sorts of accessibility relations could sometimes have fundamentally different sorts of analyses: that is, that modal flavors could involve different sorts of modal relations, not just the same modal relations to different sorts of propositions. Doing so, we could then posit a semantic entry, specified for force, with a variable ranging directly over these different sorts of accessibility relations themselves (rather than over entities intended to generate them). This would be analogous to what Recanati (2001) proposes for possessive phrases like *John's car*, where he suggests that there may be a free variable over the various relations that John may bear to this car (e.g., John may own the car, or have placed a bet on it, or be its designer, or have a particular fondness for it, etc.) (cf. Recanati 2001, 85; for a response, cf. Bach 2001, 39). Just as these relations will have fundamentally different sorts of analyses—and not a template, with variables to be filled in by context, that is capable of generating them one and all—so the same could hold for accessibility relations within the possible worlds framework.

Whether we should adopt a Recanati-style variable over accessibility relations is not the concern of this chapter, however. If the non-ambiguity project is compatible with my fundamental contention that *might* and *can* express fundamentally different sorts of relations, and not just the same relation to fundamentally different sorts of facts, then whether to do so depends on the strength of the arguments for the non-ambiguity project itself. Now I am not convinced the standard arguments for this project are all that compelling, as I will argue in the next section. However, the most important conclusion of this section is that the template view of accessibility relations and the intensional assumption, though initially motivated by Kratzer's non-ambiguity project, ultimately do not serve that project well. Any investment we have in this project should thus not lead us to favor a fact-based analysis of the difference between *might* and *can* over a relational one. Very likely, it should lead us to do quite the opposite.

## 6.5 Revisiting the arguments for Kratzer's non-ambiguity project

Kratzer's non-ambiguity project is not my primary concern in this chapter. That said, this project has been the driving force in the formal semantics for natural language

modals since its inception. Hence, it is worth addressing this project directly to see how well-motivated it is on the whole. I do this in the following sections, evaluating the individual arguments that have been offered on its behalf.

### 6.5.1 The explosion of meanings argument

In a recent paper, Viebahn and Vetter (2016) catalogue the arguments for Kratzer's non-ambiguity project. One of these they call the *explosion-of-meanings argument* (the term is due to Schaffer 2011). This is the argument from Kratzer (1977) that once we begin to multiply modal meanings to account for modal flexibility, there is no natural point at which to cease doing so before implausibly many have been posited.

We might, for example, want to hold that there is a *must*<sub>deontic</sub> which occurs in (369), and a *must*<sub>epistemic</sub> which occurs in (370).

(369) All Maori children must learn the names of their ancestors. [Kratzer's (2)]

(370) The ancestors of the Maoris must have arrived from Tahiti. [Kratzer's (3)]

However, (369) could be used to make a claim about regional social customs or instead about national laws. So we should really posit a *must*<sub>deontic-social-customs</sub> and a *must*<sub>deontic-national-laws</sub>, and specify which it is that occurs in (369). And so on. Once we have started on this road, there is no natural terminus, no non-arbitrary stopping point. Or so the argument goes.

Since it is implausible to think that a speaker will have to have access to so many meanings to be able to interpret a modal claim, however, Kratzer proposes her now-famous alternative. She argues that what is common to the meaning of the sub-types of deontic *must* can be isolated by making explicit in just what regard they differ, via *in view of* phrases like those in (371) and (372).

(371) In view of the Maori social customs, all Maori children must learn the names of their ancestors.

(372) In view of the laws of New Zealand, all Maori children must learn the names of their ancestors.

That is, we can posit a single deontic *must* with a meaning deficiency to be filled in by the contextually relevant norms. But then, Kratzer continues, we can do the same not just within the different flavors of modality, but also between them. We can isolate what is common to the epistemic and deontic readings of (373), for example, by the use of *in view of* phrases, like those in (374) and (375), that again supposedly specify the only regard in which the two readings differ.

(373) John must hate his enemies.

(374) In view of what is known, John must hate his enemies.

(375) In view of his party's bylaws, John must hate his enemies.

Kratzer concludes from this that there is just a single, neutral *must* that needs its meaning deficiency to be supplied by some contextually relevant entity: a body of evidence, a collection of norms, etc.

What is crucial to this first argument is the assumption that there is no principled distinction to be drawn between flavor-internal and flavor-external differences in apparent meaning. But this assumption is not well-founded. For example, while (376) can be used to report beliefs about deontic necessities relative to different norms, (377) *cannot* be used to report both a belief about a deontic necessity and a belief about an epistemic one.<sup>108</sup>

- (376) The village leaders think that all Maori children must learn the names of their ancestors, and so do the national legislators.
- (377) The village leaders think that all Maori children must learn the names of their ancestors, and so do the anthropologists.

This would suggest that while flavor-internal differences may be contextual, and not a part of content, flavor-external ones may be lexicalized. Indeed, while an appropriate paraphrase of (369) in both (371) and (372) is (378), for (374) and (375) different paraphrases of (373) are appropriate: namely, (379) and (380).

- (378) It must come about that all Maori children learn the names of their ancestors.
- (379) It must be true that John hates his enemies.
- (380) It must come about that John hates his enemies.

It is thus plausible that *X*'s value in the *in view of X* phrases exhausts the meaning differences between (371) and (372) without its also doing so for (374) and (375). As Palmer (1979) would put things, (374) and (379) could involve “the modality of propositions”, while (375) and (380) involve “the modality of events” (1979, 35).

Now it may be possible, as Palmer suggests, to “argue that [the modality] of propositions and [the modality] of events can be subsumed under a more general notion of [modality]” (1979, 35). But if the two modalities in (379) and (380) *are* to be subsumed, this will require factoring out some further difference between the two that may be reflected in, but is certainly not exhausted by, the difference between a body of evidence and a collection of norms. Thus, if we do try to subsume the two modalities within a single semantic entry, we will have to posit distinct contextual parameters to account for the flavor-internal differences and the flavor-external ones, as Nauze suggests. And this is consistent with there being a principled distinction between the two, such that positing an ambiguity to account for the latter need not lead us to do the same for the former. We do not face the prospect of an explosion of meanings if we encode modal flavor directly into a semantic entry instead of having a parameter for flavor-external differences or Recanati-style variables over accessibility relations (cf. also Viebahn and Vetter 2016).

### 6.5.2 The common kernel of meaning argument

The second argument that Viebahn and Vetter consider they call the *common-kernel-of-meaning argument* (a term due to Kratzer 1977). This argument, also from Kratzer

---

<sup>108</sup> This style of example is due to Alexander Williams, p.c.

(1977), closely relates to the first. It is that there is “kernel of meaning which seems to stay invariable whenever the verb *must* is used” (1977, 341). And this kernel of meaning should be represented in a single semantic entry, the argument goes, with all that is variable being made the value of a contextual parameter (or parameters).

Viebahn and Vetter’s response to this argument is that it fails to distinguish *polysemy* (a form of ambiguity in which there are distinct but related meanings) from *context-sensitivity* (a case in which a single, underspecified meaning can be used in different contexts to refer to different things). Viebahn and Vetter use adjectives like *healthy* and *long* to illustrate this distinction. *Healthy*, for example, may be used in different contexts to refer to different sorts of individuals, depending on the relevant reference class in those contexts. (John may be considered healthy for a hospital worker but not for a member of the general public.) However, it may also be used to refer to different sorts of entities altogether: to diets, lifestyles, or environments. The former is arguably a case of context-sensitivity: of a single meaning underspecified with regard to a reference class, which has to be supplied by context. The latter is arguably a case of polysemy: of related but distinct meanings, some of which are even defined in terms of the others. (A healthy diet, for example, is defined in terms of its conduciveness to being, or to becoming, a healthy individual.)

Now the correct analysis of polysemy is a matter of much dispute, and it is not my intent to enter into that debate here (cf. Nunberg 1979; Pustejovsky and Boguraev 1996; Ravin and Leacock 2000; Asher 2011; Pietroski 2005; 2018). What matters for our purposes is that not all commonalities of meaning need be represented by a common, context-sensitive meaning. Of course, we could, if we wanted, posit a single, doubly-underspecified meaning for a term like *healthy*: e.g., *relating to a well-functioning individual, via some relation R and relative to some reference class C*. Yet we could also just as well posit multiple, distinct meanings that specify this relation *R* while leaving the reference class *C* unspecified. In fact, given that the possible values of *R* will be highly constrained, while the possible values of *C* will not, this second route is likely to be the better one. Likely what we are dealing with is a constrained class of distinct but related meanings, rather than a single, underspecified meaning admitting only a highly constrained range of contextual values (cf. Viebahn and Vetter 2016, 8).

Something similar holds for other terms too, including nouns, like *observation*, that demonstrate what is known as the product/process ambiguity. *Observation*, for example, can be used to refer either to the process of observing or to its product, and there is obviously a kernel of meaning common to both these uses. For each noun like *observation*, then, we could try to abstract away the differences between process and product and to represent their commonalities via individualized, context-sensitive meanings. Yet there is also no especially good reason to hold that there *must* be a process/product parameter for which the context supplies one of two possible values whenever one of these nouns is used. In fact, it seems we do better to appeal to the distinct but related meanings that we can already easily specify rather than to factor out the differences between the two as the value of so highly constrained a parameter.

As with *healthy* and *observation*, so with *must*, at least when it comes to the modal flavors. *Must* could have meanings that are fully specified for flavor, even if they are underspecified in other, flavor-internal ways. This, in fact, is the entire argument of

Viebahn and Vetter (2016). The common kernel of meaning displayed when a modal like *must* is used to make claims of different flavors does not require there to be a context-sensitive meaning common to those flavors.

### 6.5.3 The cross-linguistic argument

The third and final argument that Viebahn and Vetter consider they call the *cross-linguistic argument*. Attributed to Hacquard (2010), this is the argument that modal flexibility is a widespread, cross-linguistic phenomenon, contrary to what we expect if flexible modals have multiple meanings. Why, after all, should “the same lexical accident be found in language after language?” (Hacquard 2010, 80).

Yet while it would be a lexical accident for the riverside/financial institution ambiguity evidenced with *bank* to occur in language after language, it would not be so for the process/product ambiguity evidenced with *observation*. Such a cross-linguistic phenomenon would presumably instead reveal something about how we conceptualize the world, as it could in the modal case too. Hacquard, in fact, is aware of this possibility, citing the same authors as Viebahn and Vetter do in response to her rhetorical question about lexical accidents (cf., a.o., Sweetser 1990; Traugott 1989). However, Viebahn and Vetter have not completely represented Hacquard’s concern.

Hacquard not only emphasizes that modal flexibility is a cross-linguistically systematic phenomenon, but also claims that certain flavors expressed with flexible modals correlate just as systematically with height in the syntactic tree. Hacquard is thus more than willing to grant of Sweetser’s proposal that it can explain why what she calls epistemic and root (or non-epistemic) flavors of modality “share a common form” (2010, 91). But crucially, she maintains, this proposal also “leaves unexplained why roots and epistemics systematically end up with the scopal properties they do” (2010, 91).

Here, I think, Hacquard is wrong. In fact, in the following chapter I will propose semantic entries for the different flavors of modality that would directly result in their having different scopal properties. For now, consider again the noun *observation*. Depending on if it is used to describe the process of observing, as in (381), or instead the product of such a process, as in (382), it may take complements of different sorts.

- (381) John’s observation of the man playing the violin calmed him.
- (382) John’s observation that the man was playing Liszt enraged him.

Yet there is a principled reason for this, insofar as what is observed in a process sense will be an ongoing event or state of affairs, while what is observed in a product sense will be a proposition. Given the relation between process and product, there is no real mystery for *observation* taking different sorts of complements depending on the sense in which it is used. We can posit an ambiguity in this case that is not only systematic and understandable but also results in different complement types being associated with different resolutions of the ambiguity.

The same could hold for modals like *must*, I would argue, if they can also be used to describe either the modality of propositions or the modality of events. There will no doubt be a close relation between these modalities that explains their being expressed with a common form. But if this common linguistic form does indeed take

different sorts of complements and occur at different sorts of heights, depending on the sense of the modal, this could well be attributed to the different nature of event and propositional modality. Apparent correlation between modal flavor and syntactic height could point as naturally to different semantic entries for the modal flavors as to a common, context-sensitive semantic entry that abstracts away from flavor.

I thus conclude that none of the arguments for the non-ambiguity project compel us to adopt it. If there are ever phenomena that seem to require us to abandon this project, as we will see in the following chapter, we should not be afraid to do so.

## 6.6 Conclusion

On the standard analysis, *might* and *can* differ in whether they express compatibility with an epistemically characterized set of facts, in the case of *might*, or instead a non-epistemically characterized one, in the case of *can*. This analysis is incompatible with the argument of this dissertation that *might* does not have an epistemic interpretation. However, I have argued that the motivations for this analysis turn out to be ill-served by it. This analysis does not make good sense of the contrasting behavior of *might* and *can* when it comes to modal contradictions, the initial intuitive motivation for the analysis. Nor does it turn out to be compatible with the Kratzerian non-ambiguity project, its initial theoretical motivation. There is thus no obstacle to rejecting the standard analysis of the difference between *might* and *can* and accepting the argument of the previous chapters that *might* is never interpreted epistemically.

## Chapter Seven: The Temporal Interpretation of *Might* and *Can*

Modals with a circumstantial modal base are restricted to future [temporal orientation]; epistemic modals are free to have any [temporal orientation] (past, present, or future).

~ Rullmann and Matthewson (2018, 283)

### 7.1 Aspectual differences between *might* and *can*

*Might* and *can* differ not only in their meaning—the topic of the previous chapter—but also, apparently, in their temporal properties. *Might*, for example, freely accepts progressive and perfect complements, as (383) shows. And yet, as (384) shows, *can* does not do the same (cf. Jenkins 1972; as cited in Ney 1978, 38).

(383) John might {speak/be speaking/have spoken} French.

(384) John can {speak/#be speaking/#have spoken} French.

Now the standard analysis of this difference, like the standard analysis of the previous chapter, is in terms of the difference between epistemic and circumstantial modality (cf. Matthewson 2012; Rullmann and Matthewson 2018; Thomas 2014; 2017; Klecha 2016). In this concluding chapter, I thus again have two aims.

First, I want to show that the standard analysis is likewise problematic in this case, independently of any argument that *might* expresses non-epistemic modality. Second, I want to show how a better account of the data in (383) and (384) provides the tools we need to formalize the proposal from the previous chapter about how *might* and *can* differ in meaning if not along the epistemic/non-epistemic dimension. Here is the structure my argument in this chapter will take.

In Section 7.2, I develop the standard analysis, which accounts for the restrictions on *can* in (384) in terms of *temporal orientation*. As defined by Condoravdi (2002), this is the evaluation time of the prejacent relative to the modal's own evaluation time. The progressive and the perfect seem to enforce present and past orientation, respectively, yet it has been argued that future orientation is obligatory, not only for the so-called pure circumstantial modalities that concern us here, but also for all circumstantial modalities (cf. Werner 2006). Whether for semantic reasons or for pragmatic ones, circumstantial modal bases are claimed to generally be incompatible with present and past orientation (for the stronger claim, cf. Klecha 2016; Rullmann and Matthewson 2018; for the weaker claim, cf. Thomas 2014; 2017; Matthewson 2012).<sup>109</sup> Present and past orientation, as signaled by the progressive and the perfect, are thus indicators of epistemic modality.

In Section 7.3, I then argue that this analysis fails on all accounts. First, with the appropriate staging, *can* will accept progressive and perfect complements, even on its pure circumstantial readings. Second, bare complements cannot always be analyzed

---

<sup>109</sup> In this chapter, I will be using the term *modal base* in yet a third way, more common amongst the authors with whom I will be engaging here: namely, to refer to functions to relevant sets of facts, rather than to those sets themselves (as in Chapters 3, 4, and 5), or to the sets of worlds determined by them (as in Chapter 6). Unless explicitly noted otherwise, I avoid these other two uses here.

as being future-oriented, neither for pure circumstantials, nor for other circumstantial modalities either. Third, obligatory future orientation could not plausibly be attributed to features of the modal base anyway. An alternative analysis is thus required.

In Sections 7.4 and 7.5, I develop my alternative to the standard analysis. I argue that *can* introduces its own event variable, and that it is used to attribute *modal states* (cf. Homer 2010).<sup>110</sup> These are actual states of individuals, locations, situations, etc., but ones that can only be specified with reference to eventualities in possible worlds other than the actual one. *Might* differs from *can* in that it is not used to attribute such states and so does not introduce its own event variable either. Instead, it is used simply to assert that some proposition is possibly true, relative to some set of facts. *Might* is an instance of what I will call a *proposition-level* modal, *can* an instance of what I will call an *event-level* one.

This difference between *might* and *can* helps to explain why the former but not the latter freely accepts the whole range of aspectual morphology. For the role of aspect is to relate an event description to a reference time, thereby helping to determine a so-called *temporal proposition* (or predicate of times). Since *might* is used to attribute modal properties to propositions, it thus makes sense that it would freely accept the whole range of aspectual morphology. Yet for the event descriptions used to specify modal states, there is often no need to impose the particular relations to particular reference times that the progressive or the perfect would. Only when there is the need to describe the potential to do two things either simultaneously or sequentially, for example, will it make sense to use the progressive or the perfect. And this explains why *can*, unlike *might*, does not freely accept either.

On a traditional view that also posits two levels of modal auxiliaries, modals are distinguished first and foremost in terms of the semantic type of their complement. For example, on their so-called *root* interpretations, modals are generally supposed to combine with predicates of events, and to yield the same, while on their *non-root* interpretations they are supposed to combine with (temporal) propositions, again to yield the same (cf. Brennan 1993, 2–5; Hacquard 2009, 292–93; a.o.).<sup>111</sup> On this view, modals have fundamentally the same semantic entry, type-shifted as necessary. In contrast, on my view, modals have fundamentally different semantic entries, distinguished by their introduction (or non-introduction) of an event variable. As a result, while modals may have inputs of the same semantic type, their outputs will be of different ones.

In Section 7.6, I discuss an implication of this view for the counterfactual readings of non-root modals, which Condoravdi attributed to aspect scoping above the modal. This explanation is unavailable to me, with my reinterpretation of non-root modals as proposition-level ones. For while event- and proposition-level modals do not differ in their ability to scope *above* aspect, on my account, they do differ in their ability to

---

<sup>110</sup> While Homer provides the inspiration for using event variables in this specific way, Hacquard (2006; 2010) provides the more general inspiration for using event variables with modals at all.

<sup>111</sup> The root/non-root distinction was originally the semantic distinction between non-epistemic and epistemic modality (cf. Hofmann 1966). However, a number of syntactic and distributional differences were noted between these semantic groupings, and Condoravdi (2002) then noted that non-root modals expressed not only allegedly epistemic but also presumably metaphysical modality. Condoravdi (2002) is also responsible for the view that the propositions that non-root modals combine with (and yield) are temporal ones.

scope *below* it. Yet I argue that we do better to attribute counterfactual readings to past tense on a modal than to aspect scoping above it. In making this argument, I thus enter into the debate about whether non-root modals scope above tense (cf. Hacquard 2010; 2011; Rullmann and Matthewson 2018; a.o.). The plausibility of the view that they cannot, I maintain, is due solely to the misinterpretation of non-root modality as epistemic modality, when in fact the modality is propositional.

Section 7.7 concludes the chapter, and the dissertation as a whole.

## 7.2 Aspectual restrictions on *can*'s complements

### 7.2.1 The standard analysis of the restrictions

Temporal orientation, as defined by Condoravdi (2002), is the relative evaluation time of a modal's prejacent. Progressive morphology seems to require that evaluation time to be simultaneous with the modal's own evaluation time, or to impose what is called *present orientation*. Perfect morphology seems to require a backward-shifted evaluation time, or to impose *past orientation*. And since *can*, as we have seen, does not initially seem acceptable with complements in either of these forms, a natural first thought is that it requires a forward-shifted evaluation time, or *future orientation*.

Future orientation, like past orientation, is generally represented with existential quantification over times. Sometimes this quantification is attributed to the modal base (cf. Klecha 2016), sometimes to the modal itself (cf. Condoravdi 2002; Werner 2006), and sometimes to a covert aspectual operator scoping under the modal (cf. Kratzer 2011; Matthewson 2012; Rullmann and Matthewson 2018; Thomas 2014; 2017). However the quantification is effected, the claim made with (385) can be glossed with (386), and given the generic logical form in (387), on the view that it has obligatory future orientation.

(385) John can speak French.

(386) At some future time in some world compatible with our present circumstances, John speaks French at that time in that world.

(387)  $[\llbracket \diamond p \rrbracket^{w,t,f} = 1 \text{ iff } \exists w' \exists t' [w' \in \cap f(w)(t) \ \& \ t < t' \ \& \ \llbracket p \rrbracket_{\epsilon}(w')(t')]$ <sup>112</sup>

This is the standard analysis of what the restrictions on *can* with regard to progressive and perfect morphology involve: namely, obligatory future orientation, or obligatory existential quantification over future times.

### 7.2.2 The standard explanation(s) for the restrictions

Why should modals like *can*, which clearly express circumstantially based modalities, require future orientation? Here a number of different answers are given by authors in the literature, though all ultimately in terms of the modal base.

Some authors propose that circumstantial modal bases are *totally realistic*: they assign to a world-time pair the total circumstances of that world up to that time (cf. Werner 2006; Laca 2012). This proposal is combined with an interpretive principle—Werner calls it the *Disparity Principle*—that requires the set of worlds determined by

<sup>112</sup> For the sake of simplicity, I will bracket in this chapter the question from Chapter 3, Section 3.3.2 whether modal relations are best analyzed in terms of the prejacent's truth at a possible world.

the value of the modal base to be diverse, or heterogenous, with regard to the prejacent's truth value at its evaluation time. In essence, in that set of worlds, there should be both ones in which the prejacent is true at its own evaluation time and ones in which it is false. Analogues of the Disparity Principle are Condoravdi (2002)'s *Diversity Condition* and Thomas (2017)'s *Modal Economy*.<sup>113</sup>

The basic idea behind these principles is to avoid the equivalence of modal and non-modal claims. Imagine, for example, that the world was deterministic. Then some event could happen if and only if it did happen, and if it did happen, then it also had to have happened. Non-epistemic possibility and necessity claims would thus be equivalent, in a deterministic world, to the claims made with their prejacent. Such equivalence makes the modals redundant, and interpretive principles like Werner's are designed to avoid that redundancy. Combined with the assumptions that the world is indeterministic, and that circumstantial modal bases are totally realistic, principles like these require the so-called root interpretations of modals—all assumed to have circumstantial modal bases—to also have future orientation.<sup>114</sup> For say the prejacent's evaluation time was either simultaneous to, or backward-shifted from, the modal's own evaluation time. Then whether that prejacent was true at that time would be just another one of the facts determined by the value of the modal base. Hence, diversity would be violated.

Totally realistic modal bases have been criticized, both in their utilization by Werner for all root interpretations of modals (cf. Portner 2009, 234–35) and in their utilization by Condoravdi for non-root, counterfactual ones (cf. Abusch 2012). I will focus on Portner's criticisms here. Portner points out that totally realistic modal bases make it difficult to account for some of the context-sensitivity of modal claims. Whether we agree with the claim made with (385), for example, may depend on whether our focus in context is on John's having learned the language, or instead on his having suffered a stroke from which he may never fully recover. This difference in context seems to be a difference in the facts relative to which we evaluate John as having the potential, or ability, to speak French. That is, it seems to be a difference in the value of the circumstantial modal base. But this difference cannot be captured if we assume that circumstantial modal bases are all totally realistic ones.

To account for this difference, Portner suggests that Werner could perhaps appeal to a difference in the value of the *ordering source*, a set of propositions used to order the worlds determined by the value of the modal base according to their compliance with certain ideals (cf. Kratzer 1981; 1991). Yet what matters in this case seems not to be any difference in the relevant ideals—for ability modals, these could *perhaps* be stereotypes—but instead a difference in the relevant facts. Anyhow, even if this is

---

<sup>113</sup> For skepticism that such principles should be stated in terms of the set of worlds determined by the value of the modal base, rather than in terms of the value of the modal base itself, cf. Chapter 3, Section 3.4.1. Condoravdi's Diversity Condition, and her totally realistic modal bases, which she referred to as *metaphysical*, were not intended for the so-called root modalities expressed by *can*. In fact, she explicitly disavowed making claims about their temporal interpretation. Yet in an odd twist of intellectual history, her analysis of the apparently obligatory future orientation of certain non-root modalities expressed by *might* has played a key role in the analysis of root modalities as having such orientation too. For this history, cf. Abusch (2012), Thomas (2014), and Rullmann and Matthewson (2018), amongst others.

<sup>114</sup> For general criticism of relying on an assumption of indeterminism, cf. Abusch (2012, 282–83).

wrong, Klecha (2016) points out that cases where the ordering source has to be empty also create problems for reliance on totally realistic modal bases (2016, 33).

In lieu of totally realistic modal bases, those who maintain that all non-epistemic interpretations of modals have obligatory future orientation have generally abandoned the attempt to explain why this should be so. Rullmann and Matthewson (2018), for example, write that “it is not our main purpose here to explain the correlation between nonfuture TO [i.e., temporal orientation] and epistemic interpretations. We simply assume ... that some restriction along the lines of the [Diversity Condition] suffices to derive the modal flavor/TO restrictions” (2018, 301). And Klecha (2016), instead of relying on such a principle, simply hardwires future orientation into circumstantial modal bases, independent of their being totally realistic or not.<sup>115</sup>

Some authors maintain, more conservatively, that only some of the non-epistemic interpretations expressed by modals like *can* have future orientation. These authors, as you might expect, have *not* abandoned attempts to explain why future orientation is favored for these interpretations. One of these explanations maintains Werner’s and Condoravdi’s general interpretive principle but limits its scope by abandoning the use of totally realistic modal bases (Thomas 2014). Another explanation converts that interpretive principle into a pragmatic preference but without necessarily limiting its scope (Matthewson 2012). I review each of these explanations in turn.

Thomas suggests that we think of circumstantial modal bases as functions to the relevant facts, as opposed to the total ones. Crucially, what the relevant facts are will vary with the flavor of modality expressed. For the so-called *priority* modalities—deontic, teleological, etc. (cf. Portner 2009)—these should not include the fact of the matter about the prejacent at the modal’s evaluation time. For a course of action does not cease to be impermissible or undesirable just because it is undertaken, nor does it cease to be required or desirable just because it is not. Hence, present orientation for the priority interpretations of modals is compatible with the idea motivating Werner’s and Condoravdi’s interpretive principles. And present orientation does indeed seem to be possible, as Thomas shows with (388) on its deontic interpretation.

(388) Sam ought to be sleeping. [Thomas’ (69)]

Things are supposed to differ for non-priority modalities, which Thomas calls *pure circumstantials*. Here Thomas suggests a principle of self-relevance is at work, so that the relevant facts must always include the facts of the matter about the prejacent’s truth up to, and at, the modal’s evaluation time (2014, 439–40). Given this principle, future orientation is required for pure circumstantials, which Thomas takes to include

---

<sup>115</sup> Oddly enough, he also writes the following:

[S]ince the circumstances in a given world are fully settled up to evaluation time, it makes sense that circumstantial modal expressions would be restricted to the future, which is the only contingent part of a set of circumstantially accessible worlds. This is motivated by the Maxim of Manner, since a circumstantial modal embedding a proposition with a past time reference could always just be replaced with the bare proposition with no change in meaning (2016, 12).

Yet this explanation presupposes the totally realistic view of circumstantial modal bases that Klecha eschews!

modals on their ability interpretations. And this requirement, Thomas argues, is borne out by the facts.

Thomas does add as a caveat, however, that past-tensed ability modals seem to be compatible with present orientation, as in (389).

(389) John was able to speak French to President Macron when he had to meet with him.

Cases like (389) involve so-called actuality entailments. In this case, attributing to John the ability to speak French on a particular occasion is supposed to entail that he *did* speak French on that occasion.<sup>116</sup> Hence, the prejacent's evaluation time would seem to be concurrent with the modal's, something we might expect for present-tense modals used to attribute abilities and potentials that are actively being realized too. Yet Thomas (2014), and then again Thomas (2017), simply saves for future research the problem that actuality entailments pose for his account of temporal orientation.

Matthewson, in contrast to Thomas, takes actuality entailments to be a good reason to convert the general interpretive principle of Werner and Condoravdi into a mere pragmatic preference (2012, 14). Yet as far as I can tell she also takes it to have wider scope, including both Thomas' pure circumstantials and Portner's priority modals. For Matthewson argues that "circumstantial modals *inherently* give rise to actuality entailments, and that [these] are removed by prospective aspect", which imparts future orientation (2012, 11; emphasis mine). Hacquard (2009), to whom Matthewson is responding, also argued that not just pure circumstantials, but also priority modals, give rise to such entailments. This wide scope for the pragmatic preference is further supported by Matthewson's joint paper with Rullmann, already cited as an example of the view that future orientation is obligatory for all circumstantial modalities. And yet wide scope encounters precisely the sort of problems with deontic interpretations that Thomas sought to avoid.

These, in sum, are some of the reasons given for why it is that circumstantially based modalities allegedly require or prefer future orientation. As is evident, there is disagreement not only about the extent to which future orientation is required—whether for all flavors of circumstantial modality or only some—but also about the nature of this requirement itself: whether it is a general interpretive principle, a mere pragmatic preference, or even a hardwired component of the modal base. Despite this disagreement, there is widespread consensus that circumstantially based modalities are subject to certain temporal constraints best understood as constraints on temporal orientation, and due, in some way, to the modal base being circumstantial. And this is supposed to explain why perfect and progressive complements are unacceptable with modals, like *can*, that express such modalities (cf. Thomas 2014, 445–48).

---

<sup>116</sup> Such entailments are apparently both clearer and also uncancellable in languages, unlike English, that have *perfective* aspectual morphology; cf. Hacquard (2009). For competing accounts, cf. Mari and Martin (2007) and Homer (2010). I will have little new to contribute to this debate in this chapter. The account of modal-aspectual interactions that I will propose is compatible with Homer's account. But I am unable to adjudicate the empirical debates between him and Hacquard about whether perfective aspect is sufficient for actuality entailments and whether such entailments arise with modal auxiliaries but not with other grammatical categories expressing modality (cf. Hacquard 2016, 48–49).

### 7.3 Against the standard analysis of *can*'s aspectual restrictions

#### 7.3.1 Progressive and perfect complements

A first problem for the standard analysis of the aspectual restrictions on *can* in (384) is that *can* does not always reject complements in the progressive and perfect forms. While (390) is odd with its progressive morphology, for example, (391) is not.

- (390) #John can be juggling and unicycling.
- (391) John can be juggling and unicycling and still carry on a trenchant conversation about Kant.

Similarly, while (392) is odd with its perfect morphology, (393) is not.

- (392) #John can have just sung the Ave Maria.
- (393) John can have just sung the Ave Maria, attended mass and confession, said five Hail Marys, and still launch into so vulgar and profane an outburst it would make even the most hardened sailor blush.

This point extends beyond *can* to flexible modals like *may*. While the progressive morphology makes a deontic interpretation of (394) very difficult, this interpretation is readily accessible in (395), and may in fact be the only interpretation available.

- (394) The parishioners may be singing.
- (395) The parishioners may be singing as they enter the church but must be silent by the time they are seated.

One possible response to these data is to note that the abilities and permissions attributed seem, in at least one sense, to be generic. They are not limited to a specific occurrence but instead are repeatable, in the case of the abilities, or ongoing, in the case of the permissions. And perhaps genericity itself imposes future orientation, somehow overriding the default past and present orientation imposed by the perfect and progressive. Or perhaps, as Thomas (2017) has suggested, genericity presents the exception to an otherwise obligatory correlation between circumstantially based modalities and future orientation.

Perhaps so. But there are problems with this appeal to genericity. For example, the adjective in (396) lexically encodes deontic modality, has a progressive complement, and can only be understood as granting a specific, one-time permission for the current moment.

- (396) The children are allowed to currently be running.

The modality would thus seem to have present temporal orientation, the prejacents' evaluation time being simultaneous to the modal's own evaluation time. Much the same holds for the semi-modal auxiliary *have to* in (397).

- (397) Bill knows the rules. He has to be doing the dishes right now. [≈ Klecha's (12b)]

Klecha (2016) has argued that (397) can only have an epistemic interpretation, not a deontic one, because of the present orientation imposed by the temporal adverb *now*. Yet (397) is perfectly fine as a response to (398), showing that Klecha is wrong to claim that a deontic interpretation is unavailable.

(398) Is Bill required to be doing the dishes at this present moment?

Modal adjectives and semi-auxiliaries with progressive complements can thus have present orientation even when they are used to attribute specific permissions and requirements. Future orientation is *not* obligatory, and so cannot explain why modal auxiliaries expressing these and other acknowledged types of circumstantial modality generally do not accept progressive and perfect complements.

### 7.3.2 Bare complements

An additional argument that future orientation is not obligatory for circumstantial modalities comes from modals with bare complements. While bare complements are compatible with either present or future orientation, it has been argued that only with future orientation can these modals be interpreted circumstantially (Werner 2006). (399), for example, is claimed on its deontic interpretation to only be able to describe the permissibility of a *future* state of affairs.

(399) Jill may be seated. [Werner's (3)]

But this empirical claim is false. (399) is fine as a response to (400), especially with a sentence-initial *yes*, showing that it can also be used to describe the permissibility of a current, or ongoing, state of affairs.

(400) Is Jill really allowed to currently be seated? The Queen *is* passing by.

More interesting for our purposes here are so-called pure circumstantial readings, on which modals are used to attribute abilities and potentials. Unlike other readings, these have been claimed to differ from allegedly epistemic ones exclusively in terms of whether the modal base is circumstantial or epistemic (cf. Kratzer 1991, 645–46). Hence, if the pure circumstantial readings of modals with bare complements do have obligatory future orientation, and if this can be traced to features of their having a circumstantial modal base, we might legitimately wonder whether *might*-claims with past or present orientation must have epistemic modal bases rather than circumstantial ones. Yet I will argue that we can no more analyze pure circumstantial (or ability) readings of modals as having obligatory future orientation than we can deontic ones.

An initial argument for this conclusion is inspired by Thomas' difficulty with past-tensed ability modals and their actuality entailments. These cases highlight the fact that many abilities and potentials that exist at the current moment can also be realized at the current moment. Hence, whenever we are attributing an ability or a potential, it seems we should allow for the possibility that it is currently being realized too, and so for the possibility of present orientation. Indeed, insofar as we can ever answer the

question whether something *can* happen by pointing out that that it *is* happening, we have to allow this possibility.<sup>117</sup> For the present is not always indicative of the future, and abilities and potentials can often be ephemeral, perishing upon realization. Hence, insofar as our answer to the question what can happen is ever a good one, the original question cannot have had obligatory future orientation.

A more developed argument comes from intuitive truth-value differences between the members of the two pairs below. In any context in which the claim made with (401) is true, it seems, the claim made with (402) will not be so.

- (401) By the end of this semester, John will be able to speak French.
- (402) John can speak French.

Similarly, while the metaphysical claim made with (403) seems intuitively as though it could be true, the ability claim made with (404) would intuitively seem to be false.

- (403) My unborn grandson might climb Mount Everest.
- (404) My unborn grandson can climb Mount Everest.

Yet these intuitive truth-value differences are hard to accommodate if we assume that modals on their ability readings have obligatory future orientation.

Take the second pair. Say that we assume that modals on their metaphysical readings have totally realistic modal bases, while on their ability readings they have circumstantial modal bases that are not totally realistic. Then the value of the modal base for (404) would be a subset of the value of the modal base for (403). Hence, the set of worlds determined by the modal base for (404) would be a superset of the set of worlds determined by the modal base for (403). It should thus be easier for a modal claim with existential force like (404) to be true than for one like (403) to be so. Indeed, the truth of the latter should entail the truth of the former. For if the total current circumstances are compatible with the world developing in such a way that my unborn grandson climbs Mount Everest at some future time, then any subset of the total current circumstances should be compatible with that developmental possibility too. And yet the truth of (403) does *not* entail the truth of (404).

Perhaps proponents of obligatory future orientation could appeal to a difference in the ordering source for modals on their metaphysical and ability readings. If (403), for example, had an empty ordering source, while (404) had a stereotypical one, then this could make the domain of quantification for (404) the smaller of the two. With (404), for example, we would look at worlds that developed in the stereotypical way, while with (403) we would look at worlds that developed in any way whatsoever.

I am skeptical. In uttering (404), I seem to be attributing an ability to my unborn grandson, not commenting on what he does at some future time in some stereotypical world. And, as I have already alluded to, Kratzer (1991, 645–46) also suggests that modals used to attribute abilities and potentials—pure circumstantials, as she too calls them—have empty ordering sources. However, the point can be waived. For even if a difference in the ordering source explains why the truth of the claim made with (403)

---

<sup>117</sup> Thomas (2014, 442) also considers such cases, using them to distinguish pure circumstantials from priority modals, but does not draw the correct conclusion from them.

does not entail the truth of the claim made with (404), similar reasoning about the modal base applies to show that the truth of the claim made with (401) should entail the truth of the claim made with (402) on the view that ability modals have obligatory future orientation. And here there is no difference in the ordering source to appeal to.

Intuitively, the ways in which the world can develop at some future time are a subset of the ways in which it can develop at the current time. The passage of time closes off possibilities, as proponents of future orientation and branching-time models of accessible worlds are apt to note (cf. Condoravdi 2002; Werner 2006). Hence, if at some future time, there is a possibility for the world to develop in such a way that John speaks French, then at the current time there should be that very same possibility too. Thus, the truth of the claim made with (401) should entail the truth of the claim made with (402), when in fact it seems to do the exact opposite.

Perhaps proponents of obligatory future orientation could hold that we are wrong to think that the values of the modal bases for (402) and (404) are subsets of those for (401) and (403). For example, maybe the value of the modal base for (402) includes the fact that John has *not* had a course in French, whereas the value of the modal base for (401) would not. This would be to no avail, however. For the fact that John has not currently had a course in French is compatible with the world developing in such a way that he does take such a course, and so compatible with the world developing in such a way that he does speak French. To block this developmental possibility, we need the fact that John has not taken a course in French to be held constant throughout future time. This, in essence, amounts to a denial of future orientation.

Perhaps future orientation could be salvaged if we held that it involved restricted quantification over future times, so that the claim made with (402) was a claim about what was possible for the *near* future. But there are obvious limitations to this response. The claim made with (402) does not become true until John has acquired a certain ability, no matter how near in the future he acquires it. This is how the modal claim made with (405) could be true at the start of the day even if that made with (406) only became true at the end.

(405) Hurry over! My toddler might {finally} walk today.

(406) Hurry over! My toddler can {finally} walk today.

The view that ability claims have obligatory future orientation thus cannot readily explain why attributions of abilities and potentials to those who do not yet possess them are false.<sup>118</sup> While abilities and potentials obviously have implications for how the future may develop, modeling these in terms of obligatory future orientation is the wrong way to proceed. Circumstantial modal claims with bare complements must, in some cases, be compatible with present orientation.

---

<sup>118</sup> Not readily. However, as already noted, Thomas (2017) has recently claimed that one important class of exceptions to the alleged correlation between circumstantial modal bases and future orientation involves the presence of covert generic operators. Such operators would likely be appealed to for (402) and could possibly also be appealed to for (404) and (406). Yet such operators could not be appealed to in the case of abilities that are complex but potentially also ephemeral: e.g., the ability to perform well Chopin's Etude in G# minor.

### 7.3.3 Circumstantial modal bases

Setting aside the empirical facts for the moment, there are also theoretical difficulties for the attempted explanations of aspectual restrictions on *can*'s complements. These restrictions could not plausibly be attributed to features of the modal base, not for priority modals, for the reason that Thomas points out, but also not for pure circumstantials, for much the same reason. Just as the truth-value of a proposition is not indicative of its deontic or teleological status, neither does the prejacent's truth necessarily indicate the presence of an ability, nor its falsity the absence of a potential (cf. Thalberg 1972). Hence, the reasoning that suggests that Werner's interpretive principle does not apply to priority modals suggests that it does not apply to ability modals either.

An additional reason to think the temporal properties of ability modals cannot be attributed to the modal base is the fact that they remain unchanged in the presence of modifiers, which have been assumed to restrict the modal base's value. There is no discernible difference in the temporal properties of (407) and the Kratzerian (408), for example. If the latter has future temporal orientation, then so does the former.

(407) Given {only} the soil, hydrangeas can grow here.

(408) Hydrangeas can grow here.

However, the facts denoted by the modifier are exclusively facts about the soil. Hence, if the value of the modal base is restricted to these facts, it should exclude the fact of the matter about the prejacent's truth value at the modal's evaluation time, and so be compatible with present orientation. As a result, if both (407) and (408) are best glossed as having future orientation, this orientation would not seem to be due to features of the modal base.

To summarize: the fact that *can* does not freely accept progressive and perfect complements has sometimes been taken as evidence that the modal claims it is used to make have obligatory future orientation, where this is attributed to the modal base for those claims being circumstantial, or non-epistemic. However, I have shown (i) that perfect and progressive complements sometimes are acceptable with *can*; (ii) that bare complements of *can* cannot always be analyzed in terms of future orientation anyway; and (iii) that future orientation cannot be attributed to the modal base's being circumstantial either. We need a better explanation for why *might* and *can* differ in how freely they accept progressive and perfect complements than the standard one in terms of the difference between epistemic and non-epistemic modality. My aim in the next two sections is to provide just this.

## 7.4 Rethinking the temporal interpretation of *can*

### 7.4.1 Modal states

Our starting point will be the difference between claims about possibilities and claims about potentials, observed above in (403) and (404). If there is a potential for my unborn grandson to *develop* the potential to climb Mount Everest, then it is true that he might climb Mount Everest. Yet it is not true, in the ordinary sense, that he can.

One possible response to this fact is to hold that *can* imposes obligatory present orientation: that for it to be true that my unborn grandson can climb Mount Everest,

there must be some accessible world in which, at the current moment, he *does* climb Mount Everest. However, since my unborn grandson presumably does not exist at the current moment in any accessible world, this condition is not met, and so the *can*-claim is false. In this way, we could hope to account for the fact that attributions of potentials and abilities are generally not true until those potentials and abilities have been acquired, even if the corresponding possibility claims are already true.

A problem for this response is that sometimes the prejacent of a *can*-claim imposes temporal constraints incompatible with obligatory present orientation. With (409), for example, the event of John's winning the race is presumably located wholly within the following day.

(409) John can win the race tomorrow.

We are not saying, that is, that in some accessible world John wins tomorrow's race today, contrary to what obligatory present orientation would require. This can be seen even more clearly with an example like (410).

(410) If the election were held tomorrow, there's no way John could win. But he can win the November election.

With (410), we are saying that John presently has the potential to win an election held in November, but not the potential to win one held tomorrow. The present potential in this case is for a specific future event, and the futurity of the event is what is key to the existence of that potential.

A seemingly better response to the contrast between *might* and *can*, then, is to abandon the attempt to account for the truth of claims about potentials and abilities exclusively in terms of the possible truth of a proposition at some present or future time. We should instead recognize that potentials and abilities are also states of individuals, locations, situations, etc. (cf. Homer 2010 on deontic states). These states are modal ones, and so may indeed have to be specified in terms of the possible truth of a proposition at some present or future time. However, the possible truth of this proposition does not itself guarantee the existence of a potential or ability, as we have seen from the cases above. Modal states are just like any other state or eventuality: they must be represented by an event variable all their own. The event variable introduced by the modal's complement, and used to represent the possible truth of a proposition, will not suffice.

If potentials and abilities are indeed modal states, attributions of them will only be true when those states inhere in the actual world, much like attributions of non-modal states. For it to be true that John loves Mary, for example, he has to have actually transitioned to that state. It does not matter how close he is to doing so. Similarly, for it to be true that John can speak French, it does not matter how close he is to having that ability: he has to have actually reached that state transition. This is in contrast to its being true that John might speak French, or my toddler finally walk today. These cases do not require those same state transitions to have occurred; they merely require other state transitions to be possible. *Can* differs from *might*, then, insofar as what it is used to attribute seem to be *modal states*.

### 7.4.2 Representing modal states

In representing modal states, two basic requirements must be met. First, the terms used to denote these states must introduce event variables to represent them, whence the name *event-level modals*. Second, these variables must be related in some way to the event variables introduced by the modals' complements. These requirements met, there is more than one way to proceed.

Homer (2010), for example, has the event-level modal take a fully propositional complement and then use its event variable to generate the set of accessible worlds, all or some or none of whose members must verify that complement depending on the force of the modal. I approach things in a slightly different manner, treating event-level modal constructions as analogous to causative and inchoative ones. When the modal introduces its event variable, it simultaneously performs existential closure on the open event variable of its complement, while also establishing a specific semantic relation between the two.

Here is an example of what this would look like, using the lambda notation and assuming, for the moment, that event-level modals combine with simple predicates of events.<sup>119</sup> (Following Hacquard (2010), I use  $\varepsilon$  to denote the type for eventualities.)

$$(411) \llbracket can_{ability} \rrbracket^{w,t,f} = \lambda P_{\langle \varepsilon, sP \rangle} . \lambda e . \\ \exists e' [\text{ABILITY-TO}(e)(e')(w) \ \& \ \exists w' [w' \in \cap f(w)(t) \ \& \ P(e')(w')]]$$

What does it mean to say that an eventuality  $e$  in a world  $w$  is an ability to  $e'$  in  $w'$ ? This is simply one way of describing a modal state, a state that can only be described by reference to some potentially non-actualized eventuality. Perhaps we do better to avoid appeal to possible worlds altogether in the description of such states (cf. Vetter 2013; or the semantics for modal objects in Moltmann 2018). However, the possible worlds framework is standard and ubiquitous. As long as we interpret it in a Kripkean fashion, rather than a Lewisian one, no difficulty arises in referencing an eventuality in another world to describe one in this (cf. Kripke 1980; Lewis 1986).

I should make clear that I have no analysis of the metalanguage predicate ABILITY-TO to offer in this chapter. Yet I do not think this is a problem. In analyzing causative or inchoative constructions, it is possible to argue that a second event variable is needed independent of any particular analysis of the metalanguage predicates CAUSE or BEGIN used to relate that variable to the first. The same holds here. What the argument of Sections 7.3.2 and 7.4.1 has shown, if I am correct, is that we cannot adequately represent the temporal properties of abilities through a relation between a modal's own evaluation time and its prejacent's. Instead of trying to encode these properties on an event variable representing what the ability is an ability *for*, we need to encode them on an event variable representing the ability *itself*. This argument for modal states does not require us to have developed a philosophical analysis of those states. Instead, just as we can rely on our pre-theoretical understanding of causes and beginnings in developing semantic analyses of causatives and inchoatives—or even treat causes and beginnings as unanalyzable primitives—so with abilities.

<sup>119</sup> Readers unfamiliar with the lambda-notation are directed to Gamut (1991, chaps. 4–5).

### 7.4.3 Revisiting the representation of modal states

In introducing a representational framework for event-level modals, I assumed for the sake of simplicity that they combined with simple predicates of events. In the case of modals with bare complements, this is initially plausible. However, modals like *can* admit not just bare complements but also, as we have seen, progressive and perfect ones. And aspectual complements are standardly assumed to be predicates of times, with some aspectual operators playing dual semantic roles: “[providing] existential quantification over the event described by the VP and [locating] its running time with respect to the time provided by tense” (Hacquard 2010, 97). On this view of aspect, there would not be any open event variables for the event-level modals to operate on. Hence, we could not convert our initially plausible entries for event-level modals with bare complements into equally plausible entries for modals with progressive and perfect ones.

There are several possible ways to proceed at this point. One possibility is to rethink the semantics of aspectual operators, limiting them to one role and assigning the existential quantification over events to an independent operator (cf. Matthewson 2012, 8). Another is to leave the semantics of aspectual operators intact, but either to adopt Homer’s method for modeling modal states mentioned above, or to hold that the apparently progressive and perfect complements of event-level modals are not, in fact, predicates of times but instead complex predicates of events.<sup>120</sup> I adopt and develop the first possibility here, insofar as it seems to me to be part of the best overall picture of modal and aspectual operators and their interactions.

Following Rullmann and Matthewson (2018), I assume that aspectual operators fall into one of two categories. Markers of *inclusion aspect* relate an eventuality’s run time to a reference time: *imperfective* aspect indicates that the run time of the event, sometimes also called its *temporal trace*, and denoted  $\tau(e)$ , includes the reference time; *perfective* aspect indicates that the opposite inclusion relation holds. Markers of *ordering aspect* then order that reference time with respect to an evaluation time: *perfect* aspect indicates that the reference time wholly precedes the evaluation time; *prospective* aspect indicates that evaluation time wholly precedes reference time.<sup>121</sup>

For Rullmann and Matthewson, both inclusion and ordering aspect are obligatory, though potentially covert, and there is a hierarchical relationship between the two that is also obligatory. Inclusion aspect performs the additional role of existential closure, converting a predicate of events into a predicate of times. Ordering aspect then takes that predicate of times as its input and yields a predicate of times as its output (2018, 287). The obligatory nature of inclusion aspect makes sense to me, as does also the obligatory hierarchical relationship between it and ordering aspect. Later I will say more to defend these claims, and also to defend the claim that ordering aspect is *not* obligatory. For now, the important point is to recast the meaning of both types of

<sup>120</sup> On this view, terms like *and still* in (391) and (393) and *as* in (395) would combine with predicates of times (the part of the modal’s complement in the progressive or perfect form) and predicates of events (the part of the complement in the bare form) to yield complex predicates of events.

<sup>121</sup> Rullmann and Matthewson (2018) posit two additional ordering aspects, *non-perfect* and *non-prospective*, which are complementary to perfect and prospective aspect.

aspectual operators once the former are no longer assigned the role of existential closure. I do this in (412)-(415), where  $i$  is used to denote the type for times.

$$(412) \llbracket PFV \rrbracket^{w,t,f} = \lambda P_{\langle e, s \rangle} . \lambda t . \lambda e . [\tau(e) \subseteq t \ \& \ P(e)(w)]$$

$$(413) \llbracket IMPFV \rrbracket^{w,t,f} = \lambda P_{\langle e, s \rangle} . \lambda t . \lambda e . [\tau(e) \supseteq t \ \& \ P(e)(w)]$$

$$(414) \llbracket PERF \rrbracket^{w,t,f} = \lambda P_{\langle i, \langle e, s \rangle \rangle} . \lambda t . \lambda e . \exists t' [t' < t \ \& \ P(e)(t')(w)]$$

$$(415) \llbracket PROSP \rrbracket^{w,t,f} = \lambda P_{\langle i, \langle e, s \rangle \rangle} . \lambda t . \lambda e . \exists t' [t < t' \ \& \ P(e)(t')(w)]^{122}$$

Given the entries in (412)-(415), the inclusion aspects now convert predicates of events into *functions* from times to predicates of events. In turn, the ordering aspects have such functions as both their inputs and their outputs. Converting our semantic entry for *can* from (411) into an entry that can accommodate aspectual complements is now a straightforward matter. We first need the event-level modal to existentially close the open time variable of its complement so that there is no conflict between the temporal information encoded on the modal's own event variable and the temporal information encoded on that of its complement (as could happen in the case of (409), for example). We then want to optionally forward-shift the evaluation time for the complement to help represent the fact that abilities and potentials are never abilities and potentials for past events but instead for present or future ones, thereby also overcoming an important lacuna in the semantic entry in (411). The result is (416).

$$(416) \llbracket can_{ability} \rrbracket^{w,t,f} = \lambda P_{\langle i, \langle e, s \rangle \rangle} . \lambda e . \\ \exists e' [\text{ABILITY-TO}(e)(e')(w) \ \& \ \exists w' \exists t' [w' \in \cap f(w)(t) \ \& \ t \leq t' \ \& \ P(e')(t')(w')]]^{123}$$

#### 7.4.4 Aspectual operators and aspectual morphology

At this point, we have two semantic entries for event-level modals like *can*. Yet if we assume that the bare complements of modals, while not morphologically marked for aspect, are nonetheless semantically marked for at least inclusion aspect, then we can reduce our two entries to one. This is a simplifying assumption on multiple counts.

First, if we maintain two entries for event-level modals, we will want to modify our first entry in (411) to account for the lacuna that (416) helped bring to light. In doing so, we will essentially have to build inclusion aspect into the meaning of the modal. But once we have built inclusion aspect into the meaning of the event-level modal, why not just hold that event-level modals always combine with complements already marked for inclusion aspect, albeit sometimes covertly so?

<sup>122</sup> The entry for the perfect operator in (414) is provisional. I will subsequently argue that we should adopt a view of the perfect as contributing the *post-state* of an eventuality (cf., a.o., Parsons 1990; Kamp and Reyle 1993), an argument that provides further support for removing existential closure as one of the functions of inclusion aspect operators.

<sup>123</sup> Optional forward-shifting correctly ensures non-past orientation for modals with complements in the bare form. We do not interpret *John can jump* as saying that John has a present ability for a past jumping event, something that the entry in (411) did not rule out. However, it does not ensure non-past orientation for modals with perfect complements, which is good. For then the incoherent attribution to John of a present ability for a past event can be held responsible for the general anomaly of *John can have jumped*. (When the post-state view of the perfect is introduced, the details of this explanation will have to be reworked, but the broad outline will remain the same.)

Second, morphologically bare, perfect, and progressive sentences occur not only as complements of event-level modal auxiliaries but also of proposition-level ones. All of them can further occur as standalone sentences in matrix clauses. Hence, if we do not appeal to covert aspectual operators for bare sentences, we will require multiple semantic entries, not only for event-level modal auxiliaries, but also proposition-level ones.<sup>124</sup> Further, since the role of inclusion aspect is to locate an eventuality within the temporal structure of the evaluation world, then something else will have to play this necessary role for matrix clauses if covert aspectual operators do not, presumably tense. Yet for tense to play this dual role we run a risk of redundancy or inconsistency whenever it co-occurs with overt aspect. That, or we have to posit multiple semantic entries, this time for each of the tense operators (past, present, etc.).

Additionally, in English, certain restrictions on the types of eventualities described by bare present tense, matrix-clause sentences are best explained in terms of inclusion aspect. But if the different inclusion aspects play a role in explaining the absence or presence of these restrictions, then by far the simplest thing to do is to appeal directly to covert aspectual operators themselves.

Here is an example of the restrictions I have in mind. (417) cannot be used to describe a current event of John's running, unless what Ogihara (2007) calls the "live sports broadcaster" manner of speech is being used (cf. Bennett and Partee 1972).

(417) John runs.

Instead, (417) must be used to state that John's running is a generic or habitual event. "Eventives" thus contrast with "statives", which do occur freely in the simple present. (418), for example, is used to describe a current state of John's being charming. And while this state may also be a habitual or generic one, if we are lucky, this is not what (418) describes, if the difference between (419) and (420) is any indication.

(418) John is charming.

(419) John is usually charming, but not now.

(420) ?John is charming, but not now.

The standard explanation for this contrast is that the English present tense denotes not an interval, but an instant, while eventives, unlike statives, describe situations that only hold at intervals (cf. Bennett and Partee 1972; Taylor 1977; Vlach 1981). To this explanation must also be added assumptions about default aspectual interpretations: perfective for eventives (Thomas 2014, 446), imperfective for statives (Rullmann and Matthewson 2018, 302). If the stative had a perfective interpretation, then the relevant situation would not only have to hold at the present instant but also to be wholly included within it. And while the transitions to and from a state may be instantaneous, it is highly unlikely that the two will occur at the same or simultaneous instants. On the other hand, if the eventive had an imperfective interpretation, then all that would be required would be for the present instant to be included within the interval at which the situation it described held. Default aspectual interpretations are responsible, then, for the fact that the stative must wholly include the present instant, which is

<sup>124</sup> But cf. Condoravdi (2002) for a creative alternative.

possible, while the eventive must be wholly included within it, which is not. And this explains why generic and habitual interpretations are generally the only available ones for eventives in the simple present tense in English.

Yet what does it mean to say that bare eventives or bare statives receive a certain aspectual interpretation by default? Insofar as eventives and statives can co-occur with overt markers of their non-default inclusion aspect (cf., a.o., Hacquard 2009; Homer 2010), I assume we do not want to build their default aspectual interpretations into the verbal predicates themselves. And I do not see how we could derive both interpretations by building aspectual meaning into the simple present tense either. The best explanation for the default interpretations would thus seem to be that, whenever eventives and statives are not overtly marked for inclusion aspect, they co-occur with certain types of covert inclusion aspect operators by default.

This argument for covert inclusion aspect operators provides a foil for an argument *against* such operators for ordering aspect. Rullmann and Matthewson propose that languages contain covert complements to their overt ordering aspects (2018, 287). In English, the perfect is the overt member of this pair, and the non-perfect the covert one. Yet if covert non-perfect aspectual operators were available, we would no longer predict the restriction on eventives in the simple present. For the non-perfect would optionally forward-shift the reference time from the present instant, thereby opening up the possibility for a future interval within which the relevant eventuality *could* be wholly included. The restriction against eventives in the simple present is thus *prima facie* evidence against covert non-perfect operators.

In positing covert aspectual operators, it is important not to lose sight of our focus on aspectual *morphology* and, specifically, why progressive and perfect morphology are much less freely available under what I am now calling event-level modals. The explanation here is relatively straightforward. For event descriptions used to specify potentials, there is often simply no need to impose the particular relations to particular reference times that the progressive or the perfect would. Only when there is a need to describe the potential to do two things either simultaneously or sequentially does it make sense to use either. That said, the details of this explanation require a closer examination of the progressive and the perfect themselves. And this is best provided in the context of an exploration of the temporal properties of modals, like *might*, that do freely accept such forms. I turn to this topic now.

## 7.5 Revisiting the temporal interpretation of *might*

### 7.5.1 A puzzle about temporal orientation

Unlike *can*, *might* freely accepts complements in the progressive and perfect forms, and seems intuitively to be compatible with the whole range of temporal orientations: past, present, and future. This might suggest that *might* places no constraints on its prejacent's evaluation time. Yet matters are significantly more complicated than this.

An initial puzzle arises with bare eventives, which occur freely as complements of modals like *might*, as in (421), but not so freely in the simple present tense, as we have seen. Assuming that *might* occurs in the simple present tense,<sup>125</sup> it follows that it cannot simply pass down its evaluation time to its prejacent. And yet *might* occurs

---

<sup>125</sup> Something sometimes claimed to be obligatory; cf. Groenendijk and Stokhof (1975).

freely with bare statives too, as in (422), suggesting that it cannot simply forward-shift the evaluation time for the prejacent either. For while (421) describes a possible future event, (422), by default, describes a possible current state.

- (421) John might win {tomorrow}.  
 (422) John might be charming.

In her seminal analysis of the temporal properties of non-root modals like *might*, Condoravdi (2002) proposed that they *extend*, rather than *shift*, the evaluation time for the prejacent. They provide an interval during which the prejacent may obtain that starts with the modal's own evaluation time and extends to the end of time (2002, 13). This allows for present orientation with bare statives, and future orientation with bare eventives. If we assume that perfect aspect simply back-shifts the evaluation time for the prejacent, then it also allows for past orientation too, as in (423).

- (423) The government might have re-established the embargo.

With (423), on Condoravdi's analysis, we require there to be some time prior to an interval stretching from the current time into infinity at which the government re-established the embargo.

Past orientation on her analysis, Condoravdi observes, is "from the perspective of the time of utterance, not some future time" (2002, 13). This is a direct result of the modal's extending, rather than shifting, the evaluation time for the prejacent into the future. And it is just this observation that has led some authors to abandon that analysis. Klecha (2016) notes that in (424), for example, we in fact need the past orientation to be from the perspective of some future time, rather than from the perspective of the utterance time (call this the past-in-the-future reading of a *might-have* sentence).

- (424) Don't book your ticket to Cuba for next summer just yet. The government might have re-established the embargo (by then). [Klecha's (69)]

Klecha concludes from this fact that we need there to be "some element distinct from the modal itself which imparts the future-shifting" (2016, 31). Klecha attributes this to the modal base. Others have posited embedded implicit tenses (McCawley 1971). And still others have proposed that this distinct element is covert prospective aspect (Matthewson 2012; Rullmann and Matthewson 2018; Thomas 2014; 2017).

Appeals to a distinct, future-shifting element do indeed solve the problem facing Condoravdi's analysis. However, they also create problems of their own. If covert prospective aspect were available for non-modal sentences, as well as for modal ones, for example, then bare eventives should be able to occur freely in the simple present without having to be coerced into a generic or habitual reading. Yet generally they cannot. Hence, if we appeal to covert prospective aspect, we have to posit that it is available in modal contexts but, for some reason, not in non-modal ones. Appeals to the modal base are similarly unpromising. Klecha, for example, wants circumstantial

modal bases to impart future orientation.<sup>126</sup> However, circumstantial modal bases, as we have seen, are also compatible with present orientation. And as for implicit tense, Condoravdi has provided a battery of arguments for why temporal orientation should not be accounted for in this way (the interested reader is directed to Condoravdi 2002, 8–11). Hence, by accounting for future orientation via a distinct, future-shifting element, we avoid the problem facing Condoravdi only at a cost of creating new ones.

How do we account for the future orientation of *might* with bare eventives then? In what follows, I will suggest that Condoravdi was right to attribute the introduction of futurity to the modal. Yet she was wrong about *how* that futurity was introduced, and also about the contribution of the perfect.

### 7.5.2 Solving the puzzle about temporal orientation

Condoravdi claims that, in the absence of perfect morphology, non-root modals “have a future orientation optionally with stative predicates and obligatorily with eventive predicates” (2002, 11). Yet this optional future orientation for statives, as Laca (2012) observes, is crucially linked to the presence of future-oriented adverbials (2012, 5–8; cf. also Condoravdi 2002, 13). Without the adverbial in (425), for example, the default interpretation for the bare stative in (426) is a present-oriented one.

(425) John may be drunk by the time we arrive. [Laca’s (6c)]

(426) John may be drunk. [Laca’s (6b)]

Now Condoravdi did not really offer an analysis of why statives default to present orientation when future orientation is also available (cf. Copley 2008; cited in Laca 2012). In contrast to eventives, she did require statives to merely overlap with the extended evaluation time afforded by a modal. Yet while this explains the possibility of present orientation, it does not explain the default to it. But Laca points out that bare statives, on their default imperfective interpretation, need to be anchored to a specific reference time that they include. In the absence of future-oriented adverbials, the only available anchors are the endpoints of the extended evaluation time provided by the modal, while the only plausible one is the initial endpoint: the present moment.

This explanation for bare statives’ default orientation extends well to progressives, Laca notes. While these receive present orientation by default, with future-oriented adverbials they can receive future orientation too, as (427) shows.

(427) John may be getting drunk {by the time we arrive}. [≈ Laca’s (7)]

Yet this optional future orientation is not the default, for the progressive functions in part to “stativize” event predicates. That is, it converts a description of an event into a description of a corresponding state. Just like bare statives, progressives thus require an anchor for the reference time that they include. And, in the absence of a future-oriented adverbial, the default anchor is the initial endpoint of the modal’s evaluation time.

---

<sup>126</sup> He even argues, in a twist of intellectual history, that what appear to be epistemic modals with future orientation are in fact circumstantial ones (cf. Klecha 2016, 34–37).

Could a similar explanation be given for why the past-in-the-future readings that are optional with perfect complements, as in (428), are likewise not the default there?

(428) John may have gotten drunk {by the time we arrive}. [ $\approx$  Laca's (8)]

Indeed, Laca argues, it can. For we can view the perfect, like the progressive, as also stativizing a predicate, in this case contributing the *post-state* of an event. On this view, the perfect indicates not just that some event culminated at a time prior to the evaluation time but also that the state of that event's having culminated, that event's *post-state*, holds at the evaluation time (cf. Parsons 1990, 235).<sup>127</sup> One possible semantic entry for the perfect on this view is given in (429), providing an additional reason not to have inclusion aspect close the event variable of its complement.

(429)  $[[PERF]]^{w,t,f} = \lambda P_{\langle i, \langle e, st \rangle \rangle} . \lambda t . \lambda e .$   
 $\exists e' \exists t' [\text{POST-STATE}(e')(t)(w) \ \& \ t' < t \ \& \ P(e')(t')(w)]$

Now say that we take the evaluation time provided by the modal to be an interval stretching from the modal's own evaluation time to the end of time, as Condoravdi does. Then a past-in-the-future reading would be impossible, since the relevant post-state would always have to hold at the modal's own evaluation time. However, say we have the modal optionally forward-shift its prejacent's evaluation time, rather than extending it, much as Abusch (2012) proposes. Then the evaluation time for both the post-state, and the eventuality it was the post-state of, could be in the future of the modal's own evaluation time, making past-in-the-future readings possible. However, the post-state, like any other stative, will need to be anchored to the specific reference time that it includes. And, in the absence of future-oriented adverbials, the only available anchor for this reference time will be the modal's own evaluation time. And this explains why past-in-the-future readings, while possible, are not the default.

On the view now being proposed, instead of appealing to past, present, and future orientation, we should, strictly speaking, appeal only to present and future orientation. The appearance of past orientation with the perfect is due to its describing a state that, while evaluated at or after the modal's own evaluation time, is defined in terms of an event that is prior to that state. Hence, given the default case of present orientation for statives, the event is also prior to the modal's own evaluation time, whence the appearance of past orientation (cf. Laca 2012, 8).

The uniformity of statives, progressives, and perfects with regard to their default and optional orientations makes sense given Laca's analysis of the progressive and the perfect as "stativizing" eventive predicates.<sup>128</sup> What the pattern in (425)-(428) thus reveals is a basic fact about statives and their contrast with eventives. Statives have present orientation by default, while bare eventives have obligatory future orientation. Future orientation is obligatory for bare eventives insofar as they cannot hold at the modal's own evaluation time, the present instant, and so require forward-shifting to

<sup>127</sup> In the case of statives, the culminating event for the perfect may be either the initial transition to the state or, more frequently, the final transition from it (cf. Kamp and Reyle 1993, 567–68).

<sup>128</sup> For the perfect, the term "stativize" is unfortunately misleading, insofar as the perfect, unlike the progressive, can occur with stative predicates.

an interval posterior to that time. And present orientation is the default for statives insofar as the modal's own evaluation time provides the most natural anchor, sometimes the only one, for the reference time that they include.

To summarize: the solution to our puzzle about the obligatory future orientation of bare eventives involves two changes to Condoravdi's solution. The change from an extended evaluation time to an optionally forward-shifted one explains how future orientation is possible with perfect complements. And the post-state analysis, and the explanation for why statives do not generally forward shift, explains why present orientation is the default. Both changes turn out to be necessary. And we are now also in better position to return to our original puzzle about the differences between *might* and *can* with regard to how freely they accept perfect and progressive morphology.

### 7.5.3 Aspectual morphology revisited

On the standard analysis, *might* and *can* share the same basic semantic entry but differ in terms of their modal base: epistemic for *might*; circumstantial for *can*. On my analysis, in contrast, *might* and *can* share a circumstantial modal base but differ in their semantic entries. *Can* is an instance of an event-level modal: a modal that introduces its own event variable and corresponding modal state description. *Might*, on the other hand, is an instance of a proposition-level modal, defined negatively as a modal that does *not* introduce its own event variable and corresponding modal state description. On standard analyses, at least in possible world semantics, *all* modals are proposition-level ones. However, part of the argument of this chapter has been that the temporal properties of some modal claims cannot be accounted for exclusively in terms of relations between the modal's and its prejacent's evaluation times. We need also to introduce an additional event variable for the modal for temporal properties to be encoded on. That is, we need some modals to be event-level ones.

Event-level modals are used to attribute modal states: these are modal properties of individuals, situations, locations, etc. In this, they contrast with proposition-level modals, which are used simply to attribute modal properties to propositions, including the property of being possibly true relative to a set of facts.

How does the distinction between proposition- and event-level modals help make sense of why *might* but not *can* freely accepts progressive and perfect complements? Recall that both the perfect and the progressive function to stativize event predicates and to impose particular relations to particular reference times: relations, respectively, of sequentiality and simultaneity. When it comes to specifying propositions, these relations will be useful. By relating an event description to different reference times, they enable different propositions to be denoted. When it comes to specifying modal states, however, these relations are useful much less often. Any analysis of potentials, for example, should hold that the potential to perform some activity entails a potential to be in the state of performing that activity, or in the post-state of having performed it (the contributions, respectively, of the progressive and the perfect). As such, while modals used to attribute modal properties to propositions can be expected to freely accept progressive and perfect morphology, modals used to attribute modal states can be expected not to. In other words, we can expect proposition-level modals like *might* to freely accept such complements, but not event-level ones like *can*.

When event-level modals do accept progressive and perfect morphology, it will be because the temporal relations they impose are in some way useful. The progressive, for example, provides an interval within which other activities at a specific reference time can be included. And the perfect, with or without the post-state view, provides a time prior to which other activities may occur. Hence, if the potential to perform an activity simultaneous with (or prior to) other activities is relevant, we can expect the progressive (or the perfect) to be acceptable, just as in Section 7.3.1.<sup>129</sup>

That said, there are other means available for indicating the potential to engage in activities simultaneously or sequentially, as (430) and (431) show, and so there may be additional factors at play when the progressive and the perfect are used.

- (430) John can juggle and unicycle even as he carries on a trenchant conversation about Kant.  
 (431) John can sing the Ave Maria and then launch into a vulgar outburst.

Pursuing these factors would take us too far afield here.<sup>130</sup> For our purposes it suffices to note that we should not expect progressive and perfect complements of event-level modals to be acceptable outside of relatively restricted contexts. By appeal to modal states, we can thus explain the restrictions on progressive and perfect complements for *can* that cannot be explained in terms of temporal orientation and modal bases.

#### 7.5.4 A uniform semantics revisited

To draw this section to a close, I would like to address an issue that I imagine will be at the forefront of readers' minds: namely, that I seem to have abandoned Kratzer's attempt to provide a uniform semantics for modals of different flavors, given my appeal to modal states and my distinguishing event-level modals from proposition-level ones. And, despite the arguments of the previous chapter, this will no doubt seem to many to be a significant theoretical cost of this chapter's proposal.

This cost is not as significant as it seems. First, we can provide a schematic version of the semantic entry for the event-level modal in (416) to cover event-level modals in general, as in (432).

- (432)  $[[can]_{\text{flavor}}]^{w,t,f} = \lambda P_{\langle i, \langle e, st \rangle \rangle} \cdot \lambda e.$   
 $\exists e' [\text{MODAL-STATE}_{\text{flavor}}(e)(e')(w) \ \& \ \exists w' \exists t' [w' \in \cap f(w)(t) \ \& \ t \leq t' \ \& \ P(e')(t')(w')]]$

<sup>129</sup> The same holds if a specific reference time itself is relevant (e.g., the present moment), though here differences between the modal flavors seem to come into play. For example, the combination of the progressive with the temporal adverb *now* can be used to denote permissions but seemingly not to denote abilities or potentials. I assume there is a pragmatic explanation for this fact. It is odd to suggest that someone has the potential to be in the state of performing an activity at the present moment, because this state requires the prior onset of that activity, something over which they now have no control (though they do have control over its current onset). In contrast, whether you are permitted to be in the state of performing an activity *could* be relevant to deciding whether you continue it.

<sup>130</sup> One possibility is that the state description carries a connotation of absorption or immersion that the event description does not. Hence, the progressive and the perfect perhaps serve to emphasize a contrast between activities that are unexpectedly able to be performed simultaneously or sequentially.

If we appeal to a Recanati-style variable over modal flavors (cf. the previous chapter), this would allow us to limit the expansion from one semantic entry to two.

Second, given the argument for optional forward-shifting from Section 7.5.2, we can adopt (433) as our semantic entry for a proposition-level modal like *might*, to which an event-level modal like *can* would then more or less simply add the second line of (432) with its event variable and corresponding modal state description.

$$(433) \llbracket \textit{might} \rrbracket^{w,t,f} = \lambda P_{\langle i, \langle e, st \rangle \rangle} \lambda t. \lambda e. \\ \exists w' \exists t' [w' \in \cap f(w)(t) \ \& \ t \leq t' \ \& \ P(e)(t')(w')]$$

This is an advantage of my representational framework for event-level modals over Homer's, which requires a special accessibility relation for event-level modals, and so cannot grant them the same core meaning as proposition-level ones.<sup>131</sup>

Adopting a modal state analysis may not actually require us to give up too much, then, by way of the Kratzerian non-ambiguity project. We may only have to posit two semantic entries for existential modals, one of which simply builds off the other. This is not a significant theoretical cost, and it seems more than justified by the expanded empirical coverage.

## 7.6 The counterfactual readings of *might*

### 7.6.1 An account in terms of aspect

To this point, I have not said anything about the counterfactual readings of non-root modals, brought to prominence by Condoravdi (2002) with examples like (434).

(434) At that point he might still have won the game, but in the end he didn't.

Previously, no one had connected these readings to the presumed epistemic readings of such modals. However, Condoravdi argued for a decompositional analysis of *might have* into independent modal and aspectual operators, and for attributing the readings to a scopal ambiguity: for the presumed epistemic readings of *might have*, the modal takes scope over the perfect; for the counterfactual readings, the scope is reversed.

Evidence for Condoravdi's view is supposed to come from examples like (435), in which the adverb *still* is replaced with *already*, and the counterfactual reading is no longer available.

(435) He might have already won the game, #but in the end he didn't.

While *still* can take direct scope over a modal, but not a perfect, the opposite holds for *already* (cf. Condoravdi 2002, 18–19). And the apparent difference in availability of counterfactual readings with these adverbs is supposed to show that the relative scope

<sup>131</sup> My approach not only makes the similarities between the two types of modals more transparent but also potentially provides an explanation for the acquisition and historical development of the two, event-level modals (or roots) in both cases generally preceding proposition-level modals (or non-roots) (cf. Hacquard 2013; 2016; and the references cited therein). On my approach, this would be due to proposition-level modals being more generic and requiring abstraction away from the differences between the more specific event-level modals. But this topic is far beyond the scope of this chapter.

of the modal and aspectual operators for the two readings has been reversed, as in (436)-(437).

- (436) PERF > STILL > MOD [counterfactual reading]  
 (437) MOD > ALREADY > PERF [“epistemic” reading]

Now on the post-state analysis of the perfect that I have adopted, the reading that results for a non-root modal with a scope like (436) is incoherent, as the reader may attest for herself. Hence, Condoravdi’s analysis of the counterfactual readings of non-root modals is not available to me. Yet this analysis seems to me to be independently problematic, and in ways that point to what I believe to be the correct analysis of the counterfactual readings that Condoravdi discussed (cf. also Abusch 2012; Laca 2012; Thomas 2017; Rullmann and Matthewson 2018).<sup>132</sup> Let me explain.

The first problem for Condoravdi’s analysis is an evidential one. The contrast between *still* and *already* that she presents is consistent with the modal and aspectual operators having the same relative scope in both the counterfactual and the non-counterfactual readings, as in (438)-(439).

- (438) STILL > MOD > PERF [counterfactual reading]  
 (439) MOD > ALREADY > PERF [“epistemic” reading]

That is, while *still* may be unable to scope directly over perfect aspect, this does not mean that it cannot scope over a modal which does so. If the counterfactual reading is distinguished from the (presumed) epistemic one by its backward-shifted evaluation time, this backward-shifting could thus be due to something other than aspect taking scope over the modal.

The second problem for Condoravdi’s analysis relates directly to the first. It turns out that *already*, like *still*, is compatible with counterfactual readings, as in (440).

- (440) At that point he might have already won the game. But in the end, incredibly, he went on to snatch defeat from the jaws of victory.

But the scope of the operators in (440) has to be as in (439). Of the six logically possible scope combinations for (440), the two with *already* scoping lowest are ruled out by its sortal restrictions against eventives (cf. Condoravdi 2002, 8–9); the two with *already* scoping directly over the modal have already been ruled out; and the one with *already* scoping highest, and directly over the perfect, is ruled out for conceptual reasons (cf. Condoravdi 2002, 23–24). Hence, the backward-shifting of the modal’s evaluation time *must* be due to something other than aspect scoping over the modal.

---

<sup>132</sup> For what it is worth, modal adjectives on their non-root interpretations do seem to be marked for aspect, as in *It had been possible that John was the murderer*. I assume that the morphological aspect in this case is not the realization of semantic aspect. However, there are other ways to make sense of this fact that, while requiring a fair amount of tinkering with the details of my proposals for event- and proposition-level modals, are nonetheless compatible both with that distinction and with its role in explaining the central phenomena of this chapter. Additionally, my argument against interpreting aspect as taking scope over modal auxiliaries like *might* is independent of whether it takes scope over proposition-level modal adjectives.

The third problem for Condoravdi's analysis confirms as much. Condoravdi claims that the perfect taking scope over the modal, which she holds to be responsible for counterfactuality, is "possible only for modals that are in the so-called subjunctive form in English, such as *might*, *would*, *should*, *ought to*" (2002, 17). But the so-called subjunctive modals are just those modals that are, morphologically, the past tense forms of other modals (cf. Palmer 2001, 13–14). Hence, if the backward-shifting of the counterfactual readings seems to be possible only with the so-called subjunctive modals, a much simpler explanation for why this is so is that the morphological past tense is also indicative, in at least some cases, of a semantic past tense. There is no need to posit a scope-reversal mechanism that, for unknown reasons, is only available for morphologically past-tensed modals.<sup>133</sup>

### 7.6.2 An account in terms of tense

Condoravdi's counterfactual readings are the result, in my view, of semantic past tense on the modal. It is well known that the use of the semantic past tense to describe a past state often generates a cancellable implicature that that state no longer obtains, as in (441).

(441) In those days, John was quite charming. {As a matter of fact, he still is.}

Something similar could hold for the counterfactual readings of modals. Here the past tense would describe a past possibility for there to be the post-state of some specified but as yet future event, as in (442).

(442) Yesterday, John might still have finished his quilt in time for tomorrow's fair.

In using the past tense, we would imply that there is no longer the possibility for there to be this post-state, and so imply that the relevant event either will not, or did not, occur. This implication is cancellable, however, as (443) and (444) show.

(443) Yesterday, John might still have finished his quilt in time for tomorrow's fair. {As a matter of fact, he still might. He'll just have to work non-stop now because of his tomfoolery yesterday.}

(444) At that point, he might still have won the game. {As a matter of fact, he did go on to win, for what it's worth. But it certainly wasn't through any skill of his own after his bumbling at that point.}

A tense- and implicature-based account of counterfactual readings is subject to two challenges. First is the argument that non-root modals scope above tense and are

---

<sup>133</sup> A similar criticism applies to Rullmann and Matthewson (2018). These authors propose a three-way taxonomy of English modals that is completely silent on the fact that their Class-III modals are all morphologically past tense forms of their Class-I modals, while their sole Class-II modal, *must*, is the one modal in English that does not have a morphological past-tense. No wonder, then, that, for some occurrences of their Class-III modals with the perfect, they have to posit special, lexicalized forms.

“evaluated at the local now” (cf. Hacquard 2010, 88). Second is the argument that the implication of counterfactuality cannot be cancelled (cf. Portner 2009, 226–27).

Start with the first challenge. The argument that non-root modals scope above tense derives its force from the view of such modals as epistemic. Epistemic states are like modal states—abilities, potentials, states of permission and obligation, etc.—in that they can change over time. However, while past tense on a flexible semi-modal auxiliary like *have to* in (445) can be used to describe a past obligation, it cannot be used to describe a past epistemic state (Hacquard 2010, 87–88).

(445) Mary had to be home yesterday. [ $\approx$  Hacquard’s (12)]

On its presumed epistemic reading, (445) instead provides evidence for the speaker’s current state about some past event, suggesting that the modal is scoping above tense. These facts are apparently even clearer in languages which are morphologically more transparent than English, and in which modals regularly inflect for tense and aspect (cf. Hacquard 2006; 2009; 2010 for French and Italian examples).

Now I do not adopt an epistemic view of non-roots. Instead, non-roots differ from roots on my view in that the latter are used to attribute modal states, which are part of the temporal structure of the world, while the former are used to attribute oftentimes atemporal modal properties to propositions. Whether an already settled proposition is a non-trivial possibility, for example, will likely not change as the facts change over time.<sup>134</sup> Given the general atemporality of proposition-level modality, it is no surprise that a flexible modal marked for past tense should generally receive an event-level reading whenever the past tense is interpreted as scoping above the modal. But this does not mean that semantic past tense is unavailable for proposition-level modals.

In fact, the general exception to atemporality for proposition-level modals comes from propositions that have *not* yet been settled at the time of evaluation. As the facts change over time, and the space of possibilities decreases, so may the modal status of these propositions change too. Hence, for proposition-level modals, semantic past tense only makes sense in a specific range of cases: viz., those with future orientation. But these are precisely the cases on which counterfactual readings can occur too, as (446) shows.

(446) At that point, it was possible that John would finish his quilt in time for the next day’s fair.

That said, for modal auxiliaries in English, there is a significant additional obstacle to interpreting past tense as taking scope over a proposition-level modal. This is the fact that the morphological past tense on a modal auxiliary is not always the realization of the semantic past. Instead, in addition to remoteness in time, past morphology may also be used to indicate remoteness in logical space (Palmer 2001,

---

<sup>134</sup> Not likely, but not necessarily impossible either. For example, if we adopt the view that “murder will out”, then it might make sense to utter a sentence like (iv).

(iv) Yesterday, {Raskolnikov/Bill Sikes} might have been innocent, given his disposition and actions. Today, he can’t be.

nn. 13–14; cf. also chs. 2.1.5, 3.2.3, 8). Or, as Palmer puts it, it may be used “to express greater tentativeness”, as in the second of the two responses to the question in (447).

- (447) Who can play the tuba for us at next week’s concert?  
 a. John can.  
 b. John could.

Add to these two interpretations of the morphological past the post-state view of the perfect. And recall that, for statives in general, future orientation is not the default, but instead generally requires an anchor or a cue. Without that anchor or cue, we can expect present orientation to be the default. But since future orientation is required for semantic past tense on a proposition-level modal to make sense, a default to present orientation translates into a default interpretation of the morphological past tense on a proposition-level modal as indicating remoteness in logical space rather than in time. Hence, while semantic past on a proposition-level modal may be available in some cases, we should not expect it to be readily accessible.<sup>135</sup>

Turn now from the availability of the semantic past to the cancellability of the counterfactuality implication. My judgments do not align with Portner’s, though I can grant that (444), a true counterfactual, is perhaps less acceptable than (443). Where judgments differ, little can be said. However, I think my tense-based account can help to make sense of Portner’s judgments. For recall that there is competition between the two available interpretations of the morphological past tense on a modal. A speaker who explicitly signals that she intends the semantic past, via her use of temporal frame adverbials, is thus likely to reinforce its general effect. She will likely make the implication of counterfactuality harder to cancel, that is, insofar as she explicitly signaled that she did not intend the alternative, default interpretation that would not have generated this same implicature.

In fact, it has been argued that past tense does not always generate a counterfactual reading. In French, for example, the conditional mood has been claimed to also be required (cf. Laca 2012, 12–18; here Laca disagrees with Hacquard 2006, 77). Yet French is morphologically more transparent than English. If Laca’s claim is correct, this may thus provide indirect support for my contention that the competition in how the morphological past tense is interpreted in English plays a role in generating the counterfactual reading. Such cross-linguistic questions are matters for future research.

---

<sup>135</sup> Since bare eventives also impart future orientation, semantic past should be available for them too. In fact, back-shifting phrases like *at that point* do turn out to be compatible with morphologically past-tensed modals in cases where they are not compatible with their morphologically present-tensed counterparts, as (v) and (vi) show.

- (v) At that point, John {#may/might} still go on to win the race. Sadly, he didn’t.  
 (vi) At that point, John {#can/could} still {go on to} win the race. Sadly, he didn’t.

This suggests that maybe the morphological past is realizing the semantic past in these cases too, and is responsible for the implication of counterfactuality. In the case of a proposition-level modal like *might*, however, the periphrastic phrase *go on to* seems to be required for this to happen, something for which I have no account at the present moment.

### 7.6.3 A second sort of counterfactual reading

As I begin to close my account of counterfactual readings, I want to address the use of *might-have* sentences in action explanations, as in (448) and (449).

- (448) That concierge is *not* a spy. But he might have been one. And that is why I hid when he came into sight.
- (449) John did *not* steal the vase. But he might have. And that is why I arrested him.

Some have argued that such examples show there can be semantically past-tensed epistemic interpretations of modals, contrary to the received view (cf. von Stechow and Gillies 2008; Rullmann and Matthewson 2018). Yet in a previous chapter I suggested that the modal sentences in such examples should be interpreted counterfactually (cf. Chapter 4, Section 4.6). I am now in position to specify how this can be done.

Unlike Condoravdi's counterfactual readings of non-root modals, I do not think that the proposed counterfactual readings of (448) and (449) involve semantic past tense. Rather, they involve present tense and a slight but subtle change of topic from the preceding sentences that relates to the post-state view of the perfect. Now that the questions of whether the concierge is a spy, or whether John did steal the vase, are taken to be settled, the question is what would be compatible with the post-states of such eventualities. The post-state of the concierge's being a spy, for example, might be compatible with his moving stealthily, which in fact he does. And the post-state of John's stealing the vase might be compatible with his looking guilty, and with his fingerprints covering the hearth, which in fact they do.

These facts are just the sort of facts that would have led us in the first place to believe and claim that maybe the concierge is a spy, and that maybe John did steal the vase, and so to act as we did, given such beliefs. Thus, their compatibility with the relevant post-states can appropriately be cited in an explanation of our actions, as in (448) and (449).

In uttering (448) and (449), we know that the relevant eventualities have not obtained, and so neither have their post-states. But in changing the topic from (the possibility of) these eventualities to (the possibility of) their post-states, we are able to ignore these facts as we determine the value of the modal base. The change of topic is key to this ability. Without it, there is not this same ability, whence the incoherence of (450), unless the speaker is interpreted as having changed his mind.

- (450) That concierge is not currently a spy. #But he might currently be a spy.

Facts that have been incorporated into the common ground, and so presumably also into the value of the modal base, cannot just be excluded again by fiat. However, the change of topic from the eventuality to its post-state allows for the exclusion of such facts. This exclusion, motivated by the assumption that the speaker does not intend to be making a trivially false claim, is responsible for generating the counterfactuality of the reading.

#### 7.6.4 Counterfactual readings and modal bases

The second sort of counterfactual reading that I have just introduced I will call a *post-state* counterfactual, in contrast to Condoravdi's *open-future* counterfactuals. From a semantic standpoint, the post-state counterfactual in (451) is indistinguishable on my view from the presumed epistemic reading of the same sentence in (452).

- (451) John did *not* steal the vase. But he might have stolen it. And that is why I arrested him.  
(452) I'm pretty sure John didn't steal the vase. But he might have stolen it.

There is no difference in lexical entry, no difference in scope, no difference in tense to distinguish the two. Importantly, there is also no difference in the type of modal base, in the traditional sense of that term. Rather, both have what would traditionally be called circumstantial modal bases, from the values of which certain facts must be excluded in order to ensure non-triviality.

For the counterfactual reading, to be sure, the facts to be excluded from the value of the modal base will be facts that had previously been added. But this is a pragmatic difference, not a semantic one. The difference between the two readings is entirely in their pragmatic overtones and is dependent on the surrounding context. Or, as Laca nicely puts it, "the difference between the epistemic and the counterfactual construal hinges on the knowledge attributed to the speaker: counterfactuality arises under the further assumption that the speaker knows which way things went" (2012, 25).

Unlike post-state counterfactuals (and presumed epistemic readings of non-roots), open-future counterfactuals (and presumed metaphysical readings) likely do not require facts to be excluded from the value of the modal base to ensure non-triviality. This has led authors like Condoravdi (2002) to propose that non-root modals with future orientation have one type of modal base, a *metaphysical* one, while non-root modals with present and apparently past orientation have another, an *epistemic* one. However, there is no reason to posit multiple modal bases here either.

What unifies proposition-level modals for me, and distinguishes them from event-level ones, is that they are all used to attribute modal properties to propositions, rather than to individuals, situations, locations, etc. In particular, non-root possibility modals fundamentally express the prejacent's compatibility with the facts. To avoid claims about trivial compatibility, we must exclude any fact of the matter about a prejacent from the value of the modal base. If the world is not deterministic, this means we will have to adopt a strategy for interpreting present- (and what appear to be) past-oriented claims that we will not need for future-oriented ones. We can interpret future-oriented claims against the background of the total facts, whereas we cannot do the same for non-future-oriented ones. This difference in background reflects our knowledge of a fundamental difference in the structure of the future when compared to the past and the present, assuming the world is not deterministic. However, there is no reason to think that it should be encoded semantically as a difference in the type of modal base. Instead, it follows for free from completely pragmatic considerations. We can grant to all four readings of our proposition-level modals the same type of modal base, in the traditional sense of that term.

## 7.7 Conclusion

### 7.7.1 Looking back

*Might* appears to differ from *can* in accepting progressive and perfect complements. This difference has been attributed to the difference between epistemic possibility (allegedly expressed by *might*) and circumstantial possibility (allegedly expressed by *can*). While progressive and perfect morphology are supposed to orient possibilities to the present and the past, circumstantial possibilities are supposed to be oriented exclusively toward the future. Yet I have shown that *can* does accept progressive and perfect complements in some cases, and that acknowledged circumstantial modalities cannot always be analyzed as being oriented toward the future either.

I then proposed that we account for the difference between *can* and *might* in how freely they accept progressive and perfect complements through a distinction between event- and proposition-level modals (for inspiration, cf. Brennan 1993; Palmer 2001; Hacquard 2009). These types of modals differ in whether they introduce their own event variables, and so also in whether they are used to attribute *modal states* (which are the modal properties of individuals, etc.) or instead to attribute modal properties to propositions. And the progressive and the perfect turn out to denote relations that are helpful in determining propositions, but not generally in describing modal states.

A distinction between event- and proposition-level modals requires us to abandon the Kratzerian attempt to provide a uniform semantics across modal flavors. But the abandonment is not as significant as it seems. We may only have to move from one semantic entry to two, a theoretical cost justified by the expanded empirical coverage. Further research may also show the two semantic entries to play a role in the attested patterns of language acquisition and historical development of the modal flavors.

The different temporal properties of *might* and *can* thus provide no reason to think that the former is distinguished from the latter in expressing epistemic possibility. In fact, they may well provide reason to think the distinction runs along some dimension other than the epistemic/non-epistemic one, as the previous chapter also suggested.

This conclusion is reinforced in every chapter of the dissertation. When it comes to explaining the truth of modal claims (Chapter 2), to disagreeing with them (Chapter 4), or to reporting attitudes expressed with them (Chapter 5), phenomena that would otherwise be anomalous make perfect sense if we assume that *might* does not have an epistemic meaning, and that the relative uncertainty that is implied with epistemic uses of *might* is implied only indirectly (Chapter 3). *Might* does not modify the force of a speaker's commitment to the content of her speech act. Nor does it make that content truth-conditionally dependent on her or on anyone else's information state. Rather, it makes that content truth-conditionally dependent on sets of circumstances. The uncertainty that is implied in an epistemic use of *might* is not a semantic function of the modal but instead a pragmatic function of making a cooperative modal claim.

### 7.7.2 Looking forward

To bring this dissertation to a close, let me point to three avenues for future research.

First, my central argument in this dissertation has been that the type of possibility expressed by *might* is circumstantial, not epistemic. And just as many in the literature who have made the epistemic assumption have refrained from offering any particular analysis of epistemic possibility, so I have refrained in this dissertation from offering

any particular analysis of circumstantial possibility too (cf. Chapter 3, Section 3.3.1). The central argument of the dissertation is independent of the details of such analyses. Yet in other work I have begun to develop an analysis of the circumstantial possibility relation (cf. Harr 2019), and this is research that needs to be continued. There are two central issues here.

The first is that possibility, whether circumstantial or epistemic, likely cannot be analyzed exclusively in terms of compatibility with circumstances, or with contents of information states (cf. Bach 2011). Instead, as Kratzer (1981; 1991) proposes, and as (453) suggests, (im)possibility should likely be analyzed in terms of (in)compatibility with the conjunctions of such sets with sets of stereotypes or defaults.

(453) [*Context: John, who has a serious gambling addiction, is not to be found. A receipt for a bank withdrawal of 5000.00 USD is lying on the table.*]

- a. [S1] John might be at the Soaring Eagle Casino.
- b. [S2] You're wrong. He can't be. The casino closed down months ago.
- c. [S1] Yes, but last month it reopened for the evenings. So he might be.
- d. [S2] But they couldn't maintain that schedule for even a week, and now they've closed down for good. So he can't be.

For (453), the relevant stereotype or default would be that casinos that close down do not reopen. Yet since institutions that close down sometimes *do* reopen, John's being at the Soaring Eagle Casino is compatible with its having closed down months ago. However, there is an important sense in which it is *not* compatible with the relevant stereotype or default about such institutions.

The second issue is how to represent this more expansive relation of compatibility. Kratzer does so in the possible worlds framework, taking circumstances to be the value of a function that determines the set of worlds quantified over, and stereotypes to be the value of a function that orders those worlds according to their compliance with the stereotypical course of events. In doing so, she represents stereotypes by means of propositions, which can be true or false of the worlds in the modal base. However, while this method works well for simple stereotypes, I argue that it does not work well for what I call *complex stereotypes*: stereotypes that characterize what happens when some other stereotype does not hold. To account for the differential effects of such stereotypes, I argue, we may need to appeal to the resources of default logic (cf. Reiter 1980; Horty 2012; 2014). And this argument brings us squarely back to questions about the role and adequacy of the possible worlds framework that I have so far for the most part avoided as tangential to my main argument in this dissertation (cf. Chapter 3, Section 3.3.2 and Chapter 7, Sections 7.2.1 and 7.4.2).

A second avenue for future research concerns my distinguishing the possibilities expressed by *might* from the potentials expressed by *can* in terms of a distinction between proposition and event modality. Since *might* is used to express that some proposition is possibly true, its *alethic* flavor of modality can only be expressed by a proposition-level modal. And since *can* is used to express that some individual, etc., has the potential to develop in a certain way, or is in a certain *modal state*, its *dynamic* flavor of modality can only be expressed by an event-level modal.

Yet for other flavors of modality, it is an open question whether they are able to be expressed only by one or the other of these two levels of modals. Brennan (1993), for example, distinguished *ought-to-be* deontics from *ought-to-do* ones, which could correspond to a distinction between a *proposition* being said to be obligatory and an *individual* being said to be in a state of obligation. Hacquard (2006; 2009) further argued, on the basis of actuality entailments, that while the latter category of deontics occur within the scope of aspect, the former do not. If correct, this would follow straightforwardly from the different semantic types I assign to event- and proposition-level modals. However, I leave it as a matter for future research whether these other modal flavors are in fact expressed by both.

A third avenue for future research, perhaps the most interesting of all, is whether there are languages in which modal auxiliaries express epistemic flavors of modality, or in which flexible modals do. In English, if I am correct, there are no clear instances of either. What seemed on first analysis to be epistemic flavors, in such cases, turned out to be better analyzed as alethic flavors of circumstantial modality. Say that other languages also turn out to be like English. Why should this be so? Is this due to a grammatical constraint? Or is it due to a sort of conceptual constraint on polysemy? One possibility, at least for modal auxiliaries, is that they are of the wrong type to have an epistemic meaning. On this chapter's proposal, their complements are not propositions, not even temporal ones, but epistemic meanings quite plausibly require propositional complements. If, cross-linguistically, modal auxiliaries never expressed epistemic flavors of modality, this proposal could make sense of that fact. But what about modal adjective like *possible*, which do take propositional complements? Are such modals ever flexible between epistemic and alethic flavors of modality? Or are they ever just flexible more generally between epistemic and non-epistemic flavors? If not, why not? These are interesting and puzzling questions that require additional research. Yet whatever these avenues of research turn up, remember this:

*“Might” is never used to make epistemic claims.*

## Bibliography

- Abusch, D. 2012. Circumstantial and temporal dependence in counterfactual modals. *Natural Language Semantics* 20: 273–97.
- Anand, P. & Hacquard, V. 2009. Epistemics with attitude. *Semantics and Linguistic Theory* 18: 37–54.
- \_\_\_\_\_. 2013. Epistemics and attitudes. *Semantics and Pragmatics* 6: 1–59.
- Asher, N. 2011. *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press.
- Austin, J. L. 1962. *How to do things with words*. Oxford: Oxford University Press.
- Bach, K. 1994a. Conversational implicature. *Mind & Language* 9: 124–62.
- \_\_\_\_\_. 1994b. Semantic slack. In S. Tsohatzidis (Ed.), *Foundations of speech act theory: Philosophical and linguistic perspectives* (pp. 267–91). London: Routledge.
- \_\_\_\_\_. 2000. Quantification, qualification and context: A reply to Stanley and Szabó. *Mind & Language* 15: 262–83.
- \_\_\_\_\_. 2001. You don't say? *Synthese* 128: 15–44.
- \_\_\_\_\_. 2005. Context ex machina. In Z. Szabó (Ed.), *Semantics versus pragmatics* (pp. 15–44). Oxford: Clarendon.
- \_\_\_\_\_. 2011. Perspectives on possibilities: Contextualism, relativism, or what? In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 19–59). Oxford: Oxford University Press.
- Barwise, J. 1981. Scenes and other situations. *Journal of Philosophy* 78: 369–97.
- Barwise, J. & Perry, J. 1983. *Situations and attitudes*. Cambridge, MA: MIT Press.
- Beddor, B. & Egan, A. 2018. Might do better: Flexible relativism and the QUD. *Semantics and Pragmatics* 11: 1–44.
- Bennett, M. & Partee, B. 1972. *Toward the logic of tense and aspect in English*. Bloomington, IN: Indiana University Linguistics Club.
- Borg, E. 2007. Minimalism versus contextualism in semantics. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* (pp. 546–71). Oxford: Oxford University Press.
- Braun, D. 2012. An invariantist theory of “might” might be right. *Linguistics and Philosophy* 35: 461–89.
- \_\_\_\_\_. 2013. Contextualism about “might” and says-that ascriptions. *Philosophical Studies* 164: 485–511.
- Brennan, V. 1993. *Root and epistemic modal auxiliary verbs*. PhD diss., University of Massachusetts, Amherst.
- Cappelen, H. & Lepore, E. 2002. Indexicality, binding, anaphora and a priori truth. *Analysis* 62: 271–81.
- \_\_\_\_\_. 2005. Radical and moderate pragmatics: Does meaning determine truth conditions? In Z. Szabó (Ed.), *Semantics versus pragmatics* (pp. 45–71). Oxford: Clarendon.
- \_\_\_\_\_. 2007. The myth of unarticulated constituents. In M. O'Rourke & C. Washington (Eds.), *Situating semantics: Essays on the philosophy of John Perry* (pp. 199–214). Cambridge, MA: MIT Press.

- Carston, R. 1988. Implicature, explicature, and truth-theoretic semantics. In R. Kempson (Ed.), *Mental representations: The interface between language and reality* (pp. 155–81). Cambridge: Cambridge University Press.
- \_\_\_\_\_. 2002. *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.
- \_\_\_\_\_. 2004. Explicature and semantics. In S. Davis & B. Gillon (Eds.), *Semantics: A reader* (pp. 817–45). Oxford: Oxford University Press.
- Chalmers, D. 2006. Two-dimensional semantics. In E. Lepore & B. Smith (Eds.), *The Oxford handbook of philosophy of language* (pp. 574–606). Oxford: Clarendon.
- Condoravdi, C. 2002. Temporal interpretation of modals: Modals for the present and for the past. In D. Beaver, S. Kaufmann, B. Clark, & L. Casillas (Eds.), *The construction of meaning* (pp. 59–88). Stanford: CSLI.
- Condoravdi, C. & Gawron, J. 1996. The context-dependency of implicit arguments. In M. Kanazawa, C. Pinon, & H. de Swart (Eds.), *Quantifiers, deduction, and context* (pp. 1–32). Stanford: CSLI Publications.
- Copley, B. 2008. Temporal orientation in conditionals. In J. Guéron & J. Lecarme (Eds.), *Time and modality* (pp. 59–78). Berlin: Kluwer.
- Crabill, J. 2013. Suppose Yalcin is wrong about epistemic modals. *Philosophical Studies* 162: 625–35.
- DeRose, K. 1991. Epistemic possibilities. *Philosophical Review* 100: 581–605.
- \_\_\_\_\_. 1998. Simple 'might's, indicative possibilities and the open future. *Philosophical Quarterly* 48: 67–82.
- Dietz, R. 2008. Epistemic modals and correct disagreement. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth* (pp. 239–62). Oxford: Oxford University Press.
- Dowell, J. 2011. A flexible contextualist account of epistemic modals. *Philosophers' Imprint* 11: 1–25.
- Drubig, H. 2001. On the syntactic form of epistemic modality. Unpublished manuscript. University of Tübingen.
- Eckardt, R. 2015. *The semantics of free indirect discourse*. Leiden: Brill.
- Egan, A. 2007. Epistemic modals, relativism and assertion. *Philosophical Studies* 133: 1–22.
- Egan, A., Hawthorne, J. & Weatherson, B. 2005. Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 131–68). Oxford: Clarendon.
- Elbourne, P. 2008. The argument from binding. *Philosophical Perspectives* 22: 89–110.
- \_\_\_\_\_. 2013. *Definite descriptions*. Oxford: Oxford University Press.
- Faller, M. 2012. Evidential scalar implicatures. *Linguistics and Philosophy* 35: 285–312.
- Fillmore, C. J. 1986. Pragmatically controlled zero anaphora. *Annual Proceedings of the Berkeley Linguistics Society* 12: 95–107.
- Frege, G. 1948. On sense and reference. *Philosophical Review* 57: 207–30.
- \_\_\_\_\_. 1956. The thought: A logical inquiry. *Mind* 65: 289–311.
- Gagnon, M. & Wellwood, A. 2011. Distributivity and modality: Where “each” may go, “every” can't follow. *Semantics and Linguistic Theory* 21: 39–55.

- Gamut, L. T. F. 1991. *Logic, language, and meaning*. Vol. 2, Intensional logic and logical grammar. Chicago: University of Chicago Press.
- Geach, P. 1962. *Reference and generality*. Ithaca, NY: Cornell University Press.
- Geurts, B. 2012. Accessibility and anaphora. In P. Portner, C. Maienborn, & K. von Heusinger (Eds.), *Handbook of Semantics* (pp. 1988–2011). Berlin: de Gruyter.
- Goodhue, D. 2017. Must  $\phi$  is felicitous only if  $\phi$  is not known. *Semantics and Pragmatics* 10: 1–27.
- Grice, P. 1989. *Logic and conversation*. In *Studies in the ways of words* (pp. 22–40). Cambridge, MA: Harvard University Press.
- Groenendijk, J. & Stokhof, M. 1975. Modality and conversational information. *Theoretical Linguistics* 2: 61–112.
- Hacking, I. 1967. Possibility. *Philosophical Review* 76: 143–68.
- Hacquard, V. 2006. *Aspects of modality*. PhD diss., Massachusetts Institute of Technology.
- \_\_\_\_\_. 2009. On the interaction of aspect and modal auxiliaries. *Linguistics and Philosophy* 32: 279–315.
- \_\_\_\_\_. 2010. On the event relativity of modal auxiliaries. *Natural Language Semantics* 18: 79–114.
- \_\_\_\_\_. 2011. Modality. In K. von Heusinger, C. Maienborn, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (pp. 1484–1515). Berlin: de Gruyter Mouton.
- \_\_\_\_\_. 2013. The grammatical category of modality. In M. Aloni, M. Franke, & F. Roelofsen (Eds.), *Proceedings of the 19th Amsterdam Colloquium* (pp. 19–26). Amsterdam: University of Amsterdam.
- \_\_\_\_\_. 2016. Modals: Meaning categories? In J. Błaszczak, A. Giannakidou, D. Klimek-Jankowska, & K. Migdalski (Eds.), *Mood, aspect, modality revisited: New answers to old questions* (pp. 45–74). Chicago: University of Chicago Press.
- Hacquard, V. & Wellwood, A. 2012. Embedding epistemic modals in English: A corpus-based study. *Semantics and Pragmatics* 5: 1–29.
- Halliday, M. A. K. 1970. Functional diversity in language as seen from a consideration of modality and mood in English. *Foundations of Language* 6: 322–61.
- Harr, Q. 2014. *Counterfactuals, epistemic modals, and semantic content*. Unpublished manuscript. University of Maryland.
- \_\_\_\_\_. 2019. *A non-monotonic interpretation of “might.”* Unpublished manuscript. University of Maryland.
- Heim, I. 1983. File change semantics and the familiarity theory of definiteness. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, use, and interpretation of language* (pp. 164–89). Berlin: Walter de Gruyter.
- Hintikka, J. 1962. *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, NY: Cornell University Press.
- Hofmann, T. 1966. Past tense replacement and the modal system. In A. Oettinger (Ed.), *Mathematical linguistics and automatic translation*. Cambridge, MA: Harvard Computational Laboratory, Harvard University.

- Homer, V. 2010. French modals and perfective: A case of aspectual coercion. In M. Washburn (Ed.), *Proceedings of the West Coast Conference on Formal Linguistics 28* (pp. 106–14). Somerville, MA: Cascadia Proceedings Project.
- Horty, J. 2012. *Reasons as defaults*. Oxford: Oxford University Press.
- \_\_\_\_\_. 2014. Deontic modals: Why abandon the classical semantics? *Pacific Philosophical Quarterly* 95: 424–60.
- Huvenes, T. 2015. Epistemic modals and credal disagreement. *Philosophical Studies* 172: 987–1011.
- Jenkins, L. 1972. *Modality in English syntax*. PhD diss., Massachusetts Institute of Technology.
- Kamp, H. 1981. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, & M. Stokhof (Eds.), *Formal methods in the study of language* (pp. 277–322). Amsterdam: Mathematisch Centrum.
- Kamp, H. & Reyle, U. 1993. *From discourse to logic*. Dordrecht: Kluwer.
- Kaplan, D. 1989. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.
- Kennedy, C. & Willer, M. 2016. Subjective attitudes and counterstance contingency. *Semantics and Linguistic Theory* 26: 1–21.
- Khoo, J. 2015. Modal Disagreements. *Inquiry* 58: 511–34.
- King, J. & Lewis, K. 2018. Anaphora. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*, Fa. 2018. Metaphysics Research Lab, Stanford University.
- Klecha, P. 2016. Modality and embedded temporal operators. *Semantics and Pragmatics* 9: 1–55.
- Knobe, J. & Yalcin, S. 2014. Epistemic modals and context: Experimental data. *Semantics and Pragmatics* 7: 1–21.
- Kölbel, M. 2002. *Truth without objectivity*. London: Routledge.
- Kratzer, A. 1977. What “must” and “can” must and can mean. *Linguistics and Philosophy* 1: 337–55.
- \_\_\_\_\_. 1981. The notional category of modality. In P. Portner & B. Partee (Eds.), *Formal semantics: The essential readings* (pp. 289–323). Oxford: Blackwell.
- \_\_\_\_\_. 1991. Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An international handbook of contemporary research* (pp. 639–50). Berlin: de Gruyter.
- \_\_\_\_\_. 2011. What “can” can mean. Talk given at SALT 21, New Brunswick, NJ, May.
- \_\_\_\_\_. 2012. *Modals and conditionals*. Oxford: Oxford University Press.
- \_\_\_\_\_. 2013. Modality for the 21st century. In S. Anderson, J. Moeschler, & F. Reboul (Eds.), *The language cognition interface: Pre-proceedings for the 19th International Congress of Linguists in Geneva* (pp. 179–99). Geneva-Paris: Librairie Droz.
- \_\_\_\_\_. 2019. Situations in natural language semantics. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*, Sp. 2019. Metaphysics Research Lab, Stanford University.
- Kripke, S. 1980. *Naming and necessity*. Cambridge, MA: Harvard University Press.

- Laca, B. 2012. On modal tenses and tensed modals. In C. Nishida & C. Russi (Eds.), *Building a bridge between linguistic communities of the old and the new world: Current research in tense, aspect, mood and modality* (pp. 163–98). Brill.
- Lasersohn, P. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28: 643–86.
- \_\_\_\_\_. 2017. *Subjectivity and perspective in truth-theoretic semantics*. Oxford: Oxford University Press.
- Lewis, D. 1976. The paradoxes of time travel. *American Philosophical Quarterly* 13: 145–52.
- \_\_\_\_\_. 1986. *On the plurality of worlds*. Oxford: Blackwell.
- \_\_\_\_\_. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74: 549–67.
- Lyons, J. 1977. *Semantics*. Vol. 2. Cambridge: Cambridge University Press.
- MacFarlane, J. 2003. Future contingents and relative truth. *Philosophical Quarterly* 53: 321–36.
- \_\_\_\_\_. 2011. Epistemic modals are assessment-sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 144–78). Oxford: Oxford University Press.
- \_\_\_\_\_. 2014. *Assessment sensitivity: Relative truth and its applications*. Oxford: Clarendon.
- Mandelkern, M. 2017. *Coordination in conversation*. PhD diss., Massachusetts Institute of Technology.
- Mari, A. & Martin, F. 2007. Tense, abilities, and actuality entailment. In M. Aloni, P. Dekker, & F. Roelofsen (Eds.), *Proceedings of the 16th Amsterdam Colloquium* (pp. 151–56). Amsterdam: ILLC/Department of Philosophy, University of Amsterdam.
- Matthewson, L. 2012. On the (non-)future orientation of modals. *Proceedings of Sinn und Bedeutung* 16: 431–46.
- \_\_\_\_\_. 2015. Evidential restrictions on epistemic modals. In L. Alonso-Ovalle & P. Menéndez-Benito (Eds.), *Epistemic indefinites: Exploring modality beyond the verbal domain* (pp. 141–60). Oxford: Oxford University Press.
- \_\_\_\_\_. 2016. Modality. In M. Aloni & P. Dekker (Eds.), *Cambridge handbook of formal semantics* (pp. 525–59). Cambridge: Cambridge University Press.
- Matthewson, L., Rullmann, H. & Davis, H. 2008. Evidentials as epistemic modals: Evidence from St’át’imcets. In J. van Craenebroeck & J. Rooryck (Eds.), *Linguistic Variation Yearbook*. Vol. 7 (pp. 201–54). Amsterdam: John Benjamins.
- McCawley, J. 1971. Tense and time reference in English. In C. J. Fillmore & T. Langendoen (Eds.), *Studies in Linguistic Semantics* (pp. 97–113). New York: Holt, Rinehart, and Winston.
- Moltmann, F. 2018. An object-based truthmaker semantics for modals. *Philosophical Issues* 28: 255–88.
- Montminy, M. 2012. Epistemic modals and indirect weak suggestives. *Dialectica* 66: 583–606.
- Moore, G. E. 1962. *Commonplace book: 1919-1953*. Edited by Casimir Lewy. C. Lewy (Ed.), . London: Allen and Unwin.

- Nauze, F. 2008. Modality in typological perspective. PhD diss., Institute for Logic, Language and Computation.
- Neale, S. 2000. On being explicit: Comments on Stanley and Szabo, and on Bach. *Mind & Language* 15: 284–94.
- Ney, J. 1978. Necessity, hypothesis and prohibition in the English modals “must” and “should.” *Journal of English Linguistics* 12: 38–49.
- Ninan, D. 2018. Relational semantics and domain semantics for epistemic modals. *Journal of Philosophical Logic* 47: 1–16.
- Nunberg, G. 1979. The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy* 3: 143–84.
- Ogihara, T. 2007. Tense and aspect in truth-conditional semantics. *Lingua* 117: 392–418.
- Palmer, F. R. 1979. *Modality and the English modals*. London: Longman.
- \_\_\_\_\_. 1986. *Mood and modality*. Cambridge: Cambridge University Press.
- \_\_\_\_\_. 2001. *Mood and modality*. 2nd ed. Cambridge: Cambridge University Press.
- Papafragou, A. 2006. Epistemic modality and truth conditions. *Lingua* 116: 1688–1702.
- Parsons, T. 1990. *Events in the semantics of English: A study in subatomic semantics*. Cambridge, MA: MIT Press.
- Partee, B. 1989. Binding implicit variables in quantified contexts. In C. Wiltshire, B. Music, & R. Graczyk (Eds.), *Papers from CLS 25* (pp. 342–56). Chicago: Chicago Linguistics Society.
- Pietroski, P. M. 2005. Meaning before truth. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 253–300). Oxford: Clarendon.
- \_\_\_\_\_. 2018. *Conjoining meanings: Semantics without truth values*. Oxford: Oxford University Press.
- Portner, P. 2009. *Modality*. Oxford: Oxford University Press.
- Price, H. 1983. Does “probably” modify sense? *Australasian Journal of Philosophy* 61: 396–408.
- Pustejovsky, J. & Boguraev, B., eds. 1996. *Lexical semantics: The problem of polysemy*. Oxford: Clarendon.
- Ravin, Y. & Leacock, C., eds. 2000. *Polysemy: Theoretical and computational approaches*. Oxford: Oxford University Press.
- Recanati, F. 2001. What is said. *Synthese* 128: 75–91.
- \_\_\_\_\_. 2002. Unarticulated constituents. *Linguistics and Philosophy* 25: 299–345.
- \_\_\_\_\_. 2004. *Literal meaning*. Cambridge: Cambridge University Press.
- \_\_\_\_\_. 2007. *Perspectival thought: A plea for (moderate) relativism*. Oxford: Oxford University Press.
- Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13: 81–132.
- Richard, M. 2015. What is disagreement? In *Truth and truth bearers*. Vol. 2 of *Meaning in context* (pp. 82–114). Oxford: Oxford University Press.
- Rullmann, H. & Matthewson, L. 2018. Towards a theory of modal-temporal interaction. *Language* 94: 281–331.

- Schaffer, J. 2011. Perspective in taste predicates and epistemic modals. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 179–226). Oxford: Oxford University Press.
- Schnieder, B. 2010. Expressivism concerning epistemic modals. *Philosophical Quarterly* 60: 601–15.
- Sennet, A. 2008. The binding argument and pragmatic enrichment, or, Why philosophers care even more than weathermen about “raining.” *Philosophy Compass* 3: 135–57.
- Silk, A. 2017. How to embed an epistemic modal: Attitude problems and other defects of character. *Philosophical Studies* 174: 1773–99.
- Sperber, D. & Wilson, D. 1986. *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Stalnaker, R. 1978. Assertion. In P. Cole (Ed.), *Syntax and semantics*. Vol. 9, *Pragmatics* (pp. 315–32). New York: Academic Press.
- Stanley, J. 2000. Context and logical form. *Linguistics and Philosophy* 23: 391–434.
- Stephenson, T. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30: 487–525.
- Swanson, E. 2006. *Interactions with context*. PhD diss., Massachusetts Institute of Technology.
- \_\_\_\_\_. 2011. How not to theorize about the language of subjective uncertainty. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 249–69). Oxford: Oxford University Press.
- Sweetser, E. 1990. *From etymology to pragmatics: Metaphorical and cultural aspects of semantic structure*. Cambridge: Cambridge University Press.
- Tancredi, C. 2007. *A multi-model modal theory of I-semantics*. Unpublished manuscript. University of Tokyo.
- Taylor, B. 1977. Tense and continuity. *Linguistics and Philosophy* 1: 199–220.
- Teller, P. 1972. Epistemic possibility. *Philosophia* 2: 303–20.
- Thalberg, I. 1972. How is ability related to performance? In *Enigmas of agency: Studies in the philosophy of human action* (pp. 115–42). London: Allen and Unwin.
- Thomas, G. 2014. Circumstantial modality and the diversity condition. *Proceedings of Sinn und Bedeutung* 18: 433–50.
- \_\_\_\_\_. 2017. *Revisiting the temporal orientation of circumstantial modals*. Unpublished manuscript. University of Toronto.
- Traugott, E. 1989. On the rise of epistemic meanings in English: An example of subjectification in semantic change. *Language* 65: 31–55.
- Vetter, B. 2013. “Can” without possible worlds: Semantics for anti-Humeans. *Philosophers’ Imprint* 13: 1–27.
- Viebahn, E. & Vetter, B. 2016. How many meanings for “may”? The case for modal polysemy. *Philosophers’ Imprint* 16: 1–39.
- Vlach, F. 1981. The semantics of the progressive. In P. Tedeschi & A. Zaenen (Eds.), *Syntax and semantics*. Vol. 14, *Tense and aspect*. New York: Academic Press.
- von Fintel, K. & Gillies, A. 2007. An opinionated guide to epistemic modality. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. Vol. 2 (pp. 32–62). Oxford: Oxford University Press.

- \_\_\_\_\_. 2008. CIA leaks. *Philosophical Review* 117: 77–98.
- \_\_\_\_\_. 2010. “Must”... stay... strong! *Natural Language Semantics* 18: 351–83.
- \_\_\_\_\_. 2011. “Might” made right. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 108–30). Oxford: Oxford University Press.
- von Fintel, K. & Heim, I. 2011. *Intensional semantics*. Unpublished manuscript. Massachusetts Institute of Technology.
- Weatherson, B. & Egan, A. 2011. Introduction: Epistemic modals and epistemic modality. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 1–18). Oxford: Oxford University Press.
- Werner, T. 2006. Future and non-future modal sentences. *Natural Language Semantics* 14: 235–55.
- Westmoreland, R. 1998. *Information and intonation in natural language modality*. PhD diss., Indiana University.
- White, A. 1975. *Modal thinking*. Ithaca, NY: Cornell University Press.
- Willer, M. 2013. Dynamics of epistemic modality. *Philosophical Review* 122: 45–92.
- \_\_\_\_\_. 2015. An update on epistemic modals. *Journal of Philosophical Logic* 44: 835–49.
- Williams, A. 2012. Null complement anaphors as definite descriptions. *Semantics and Linguistic Theory* 22: 125–45.
- Wright, C. 2007. New age relativism and epistemic possibility: The question of evidence. *Philosophical Issues* 17: 262–83.
- \_\_\_\_\_. 2008. Relativism about truth itself: Haphazard thoughts about the very idea. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth* (pp. 157–85). Oxford: Oxford University Press.
- Yablo, S. 2011. A problem about permission and possibility. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 270–94). Oxford: Oxford University Press.
- Yalcin, S. 2007. Epistemic modals. *Mind* 116: 983–1026.
- \_\_\_\_\_. 2011. Nonfactualism about epistemic modality. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 295–332). Oxford: Oxford University Press.
- Yanovich, I. 2014. Standard contextualism strikes back. *Journal of Semantics* 31: 67–114.