

© 2016

Una Stojnić

ALL RIGHTS RESERVED

CONTEXT-SENSITIVITY IN A COHERENT DISCOURSE

BY

UNA STOJNIĆ

A dissertation submitted to the
Graduate School—New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
Graduate Program in Philosophy

Written under the direction of
Jeffrey C. King and Ernest Lepore
and approved by

New Brunswick, New Jersey

October, 2016

ABSTRACT OF THE DISSERTATION

Context-sensitivity in a Coherent Discourse

By Una Stojnić

Dissertation Directors:

Jeffrey C. King and Ernest Lepore

What I communicate with ‘Give me that’, pointing at a book, differs from what I communicate by it pointing at a cup. Your actions and my expectations likewise differ in these two cases. At the same time, the referent of any particular use of ‘that’ is typically unambiguous and recovered effortlessly. What determines the referent of ‘that’ and which resources permit us to recover it so easily? Though everyone agrees that what we can communicate is constrained by grammar, most believe that the role of grammar is very limited, and that interpretation largely relies on general reasoning about the communicative situations and intentions of the speaker. Philosophers frequently identify such (purported) context-sensitivity within philosophically interesting expressions like ‘know’, or ‘good’, and appeal to it to shed light on problems involving the concepts these expressions denote; but they also assume that these expressions get their values in context by way of common-sense reasoning about speaker intentions. This has led to many radical conclusions. To take just one concrete example, the behavior of context-sensitive expressions describing possibilities and necessity (‘must’ and ‘might’) *prima facie* gives rise to apparent failures of classical patterns of inference, like *Modus Ponens*, which has led many philosophers throughout the history of analytic philosophy, operating under the traditional assumptions about the resolution of context-sensitivity, to conclude that there is a deep incompatibility between classical logic and natural language.

Against this tradition, I argue, drawing on resources from philosophy, linguistics, mathematical logic and computer science, that the reason we can interpret context-sensitivity so effortlessly is

because grammar itself is much more subtle and pervasive than has been assumed, and that resolution of context-sensitivity is entirely a matter of linguistic convention. Thus, linguistic rules render a particular object prominent in a given context, and this is what determines what ‘that’ picks out in that context. Moreover, in recovering this referent it is this narrow set of linguistic cues that we exploit.

The conventions that I argue govern the resolution of context-sensitivity have gone unnoticed because their principal domains are entire discourses and not just their constituent words and sentences. While it is universally accepted that the way individual sentences are constructed depends on conventions of syntax and semantics, which specify the rules by which individual expressions combine, I argue for rules–discourse conventions–that specify how individual sentences combine to form a discourse. These conventions govern how speakers organize utterances into larger units that address particular topics and answer questions about them; I argue these rules determine the resolution of context-sensitivity. The move to investigating discourses has far reaching consequences: I show that a host of contextualist arguments that resort to context-dependence rest on a flawed conception context, particularly in the debate about the relation between the natural language and logic, and argue that context-dependence will have to be invoked quite differently than has become customary.

Acknowledgements

First and foremost, I owe gratitude to my dissertation advisers: Jeffrey C. King and Ernie Lepore. Seminar after seminar, discussion after discussion, draft after draft, they have influenced my thinking about language, in ways that were critical for the development of this dissertation. I'm indebted to them for their generous feedback, and for their support.

It was during Jeff's seminar I took in my first semester of graduate school that I first started thinking about context. Ever since, Jeff's sharp critical sense for detail has forced me to make my ideas clearer and arguments sharper. When we agreed it was great; when we disagreed it was even better, challenging me in ways that were crucial for the shaping of the present project. It is indeed hard to overstate how much I learned from Jeff.

Ernie's guidance has been essential every step of the way. His generous feedback was critical in refining some of the key aspects of my project, including my thinking about the fundamental questions concerning meaning and context, and the interaction between thought, communication, and action. Countless conversations with Ernie influenced my thinking more broadly: it was he who reminded me not to miss the forest for the trees, and to pursue those questions for which I can in principle see the shape of an answer. Moreover, Ernie is also the driving force behind the interdisciplinary semantics group at Rutgers that provided the uniquely inspiring and stimulating environment in which this dissertation took form.

Apart from Jeff and Ernie, I owe special gratitude to Matthew Stone, who directed my Cognitive Science Project, and who, along my advisers, also served as a mentor in every sense of the term. He has helped me learn how to do interdisciplinary work, navigating and bridging gaps between philosophy, linguistics, logic, computer science and psychology, in ways that let me exploit the tools of each of these disciplines to gain insights that transcend their boundaries. He has also taught me how (and when) to pursue formal rigor in support of more fundamental ideas. Finally, Matthew's patience and encouragement to push forward even (or especially) at times when details

seemed hairy, drafts got long, conversations lengthy and waters murky, were critical for the success of the project.

Ernie and Matthew have also been my collaborators. I first started thinking more closely about some of the ideas that materialized in this dissertation in the summer of 2012. I was working as a research assistant in Matthew's lab, on the project 'Collaborative Reference in Open Domains', (supported by NSF IIS-1017811). Our goal was to build meaning representations for situated demonstrative utterances. At the same time, through discussions and collaboration with Ernie, we started thinking about the logical forms of stretches of discourse that extend beyond the boundaries of a sentence. From these two lines of thought several co-authored papers were born, one of which—'Discourse and Logical Form: Pronouns, Attention and Coherence'—became the second chapter of this dissertation.

I owe gratitude to the rest of my committee. Liz Camp has offered long and detailed comments on numerous versions of this manuscript. Discussions with her have led me to position my account more clearly against the alternatives. Andy Egan has provided insightful comments on various drafts, which led to important clarifications and refinements, especially in Chapter 3. Thony Gillies has offered pointed criticism, which helped keep the details in check. And the project has benefited greatly from the discussions with John Hawthorne, including during my research visit under his supervision at the University of Oxford in 2013. I owe him gratitude for these stimulating discussions, and for making philosophy fun.

In addition, I owe thanks to all the members of the interdisciplinary semantics group at Rutgers, and to fellow students in the language group: Alex Anthony, Sam Carter, Nico Kirk-Giannini, Peter van Elswyk, Simon Goldstein, Verónica Gómez, Lucy Jordan, Kelsey Laity-D'Agostino, Jess Law, Ang Li, and Morgan Moyer.

I also owe thanks to Gilbert Harman, and Jim Pryor, who both read various drafts of my work, and provided insightful and detailed comments and discussion, to Alex Lascarides, for stimulating conversations during my research visit at the University of Edinburgh, Institute for Language, Cognition and Computation, and to Daniel Altshuler, with whom I pursued a further application of the account of context-sensitivity and prominence developed in this dissertation to the semantics of the temporal indexical 'now'.

I have benefited from discussions with a number of colleagues. I would like to thank Luvell Anderson, Josh Armstrong, Nicholas Asher, Chris Barker, David Beaver, Maria Bittner, Justin Bledin, David Braun, Fabrizio Cariani, Nate Charlow, Simon Charlow, Jonathan Cohen, Sam Cumming, Cian Dorr, Branden Fitelson, Graeme Forbes, Michael Glanzberg, Jeremy Goodman, Gabe Greenberg, Alan Hájek, Dan Harris, Andy Kehler, Angelika Kratzer, Karen Lewis, Barry Loewer, Robert May, Brian McLaughlin, Matt Moss, Sarah Murray, Stephen Neale, Barbara Partee, Paul Pietroski, Brian Rabern, Hannah Rohde, Jessica Rett, David Rosenthal, Daniel Rothschild, Craige Roberts, Jonathan Schaffer, Philippe Schlenker, Ted Sider, Adam Sennet, Mandy Simons, Jason Stanley, Will Starr, Gala Stojnić, Kristen Syrett, Rich Thomason, Brian Weatherson, Timothy Williamson, Seth Yalcin, and numerous others whom I must have failed to mention.

I thank my professors at the University of Belgrade for providing first guidance in philosophy, and encouraging me to pursue graduate education. Thanks especially to Miloš Arsenijević, Mašan Bogdanovski, Vojislav Božićković, Kosta Došen, Vladan Đorđević, Leon Kojen and Živan Lazović.

Last, but definitely not the least, I thank my family. To my parents, Jasna and Danko, and my sisters, Aneta and Gala: thank you for everything.

Dedication

Dedicated, with unconditional love and admiration, to Kiká and Mecha.

Table of Contents

| | |
|---|--------|
| Abstract | ii |
| Acknowledgements | iv |
| Dedication | vii |
| Table of Contents | viii |
| Introduction | 1 |
| 1. Discourse, Context and Coherence: The Grammar of Prominence | 5 |
| 1.1. Introduction | 5 |
| 1.2. Coherence | 9 |
| 1.3. Coherence and Context-Sensitivity | 12 |
| 1.4. Coherence and Grammar | 21 |
| 1.5. Discourse Structure, Information Structure and Grammar | 25 |
| 1.6. Conclusion | 30 |
| 2. Discourse and Logical Form: Pronouns, Attention and Coherence | 32 |
| 2.1. Introduction | 32 |
| 2.2. Attention | 36 |
| 2.3. Refinements | 43 |
| 2.4. Coherence | 48 |
| 2.5. Conclusion | 57 |
| 3. One's <i>Modus Ponens</i>: Modality, Coherence and Logic | 59 |
| 3.1. Introduction | 59 |
| 3.2. Part I | 61 |

| | | |
|-----------|--|------------|
| 3.2.1. | The Challenge to <i>Modus Tollens</i> | 61 |
| 3.2.2. | Modals and Pronouns | 66 |
| 3.2.3. | Coherence | 67 |
| 3.2.4. | Modals as Pronouns | 71 |
| 3.2.5. | Yalcin’s Counterexample Explained Away | 78 |
| 3.2.6. | Conclusion to Part I | 80 |
| 3.3. | Part II: A New Semantics for Modality | 81 |
| 3.3.1. | Anaphoricity and Truth-Conditions | 81 |
| 3.3.2. | Truth-conditions and the Dynamics of Context-change | 85 |
| 3.3.3. | Conclusion to Part II | 100 |
| 4. | Content in a Dynamic Context | 101 |
| 4.1. | Introduction | 101 |
| 4.2. | Modality and Propositional Content | 103 |
| 4.3. | Modals and Context-Change | 106 |
| 4.4. | Anaphora and Modality | 113 |
| 4.5. | Content in a Dynamic Context | 126 |
| 4.6. | Conclusion | 132 |
| | Appendices | 134 |
| A. | Formal Definitions for the Attention-Coherence Approach | 135 |
| A.1. | Appendix: Formal definitions | 135 |
| B. | Language for Modality with Coherence | 142 |
| B.1. | Appendix: Formal definitions | 142 |
| B.1.1. | Syntax: | 142 |
| B.1.2. | Models: | 144 |
| B.1.2.1. | Truth-conditional contributions of modals and conditionals: . . . | 144 |
| B.1.3. | Describing operations on stacks (sequences of worlds) and sets of stacks. | 145 |

| | |
|---|-----|
| B.1.4. Semantics: | 146 |
| B.1.5. Truth, validity, entailment. | 148 |
| B.1.6. Modus Ponens (MP) and Modus Tollens (MT): | 149 |
| B.1.6.1. MT | 149 |
| B.1.6.2. MP | 151 |
| B.1.7. Relation between the dynamic (B.1.7.2) and classical (B.1.7.3) interpretations | 153 |
| B.1.7.1. Abstract Logical Forms (ALFs) for a Fragment of English: . . . | 153 |
| B.1.7.2. Dynamic Interpretation: | 154 |
| B.1.7.2.1. Dynamic Translations | 154 |
| B.1.7.3. Classical Interpretation | 155 |
| B.1.7.3.1. Classical Semantics: | 155 |
| B.1.7.3.2. Classical Translation: | 156 |
| B.1.7.4. Proof. | 157 |
| B.1.8. Yalcin’s Counterexample: | 158 |
| Bibliography | 161 |

Introduction

Though the rules of language constrain what we can communicate with a given utterance, it is widely assumed that the role of grammar in determining semantic content is quite limited. What ‘that’ picks out on an occasion of use depends on what the speaker intends and what the utterance situation is like, and audiences exploit whatever epistemic cues a speaker and her situation render available to discern her intentions. Philosophers frequently identify (purported) context-sensitivity within philosophically interesting expressions like ‘know’, or ‘good’, and appeal to it to shed light on problems involving knowledge, or goodness, but in doing so, they assume that these expressions get their values in context in virtue of the speaker’s intentions, and the features of the utterance situation.

Against this tradition, I argue that the reason we interpret context-sensitive expressions so effortlessly is that the rules of language are much more subtle and pervasive than has been assumed. In particular, I argue that context-sensitivity is resolved by linguistic conventions that have gone unnoticed because their principal domains are entire discourses, rather than their constituent words and sentences. While it is universally accepted that the conventions of syntax and semantics specify the rules by which individual expressions combine to form sentences, drawing on linguistic research on discourse structure and information structure, I similarly argue for rules–discourse conventions—that specify how individual sentences combine to form a discourse. These conventions govern how speakers organize utterances into larger units that address particular topics and answer questions about them.

To illustrate, speakers tend to understand an utterance of ‘Phil tickled Stanley. Liz poked him,’ spoken out of the blue, in one of two ways: either that Phil tickled Stanley, and as a result, Liz poked him, or that Phil tickled Stanley and similarly, Liz poked him.¹ This is captured in the literature on the so-called Discourse Coherence Theory by positing that in one case, the discourse is organized

¹The example is due to Smyth (1994).

around an event-result relation, and in the other around a relation of parallelism. As a result, in the former case ‘him’ is understood to refer to Phil, and in the latter to Stanley. I argue that the discourse relations that organize utterances into coherent discourses are a part of the grammar of a language, and that they guide the interpretation of context-sensitivity. This in turn means that a host of philosophical arguments that resort to context-dependence rest on a flawed conception of context, and that context-dependence will have to be invoked quite differently than has become customary.

I begin my dissertation by laying out the grounds for a general theory of context-sensitivity resolution, providing empirical evidence that the resolution is governed by discourse conventions. Drawing on empirical linguistic research, I identify mechanisms that I argue constitute these discourse conventions, and which set the values of contextual parameters that govern the resolution of context-sensitive expressions. The general theory in turn provides grounds for re-thinking of the role of context-sensitivity in philosophical arguments. To this end I apply it to the basic class of context-sensitive expressions—pronouns (‘he,’ ‘that’)—and to modal expression (‘might,’ ‘must’). I argue that these expressions are importantly similar in the kind of context-sensitivity they exhibit, and in the ways the discourse conventions determine their resolution. In turn, I show that my general theory has important repercussions for three larger debates in which the context-sensitivity of these expressions has played a central role: contrary to the growing consensus, the context-sensitivity of these expressions neither corroborates the intentionalist accounts of meaning, sketched above, nor the alleged tension between natural language and classical logic, nor the non-truth-conditional accounts of semantic content.

Chapter 2, co-authored with Matthew Stone and Ernie Lepore, develops a general meta-semantics of pronoun resolution, arguing that the resolution of a pronoun is determined by what is most prominent, i.e. at the center of attention in a coherent discourse, where the notion of attention and coherence are made formally precise by appeal to discourse conventions. By characterizing pronoun resolution in linguistic terms, the theory gives a uniform analysis for pronouns, covering its demonstrative, bound, and discourse bound (E-type) readings in a simple and uniform way. The resulting account avoids the appeal to rational or psychological processes of intention recognition that are emphasized in the broadly Kaplanian and Gricean meta-semantics tradition. These theories have mistakenly taken proper linguistic rules to be reflexes of general pragmatic processes.

Chapter 3. deals with an alleged tension between natural language and classical logic. Several authors have independently touted counterexamples to some of the most entrenched classical rules of inference; viz., *Modus Ponens* and *Modus Tollens*. Since each counterexample involves modal language, critics believe they have uncovered a tension between the standard semantics for modality and classical logic. The growing consensus among philosophers and linguists is to abandon the standard semantics for one that invalidates these patterns. The alternative semantics typically abandon the idea that modals express propositional content. On the contrary, I argue that the apparent tension between modal vocabulary and classical logic is a result of a failure to appreciate the import of discourse conventions on the interpretation of modals. The semantic theory I develop specifies these conventions and captures precisely how they affect the interpretation of modals in a context, in a way that explains the appearance of the putative counterexamples, maintains the idea that modals express propositional content and, as I prove, preserves classical logic.

The theory of context-sensitivity I have been advocating raises a general question about the nature of semantic content in this kind of a framework, and its interaction with context. A long-standing tradition in theorizing about meaning, going back to at least Frege (1892, 1918) identifies the meaning of a (declarative) sentence with the proposition it expresses, where a proposition is, or at least determines, the truth-conditions in a given context. But a recent trend in philosophy of language and linguistics has advocated a departure from this tradition. The key data that fueled this trend involve a certain kind of puzzling behavior of modal vocabulary in discourse, which seems to evade propositional treatment.

The non-propositionalism has received different implementations in relativist semantics (Kolodny and MacFarlane, 2010), expressivist semantics (Moss, 2015; Swanson, 2006; Yalcin, 2012b) and certain kinds of dynamic update semantics (Veltman, 1985; Gillies, 2010, 2004; von Fintel and Gillies, 2009). They maintain that the key aspect of interpretation of modal claims is the characteristic dynamic effect they have on the context of utterance. Against this trend, I argue in the final chapter of the dissertation that pessimism about truth-conditional accounts arises from a failure to appreciate the import of linguistic mechanisms that systematically affect the interpretation of modal vocabulary in a context. While I agree with the critics that an important aspect of modal meaning has to do with the dynamic effects modals have on the context, I argue that the revisionary accounts

have mischaracterized the nature and the complexity of these effects.

I propose a theory of semantic content that results from the theory of context-sensitivity I have been developing, according to which utterances, as a matter of meaning, make a two-fold semantic contribution: on the one hand, they express propositional content, and on the other hand, they are associated with dynamic content that actively changes the context of utterance. Both aspects of meaning are crucial and they are interrelated: the dynamic content captures the way in which utterances affect the context, and in turn, such a dynamically changing context determines the propositional content of an utterance. Conversely, the propositional content is required to correctly capture the dynamic effect an utterance has on the context. Since, the dynamic aspect of meaning is crucial for determining the proposition expressed by an utterance in a context, far from being incompatible with propositional meaning, it is precisely what allows us to systematically predict the correct propositional content of a given utterance.

The chapters of this dissertation are conceived as independent papers. A version of Chapter 3. is forthcoming in *Philosophy and Phenomenological Research*.

Chapter 1

Discourse, Context and Coherence: The Grammar of Prominence

1.1 Introduction

What we say when we speak depends not solely on the words and expressions we use, but also on the context in which these words and expressions are uttered. If I say “I am hungry” and you also say “I am hungry”, we have said two different things. If I say “She is happy” pointing at Ann, and you say “She is happy” pointing at Sue, we have said two different things. This phenomenon has to do (at least in part) with expressions such as ‘I’ and ‘she’.¹ The unifying feature of these expressions is that their interpretation varies with a context of utterance: they are context-sensitive. Context-sensitivity permeates all aspects of our linguistic communication. Yet, we resolve it effortlessly, “on the fly”.²

How is this achieved? What are the means by which the interpretation of a particular context-sensitive expression is determined on an occasion of use? And what sorts of interpretive resources do interlocutors appeal to when settling on an interpretation of a context-sensitive expression? It is quite common in the philosophical and linguistic literature to appeal to context-sensitivity of a particular item, while leaving the question of how context-sensitivity is resolved more or less open. For certain purposes this strategy might be warranted. However, ultimately it is important to provide an adequate answer to this question. An adequate theory of linguistic communication

¹I say “in part” because the fact that the information conveyed varies with context of utterance is partly tied to considerations concerning other context-sensitive elements, such as tense, etc.

²Which expressions are context-sensitive is a matter of great controversy; some believe that the class is rather small, containing only the so-called indexical expressions (like ‘I’) (Lepore and Cappelen, 2005), others believe that the class is so large that it comprises virtually all expressions (Travis, 1989). Most, however, believe that the class contains more than just a few, but fewer than (almost) all expressions. But even so, they still disagree about which particular expressions are context-sensitive (and in what way). So, for example, we find debates over context-sensitivity of modals (Egan, Hawthorne, and Weatherston, 2005; von Stechow and Gillies, 2008; Kolodny and MacFarlane, 2010; Yalcin, 2007), predicates of personal taste (Egan, 2010; Cappelen and Hawthorne, 2009), knowledge ascriptions (Hawthorne, 2004a; Stanley, 2005; Schaffer and Szabó, 2013; DeRose, 2009, 1995; Cohen, 1998), just to mention a few.

has to provide an account of how context-sensitivity is integrated within the semantics-pragmatics interface. Moreover, philosophers often posit context-sensitivity of particular philosophically interesting expressions (e.g. ‘know’, ‘good’, ‘truth’), and then use this context-sensitivity to motivate various philosophical conclusions about the underlying phenomena (e.g. knowledge, goodness, truth). There has virtually been no area in philosophy where the contextualist move hasn’t been exploited in one way or another. It is important, in assessing these arguments, to determine whether context-sensitivity is capable of the kind of patterns of resolution they predict and rely on.

So, how is context-sensitivity resolved?³ Linguistic communication is obviously constrained by grammar, but it also seems to depend on various features of an utterance situation, for example, which gestures accompany my utterance, which contribution the speaker intended to get across, and whether it was intended literally or figuratively. Accordingly, most theorists hold that meaning largely relies on communicative situations and intentions, and the interpretation involves general reasoning about these utterance features.⁴ This is manifest in the standard assumptions about the resolution of context sensitivity. To see this, it is useful to recall Kaplan’s (1989a) distinction between ‘pure indexicals’ and ‘true demonstratives’. Pure indexicals, in Kaplan’s sense, are context-sensitive expressions whose linguistic meaning fully governs their resolution. The standard example is the first person pronoun ‘I’. The linguistic meaning of ‘I’, its *character*, in Kaplan’s terminology, is roughly, *the speaker of the context*. This character is sufficient to fully determine, given a context, what the referent of ‘I’ is; one needn’t look into the speaker’s mind in order to determine this. On the other hand, we have true demonstratives, expressions like ‘he’ or ‘that’, whose character does not completely determine their meaning given a context. If I say ‘That is lovely’, it doesn’t suffice that you know that the character of ‘that’ is something like *the salient object in the context*. Some

³Sometimes, the theory of how context determines interpretation is called *metasemantics* (Kaplan, 1989b,a; King, 2014a; Glanzberg, 2007a), though one should bear in mind that the term ‘metasemantics’ has been used in different ways by different authors. For instance, it is sometimes used to denote *pre-semantic* processes that need to take place before the mechanisms of semantic interpretation can even kick in (e.g. disambiguation), and sometimes to denote the theory of how semantic facts get fixed (Kripke, 1980). The present paper is *not* concerned with metasemantics, in the latter two senses. Nor is my project concerned with metasemantics in the sense of Yalcin (2015), who understands it, roughly, as a theory of the grounds of the semantic properties postulated by the first-order semantic theorizing.

⁴The idea that interpretation relies on mind-reading goes back at least to Locke (1689), and, with important modifications, continues to be the dominant model through Grice (1975) and Kaplan (1989a). However, whereas Locke thought communication is entirely a matter of encoding and decoding content through linguistic forms, currently it is almost universally accepted, as Grice and Kaplan would have it, that meaning is at least partly determined through non-linguistic features of utterances. It is the Gricean/Kaplanian model that has dominated the discussions on context-sensitivity resolution.

supplementation is needed to tell us how salience is determined. Crucially, this supplementation has traditionally been thought to go beyond the scope of grammar. What ‘that’ picks out depends on what the speaker intends and what the utterance situation is like; correlatively, the audience exploits whatever epistemic cues a speaker and her situation render available to discern her intentions.

Most context-sensitive expressions have traditionally been thought to pattern as true demonstratives. So, the standard answer to the question of how context-sensitivity is resolved, going back to at least Kaplan (1989a) and Grice (1975), is that the resolution of context-sensitive items depends on non-linguistic features of utterances. It is ultimately the speaker’s (referential and/or communicative) intentions that determine the resolution of context-sensitive items, perhaps aided by certain contextual features that serve as epistemic cues that help make the relevant intention manifest.⁵ This model predicts significant flexibility in the effects of context on the interpretation of context-sensitive items.

By contrast with the tradition, I shall suggest that all context-sensitive expressions are best thought of as pure indexicals in the sense that their character determines their content given a context. To illustrate, take the paradigm example of a true demonstrative, a demonstrative pronoun, ‘she’. The character of ‘she’ is, roughly, the most prominent individual in the context, satisfying the number, gender and person requirements associated with the pronoun. What I shall argue is that given this character, the semantic content of ‘she’ is determined automatically, as a function of context, much like the semantic value of ‘I’ is. One does not need to appeal to extra-linguistic features of context (e.g., world-knowledge and speaker’s mental states), as determinants of semantic

⁵There are many proponents of the view that in some form or other the key meaning-determining role is played by speaker intentions. For instance, see e.g. Cohen (1998); Dowell (2011); Neale (2004, 1990); King (2014a,b); King and Stanley (2005); Reimer (1992). There are important differences among the proponents of different views in this tradition, but they all in one way or another emphasize the role of speaker intentions. One should note that most authors reject the naïve form of intentionalism, according to which the semantic content of a context-sensitive expression simply is what the speaker intended it to be. To illustrate the problem, suppose that the speaker points at John, and utters ‘He is happy’, while in fact intending to refer to Tim, her friend, who is neither present nor in any other way salient to the interlocutors. It seems implausible that her utterance expresses that Tim is happy. In response to this problem, some authors have noted the need for additional constraints on the communicative/referential intentions, requiring e.g. that they be appropriately manifest, as well as putting constraints on what can be reasonably intended. See in particular King (2014a,b) who offers a sophisticated intentionalist account that puts constraints both on intentions and what can be reasonably intended, and on the features of speakers and audience. Others have argued for what is sometimes labeled ‘objective’ meta-semantics, that does not give a privileged role to speaker intentions, but rather maintains that (at least in certain cases), more objective features of context and general facts about the world fix the resolution of context-sensitive expressions (Wettstein, 1984; Glanzberg, 2007b; Lewis, ms.). The view I shall defend differs from both of these types of view, since it maintains that it is meaning conventions, rather than referential/communicative intentions, or general non-linguistic facts about the world or context, that play the content-determining role for context-sensitive expressions.

content. This is because, as I shall argue, the prominence of potential referents in a context is itself maintained through a set of linguistic rules. The content of a pronoun is automatically determined as whatever is the most prominent referent in the context satisfying its character, at the point in at which the pronoun occurs. So, it is the linguistic meaning of the term, together with linguistic mechanisms governing prominence, that determine the referent; no extra-linguistic supplementation is required.

I shall defend this kind of a meta-semantic account not only as a meta-semantic account of pronoun resolution, but as a general meta-semantics of context-sensitivity. In particular, I shall argue that contextual parameters governing the resolution of context-sensitive expressions are set to particular values at particular points in discourse through linguistic mechanisms that govern how particular values are made prominent, or demoted, as the discourse evolves. Though my account treats all context-sensitive expressions as pure indexicals, it preserves an important insight behind the original Kaplanian distinction between pure indexicals and true demonstratives. Namely, most context-sensitive expressions have been thought of as true demonstratives, because, unlike the character of ‘I’, their character is sensitive to the *prominence* of a particular value of a contextual parameter determining its meaning. I will call such expressions *prominence-sensitive* expressions.⁶ However, by contrast with the tradition which maintains that prominence is largely a matter of non-linguistic features of utterance, I shall argue that the dynamics of contextual prominence is governed by linguistic mechanisms.

⁶This is not to say that the value of a contextual parameter governing the meaning of the first person pronoun is not dynamically changing as the discourse evolves. After all, interlocutors normally take turns speaking. However, the character of a prominence-sensitive expression is directly sensitive to a prominence ranking, having roughly the following form: ‘the most prominent value satisfying the property X’, where *the property X* is whichever property captures the linguistic constraints on the potential value of the term in question (e.g. phi-features of a pronoun). The character of ‘I’ does not reference prominence in this way. Another way to get at this distinction is to note that prominence sensitive expressions typically allow for discourse and operator bound anaphoric readings—that is, their interpretation can co-vary with an expression introduced antecedently in the discourse. Expressions like the English first person pronoun ‘I’, do not allow for such readings. Their interpretation is sensitive to a parameter that plays a particular contextual role that is not sensitive to prominence (e.g., the character of ‘I’ is just ‘the speaker of the context’, not ‘the prominent individual speaking’). Note that, interestingly, ‘here’ and ‘now’, thought they are often classified as pure indexicals, together with ‘I’, do allow for readings on which they select a prominent value from the context. This explains why they can be bound, or discourse bound (unlike ‘I’). While traditionally this has been thought of as a challenge to the claim that these expressions are pure indexicals, we shall see that all such examples demonstrate is that these expressions are prominence-sensitive. As we shall see, this is perfectly compatible with them being pure indexicals. For a detailed development of an account of prominence-sensitivity of ‘now’ along the lines of the theory I shall propose here, see Altshuler and Stojnić (2016).

1.2 Coherence

I begin with a perhaps obvious, yet often neglected insight that comes from the literature on *discourse coherence*: a discourse is more than just a random sequence of unrelated sentences (Hobbs, 1979; Asher and Lascarides, 2003; Kehler, 2002). The stark difference between (1) and (2), from Hobbs (1979), nicely illustrates the point:

- (1) John took a train from Paris to Istanbul. He has family there.
- (2) John took a train from Paris to Istanbul. He likes spinach.

While (1) is perfectly fine, (2) sounds bad. Why is this? Crucially, (1) is not a mere sequence of random, unrelated facts about John. The two bits of discourse are connected in a meaningful way—naturally, (1) is understood as conveying that John took a train from Paris to Istanbul *because* he has family there. This explanatory connection is a part of understanding (1); a failure to grasp this is a failure to fully understand the speaker's contribution in (1).

Meanwhile, in (2), a failure to establish an appropriate connection between the two sentences results in a failure to fully understand the discourse; the interlocutors are left searching for an explanatory connection. Unable to affirm it, they cannot fully grasp the speaker's contribution. The same interpretive effort is invoked in interpreting both discourses, which suggests an implicit organization of the discourse that is recoverable by the interlocutors according to shared standards.

Coherence theorists materialize these observations by positing an implicit organization of the discourse that establishes coherence relations—such as the explanatory connection in (1)—among successive utterances, signaling how the speaker is organizing her contributions (Hobbs, 1979; Asher and Lascarides, 2003; Kehler, 2002). These relations signal what the speaker is doing with a given utterance—what kind of a discourse move she is making; she might be providing an explanation, continuing a narrative, elaborating on a previously mentioned event, etc.⁷

Importantly, the tasks of establishing coherence in a discourse, and of pronoun resolution are

⁷What kinds of inferential connections can we expect to find in a discourse? Kehler (2002) suggests that coherence relations cluster in three broad categories. A discourse can be organized around a sequence of unfolding events, or an extended elaboration of a particular event, illustrating the class of what Kehler calls *contiguity* relations, as in (i):

- (i) Max spilt a bucket of water. He spilt it all over the rug. (Kehler, 2002)

mutually correlated, not independent tasks. To illustrate, consider the following example from Smyth (1994):

- (3) Phil tickled Stanley, and Liz poked him.

There is at least two ways one can understand (3). According to one interpretation, the second clause is describing the result of the event described in the first clause—Phil tickled Stanley, *and because of that* Liz poked him. Alternatively, one can understand (3) as comparing and contrasting two events—Phil tickled Stanley, and *similarly*, Liz poked him. Crucially, when (3) is understood as organized by the Result relation, the referent of ‘him’ in (3) is Phil, and when the organizing relation is Parallel, the referent is Stanley.⁸

Elsewhere I have argued that these two tasks are more than merely correlated (Stojnić, Stone, and Lepore, 2013, 2014). In particular, I have argued that coherence relations, as a matter of a linguistic convention, make certain referents prominent for anaphora resolution, demoting others. Of course, it might seem that given a particular choice of a coherence relation, only a certain resolution makes sense. For instance, if the pronoun occurs in a narrative flow about a particular person, it is only natural that it refers to that person. However, as I argued, the effects of coherence relations on pronoun resolution are present even when that leads to an overall implausible, and even incoherent reading, and moreover even when a perfectly plausible alternative interpretation is available. The following example from Kehler (2002) illustrates the point:

A discourse can also be organized around the cause-effect (or event-result) relations:

- (ii) Max spilt a bucket of water. He tripped over his shoelaces. (Kehler, 2002)

Finally, a discourse could be organized around the so called *resemblance* relations, comparing and contrasting two events along a certain dimension of (dis)similarity, as in:

- (iii) Max spilt a bucket of water. John dropped a jar of cookies.

Each of these three broad classes of coherence relations contains further more fine-grained instances, and the research suggests that interlocutors draw on a wide variety of specific relations. (See, e.g., Mann and Thompson (1988); Knott (1996); Asher and Lascarides (2003).) Yet the broad classification is illustrative of the kind of dependencies we might expect to find.

⁸That pronoun resolution and the choice of a coherence relation are thusly correlated has been confirmed by a number of empirical studies. See e.g. Keiser (2009); Wolf, Gibson, and Desmet (2004); Kehler et al. (2008).

- (4) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

As Kehler reports, English speakers find (4) infelicitous, interpreting it as if the speaker has mistakenly used ‘her’ to refer to Reagan. While there is a perfectly accessible, suitable antecedent for the pronoun ‘her’ that is a well know object of Bush’s admiration, namely, Thatcher, the interlocutors do not opt for this interpretation, but rather recover an infelicitous one, that requires attributing an error to the speaker. Why is this? Note that the ‘admires’ in the first sentence is followed by a stronger alternative, ‘worships,’ thus signaling that the discourse is organized around a comparison of Thatcher’s and Bush’s respective attitudes toward Reagan. This signals Parallel relation between the two conjuncts, that in turn makes Reagan prominent for subsequent reference resolution. This effect of Parallel relation cannot be simply overridden by common-sense, world-knowledge and plausibility considerations: the infelicitous interpretation is recovered even though there is a perfectly plausible nearby interpretation. This suggests that the effect of Parallel on pronoun resolution is conventionalized.

Moreover, though coherence relations have a certain effect on reference resolution *in English*, these effects differ across languages. Languages with overt markers for shifting attentional prominence, and signaling preference for reference resolution, tend to be much more constrained in what sorts of shifts they allow to happen implicitly, as an effect of discourse coherence. Since common-sense reasoning, and real world and plausibility considerations are presumably the same the world over, such cross-linguistic variations point to conventionality as the mechanism underlying these effects.⁹

I shall argue that discourse structuring mechanisms affect, and indeed, determine, the resolution of context-sensitivity quite *generally*, and do so as a matter of linguistic conventions, not as a matter of general common-sense reasoning in search of the overall most plausible interpretation. In particular, we shall see that the kind of effects of discourse structuring mechanisms we see on pronouns carry over to other types of context-sensitive expressions, and that, as with pronouns, such effects exhibit the hallmark of conventionality.

⁹See Stojnić, Stone, and Lepore (2014) for discussion and further examples.

1.3 Coherence and Context-Sensitivity

As the examples from the previous section illustrate, mechanisms of discourse coherence make certain referents prominent as antecedents for subsequent pronouns. One way to think about the affects of discourse coherence on pronoun resolution, the one I have argued for elsewhere (Stojnić, Stone, and Lepore, 2013, 2014), builds on the idea that we can rank candidate referents in a discourse according to their relative prominence, those higher in the ranking being preferred resolutions to those lower in the ranking.¹⁰ We can think of coherence relations as affecting prominence by updating the prominence ranking of candidate referents—making certain referents prominent and demoting others. A pronoun is then simply resolved to the highest ranked candidate referent that satisfies its character; for instance, ‘she’ denotes the top-ranked female, third person, singular referent at a given point in a discourse. Thus, insofar as we think of the character of a pronoun as determining its semantic value as a function of context, we can think of the prominence ranking as a parameter of the context the character is operating on.

In a Kaplanian framework, a context is a tuple comprising the parameters that affect the resolution of context-sensitive expressions (a tuple comprising the speaker, world, time and location of the utterance). But we can generalize this idea, by envisioning context as a *conversational record*, in the sense of Lewis, an abstract scoreboard that keeps track of the relevant information about the conversation. In particular, the scoreboard keeps track of the standard parameters fixing the meaning of the of indexicals, such as who is speaking, at what time, in which world, who is the addressee, and so forth. More importantly, it is also keeping track of the information about how the relevant contextual parameters change as the discourse evolves. It tracks which moves have been made in the conversation, what has been said, which propositions have been mentioned, which individuals have been made prominent. In this way, we can think of the prominence ranking of candidate referents as of one aspect of a conversational scoreboard. Thus, unlike Kaplanian context, which is static, the conversational scoreboard keeps track of the dynamic change in relevant contextual parameters as the discourse evolves. Given this notion of a context, a character of a demonstrative pronoun determines its semantic value as a function of context—e.g., ‘she’ selects the top-ranked candidate

¹⁰The idea that we can rank candidate referents for anaphora resolution is exploiting the resources from Centering Theory (Grosz, Joshi, and Weinstein, 1995; Sidner, 1983; Bittner, 2014).

referent, which is singular, third person, female.¹¹ Coherence relations can then be seen as affecting prominence by promoting certain referents, while demoting others. More precisely, a part of the linguistic contribution of a coherence relation is a prominence re-ranking update.

Moreover, as I shall argue presently, coherence relations can be seen as changing parameters of the scoreboard that determine the resolution of context-sensitive items more generally: the contextual prominence of values of parameters governing the resolution of prominence-sensitive expressions generally is affected by the same mechanisms which affect prominence of candidate referents for pronoun resolution. Consider the case of an obviously prominence sensitive class of expressions: that of quantifiers. An utterance of ‘Everyone had fun today’ does not typically mean that *everyone in the universe* had fun on the day of utterance. Typically, it conveys something more restricted, perhaps that everyone that attended some particular event had fun. The restriction on the domain of quantification is normally thought to be contextually determined.¹² Moreover, just as pronouns are prominence-sensitive insofar as they are looking for an antecedent that can be provided either from the non-linguistic context, or from a prior discourse, quantifiers too are looking for a prominent restrictor either from the non-linguistic context or from a prior discourse.

For instance, just as one can use (5) to refer to a prominent female in a non-linguistic context, so too one can use (6) to refer to a prominent set of individuals.¹³

(5) (Referring to a certain significant female) She left me. (Partee, 1984)

(6) (Looking around the classroom about to start my lecture) Everybody is here, so we can start.

But a restrictor can be provided by the prior discourse, as well:

(7) Half of the students came to the party. Some students brought cookies.

¹¹It is worth pointing out that this type of an account allows for an implementation which captures both referential and bound (including discourse bound, or E-type) readings of pronouns. In particular, the semantics will treat all uses of pronouns as co-varying with an antecedent expression, and thus, as bound variables. For formal details, see Stojnić, Stone, and Lepore (2014). For the present purposes, the informal sketch of the view will suffice.

¹²For more on quantifier domain restriction, see e.g. von Stechow (1994); Stanley and Szabo (2000).

¹³Whether the domain restriction comes in terms of sets, or functions from objects to sets, or properties or perhaps situations is a matter of debate (Stanley and Szabo, 2000; von Stechow, 1994). Though specific details are important, we need not worry about them presently. For a more detailed account, see Stojnić (2016).

Here the second quantifier is understood to quantify over a subdomain of the first one, thus receiving a more restricted reading—some *of the students who came to the party* brought cookies. What I shall now argue is that the prominence of potential restrictors on the domain of quantification is governed by mechanisms of discourse coherence, just as the resolution of pronominal anaphora is. We can readily appeal to these mechanisms to predict the intuitively correct interpretation of (7). The first sentence describes a situation about the group of students at a party. The second sentence is normally understood as elaborating on the first one; the set of students introduced in the first sentence is thus made prominent, and ‘Some students’ is restricted to quantify over that set of students. So, we understand (7) to mean that half of the students came to the party, and some of those students who came to the party brought cookies.

Perhaps a more interesting example is the following one:¹⁴

(8) If there is fire, the alarm always rings.

(9) But the alarm doesn’t always ring.

In (8), the consequent of the conditional is understood as elaborating on the information described by the antecedent, and the quantifier, ‘always’, is understood as restricted by those situations described in the antecedent, namely the ones in which there is fire. In (9), however, ‘always’ in the second sentence is *not* thusly restricted—we cannot, for instance, conclude from (8) and (9) that there is no fire. Assuming that mechanisms of discourse coherence affect the prominence of potential restrictors, we can explain this difference in domain restriction between (8) and (9). The second sentence stands in Contrast relation to the first one. The two sentences are thus understood as contributing contrasting bits of information about some body of information.¹⁵ Contrast makes the situations described by that body of information prominent. As a result, ‘always’ in the second sentence is understood as quantifying over the times of those situations, not just the ones in which

¹⁴The example is from (Yalcin, 2012b), though he uses it for different purposes, and does not exploit it in the context of domain restriction.

¹⁵More precisely, where *p* and *q* stand in a Contrast relation, they are required to provide contrasting information about some body of information, regarding some common topic. The relevant topic is typically signaled by the cues in the information structure, the dimension of meaning that signals how information conveyed by an utterance relates to the previous contributions (by dividing it into the topic and comment, or focus and background), how it contrasts with alternative possibilities in the discourse, and which role it plays as a move in the dynamics of conversation. In English, we exploit prosodic accents as signals of a particular information structure (Rooth, 1992; Roberts, 1996a).

there is fire.¹⁶ So we understand (8) as conveying that it's not the case that the alarm rings in all the situations compatible with our overall body of information. Compare this case with the following:

(10) If there is fire, the alarm always rings.

(11) The firefighters usually arrive within five minutes.

Superficially, (10) and (11) seem to have exactly the same shape as (8) and (9). However, they are not naturally understood as standing in the Contrast relation. Rather, (11) is understood as Elaborating on the situation described by the consequent of (10) (which in turn is itself understood as elaborating on the antecedent). As a result of Elaboration, the situation described by the consequent—the one in which there is fire, and the alarm rings—is made prominent, and the quantifier in (11) is restricted by it. So we get the reading that *at all the times when there is fire*, the alarm always rings, and the firefighters usually arrive within five minutes.

Unsurprisingly, a similar kind of behavior is exhibited by modal expressions, which are traditionally understood as quantifiers over a contextually determined set of possibilities.¹⁷ Note first that just like quantifiers (and pronouns), modals too can be restricted by a possibility either made prominent in the non-linguistic context, or mentioned earlier in a discourse.¹⁸ For instance, one can use (12) to describe a hypothetical scenario prominent in the non-linguistic context:

(12) (Looking at a high-end stereo in an electronics store) My neighbors would kill me. (Stone, 1997)

In turn, in (13), the restriction is provided by the prior discourse:

(13) A wolf might walk in. It would eat you. Roberts (1989)

So, we see that, like pronouns and quantifiers, modals are likewise prominence-sensitive, searching for a prominent possibility restricting their domain of quantification. And again, prominence

¹⁶For a formalization of this effect of Contrast, see Stojnić (forthcoming.).

¹⁷I have defended this account of modality in detail elsewhere. See Stojnić (forthcoming.) and Chapter 4.

¹⁸See (Stone, 1997, 1999), for a detailed argument for the anaphoricity of modal expressions.

is sensitive to mechanisms of discourse coherence. For instance, in (13), we understand the second sentence to elaborate on the first one. Elaboration makes the possibility described by the first sentence prominent. The second sentence is then understood as quantifying over the (epistemically accessible) worlds in which this possibility obtains, and so we get the intuitively correct reading: a wolf would eat you, given that one walks in. Again, as before, the affect of coherence on prominence is crucial, as is illustrated by the following example:

- (14) If a wolf walks in, it would eat you. But one probably won't walk in. So, you are probably safe.

The consequent of the first sentence in (14) elaborates on the hypothetical scenario described by the antecedent. Elaboration makes the hypothetical scenario introduced by the antecedent prominent, and as a result, 'would' is understood as restricted by it—a wolf would eat you given that one walks in. Crucially, however, the second sentence does not further elaborate upon a scenario described by the two modals in the first sentence. The two sentences stand in Contrast relation, signaled by the discourse marker 'but', and are understood as contrasting two hypothetical scenarios—one in which a wolf walks in, and one in which one does not. Contrast requires that the first and the second sentence provide contrasting information about some body of information. Assuming that the conditional is uttered discourse initially, the first sentence is interpreted as contributing information about the set of epistemically accessible worlds determined by the context discourse initially, describing what might be the case if a wolf walked in, *given this overall body of knowledge*. The second sentence then has to provide a contrasting bit of information, regarding the possibility of a wolf's entrance, about this body of information. Contrast makes this body of information prominent, and as a result, 'would' is understood as quantifying over it. Thus, we get the intuitively correct interpretation that *given this overall body of knowledge*, a wolf probably won't come in.

Similarly, consider (15):

- (15) A wolf might walk in. But, then again one will probably not walk in. It would eat Harvey.

In (15), we cannot understand 'would' in the final sentence as restricted by the proposition describing the epistemic possibility of a wolf walking in. The first two sentences stand in the Contrast

relation, and the two modals are understood as contributing contrasting bits of information about some body of information (e.g. the epistemically accessible worlds determined by the context discourse initially). The third sentence, however, is naturally understood as further elaborating on the second. Elaboration makes the possibility described by the second modal prominent, and that is why the modal in the third sentence, ‘would’, is understood as quantifying over it. Note that this interpretation is persistent, despite the intuitive infelicity of (15).

Similar behavior is exhibited by other prominence-sensitive expressions. For instance, we see this with the so called incomplete definite descriptions. To see what is at stake, consider the following example:

(16) The table is covered with books.

Definite descriptions typically are thought to require some form of uniqueness (so, that ‘the *F*’ denotes the unique *x* satisfying *F*).¹⁹ But examples like (16) seem to go against this requirement. Even though there are many tables in the world, we can imagine (16) uttered in a room with exactly one table, to successfully convey that the table in the room is covered with books.²⁰ So, ‘the table’ in (16) is an incomplete description—the context supplies an extra bit of content that is left implicit. The uniqueness is merely required within a contextually restricted domain. Thus definite descriptions, like quantifiers, and modals, require a contextually prominent restriction—they are prominence sensitive. As before, the restriction can be provided by the non-linguistic context, as in (16), or by the prior discourse, as in (17):

(17) John has a cat and a dog. He walks the dog twice daily and lets the cat out at night. Roberts (2011)

Once again, I suggest, prominence is affected by discourse structuring mechanisms. In (17), we understand the second sentence as elaborating on the situation described by the first. Similarly as

¹⁹This observation goes back to at least Frege (1892) and Russell (1905).

²⁰As Roberts (2011) points out, even a weak requirement that the referent be unique merely within the common ground—i.e. within the body of information the interlocutors mutually accept for the purposes of the conversation—won’t help us get around the problem. A typical common ground will normally entail the existence of many tables.

before, Elaboration makes this situation prominent. The situation in turn restricts the domain relevant for the interpretation of the definite descriptions in question: the two definite descriptions are understood as denoting the unique dog and cat in the (minimal) situation described by the previous utterance, namely, the one in which John has a cat and a dog.²¹

A similar, though slightly more complicated, example is the following one from Lewis (1979):

- (18) (While Bruce, the cat, is in the room, making itself salient by dashing madly about): The cat is in the carton. The cat will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he'll stay, because Miriam would be sad if the cat went away.

Lewis himself uses this example to illustrate how definite descriptions are sensitive to prominence, arguing that 'the *F*' denotes the unique *F* that is most salient in the domain of discourse. Assuming that prominence is governed by discourse coherence, here's how we can explain (18): the first sentence summarizes a perceptually present situation, and as a result, the description is understood as denoting the unique cat within this situation. The first conjunct of the second sentence provides a further elaboration about the cat in the first sentence, making the situation described by the first sentence prominent: the cat in the perceptually present situation will never meet the other cat; in turn, the situation described by the second conjunct provides an explanation of the first one—it is because the other cat lives in New Zealand that the first cat will never meet it—and, the last two sentences provide further elaboration of this explanation. It is due to this chain of Elaboration relations that the situation described by the second conjunct of the second, and the third sentence is made prominent, and consequently, we understand 'the cat' in the final sentence to be the unique cat in this situation. Thus, we get the interpretation that the New Zealand cat will stay, because Miriam would be sad if it went away.

Similar kinds of considerations generalize to other kinds of context-sensitive expressions as well. Consider, for instance, a gradable adjective, like 'tall'. Gradable adjectives are typically

²¹There is a long-standing tradition of treating domain restriction by exploiting possible situations (Barwise and Etchemendy, 1987; Kratzer, 1998, 2014; Recanati, 1996; Elbourne, 2008). Though this approach is elegant, it's not essential for the present account. We could equally well exploit reference to sets, or properties as mechanisms of domain restriction (von Fintel, 1994; Stanley and Szabó, 2000).

treated as expressing relations between individuals and degrees, where the individual satisfies the property just in case it possesses the property to at least a contextually specified degree.²² Thus, gradable adjectives are looking for a contextually prominent degree or threshold. What this degree in the context exactly is, is typically left to pragmatic reasoning to determine. However, examples (19)—(20) again suggest that the mechanisms of discourse coherence affect the possible values of this contextual parameter, making certain resolutions prominent.

(19) John is tall! His height is 2m.

(20) We should put John on our basketball team. He's tall.

In (19), the second sentence is related to the first one by Specification, thus providing more detail about how one should understand the first sentence.²³ In particular, the second sentence makes it clear that being two meters tall suffices to count as tall in the given context. In (20), in turn, Explanation relation ties the two bits of discourse together. This again specifies the relevant degree of height—in particular, John has to be at least as tall as an average basketball player.

And the same point extends to relational expressions, like 'local', that require a contextually specified 'implicit argument'. Note, first, that an utterance of (21) can convey that John went to a bar local to him, or local to some other, contextually specified person or place. Relational expressions, like 'local', are searching for an implicit argument prominent in the non-linguistic context or in the prior discourse.

(21) John went to a local bar.

Once again, in the case of relational expressions, the prominence of potential arguments is sensitive to discourse structuring mechanisms. Consider the following example:

(22) The Maryland Food Bank launched its "Save a Seat at Your Table" campaign in November, a large-scale effort encouraging Marylanders to pledge \$1 per day to help feed the 757,000

²²See e.g. Kennedy (2007); Hamann (1991); Klein (1991); Barker (2002). We need not concern ourselves with specific formal details of the semantics of gradable adjectives for our present purposes.

²³For a similar point involving Specification, see Stone and Stojnić (2015).

people who are food insecure in the state. Now, the nonprofit organization is recruiting local restaurants to join the movement.²⁴

Here, ‘local’ occurs within a Narrative about Maryland and its Food Bank, and is thus understood as local-to-Maryland. Crucially, ‘local’ cannot just freely be used to denote a location local to any place or person made salient in the context. Consider the following:

- (23) John was in Paris in May, but on June 1st he arrived in Belgrade. He checked into a hotel, and proceeded to explore the city. He particularly enjoyed a local bar near the Eiffel Tower.

Interestingly, in (23), we naturally interpret the final sentence as conveying information that John enjoyed a local bar in Belgrade, near the Eiffel Tower. Even though the audience may well be aware that there is no Eiffel Tower in Belgrade, and might know perfectly well that the Eiffel Tower is in Paris, and moreover even though Paris has been explicitly mentioned earlier in the discourse, we do not understand the final sentence as conveying that John particularly enjoyed a local bar in Paris. Assuming that prominence of a particular resolution of ‘local’ is dictated by discourse coherence, we can predict this straightforwardly: the final sentence in (23) is introduced in a chain of Elaboration about the Belgrade trip. The Elaboration makes Belgrade prominent, and as a result, ‘local’ is understood as being local to Belgrade.

What all these examples illustrate is that discourse structuring mechanisms dictate the prominence of potential values of contextual parameters governing the resolution of context-sensitive, prominence-sensitive expressions quite generally. We see similar effects across a range of examples—pronouns, quantifiers, modal expressions, definite descriptions, gradable adjectives and relational expressions. Though of course the set of examples considered here is limited, and thus the discussion provisional, what seems to transpire is that these kinds of effects on prominence are not isolated to a small sample of expressions, but are rather general and pervasive. In all these cases discourse coherence governs contextual prominence in analogous ways.

²⁴<http://www.baltimoresun.com/entertainment/dining/baltimore-diner-blog/bal-local-restaurants-join-in-maryland-food-bank-s-save-a-seat-campaign-20160311-story.html>

1.4 Coherence and Grammar

How does the discussion of the previous section bare on intentionalism and, more generally, on the traditional pragmatic accounts of context-sensitivity resolution? If the impact of coherence on prominence of a particular value for a contextual parameter were merely a byproduct of general pragmatic reasoning—that is, if a particular coherence relation simply served as a cue manifesting (defeasibly) a particular speaker intention, or if it were merely one among many, potentially extra-linguistic cues that helped fix the supplementation required for the interpretation of context-sensitive expression given its incomplete character—then the intentionalist, and pragmatic meta-semantics would be compatible with the considerations raised in § 1.3. Most context-sensitive expressions would then, indeed, behave like true demonstratives. If, however, the impact of coherence on the prominence of potential values of contextual parameters were a matter of an underlying linguistic convention—a rule of language—then context-sensitive expressions, quite generally, would behave like pure indexicals. Given their character, no extra-linguistic information would be required to determine their semantic content in a context. More precisely, though the linguistic meaning of a particular context-sensitive expression might require contextual supplementation, if this supplementation is itself determined by linguistic constraints, then it is the rules of language—those governing the use of the expression together with those governing the prominence of potential resolutions—that determine the semantic content of the expression in a context. As foreshadowed in the previous sections, I shall argue that evidence favors the latter view.

We have already seen some evidence that the effects of coherence on pronoun resolution are a matter of underlying linguistic conventions. Recall (4), repeated in (24):

(24) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

We have seen that, even if we suppose that the speaker uttered (24) with an intention to refer to Thatcher with the pronoun ‘her’, (24) is still infelicitous. The choice of a coherence relation—in this case Parallel—requires that the antecedent of the pronoun be introduced in the object position. This is what underwrites the judgement of infelicity. And this is so, even though in this particular case, all the plausibility, relevance and charity considerations would favor a different interpretation: there is an available antecedent that matches the grammatical gender of the pronoun, namely,

Thatcher, which is, moreover, a well-known object of Bush's admiration. The interpretation actually recovered, according to which the speaker has mistakenly uttered 'her' to refer to Reagan, is uncharitable. That it is nevertheless the naturally recovered interpretation, despite all these pragmatic considerations, suggests that Parallel dictates a particular interpretation of a pronoun as a matter of an underlying linguistic rule, which is hard to override by such pragmatic considerations.²⁵

Evidence suggests that the effects of coherence on prominence more generally exhibit hallmarks of conventionality: these effects seem to be a part of our linguistic repertoire, inert to considerations of common-sense reasoning and plausibility, and variant across different linguistic communities. We have already seen examples which suggest that this is the case with the effects of coherence on domain restriction of modals. For instance, in (15), even though the possibility of a wolf walking in has been explicitly mentioned, and left open in the discourse, the third sentence cannot be understood as restricted by this possibility. It is understood as elaborating on the second sentence, and is thus restricted by the possibility described by the second sentence. This is so despite the fact that this leads to an infelicitous interpretation. So, just as in the case of pronouns, the mechanisms of discourse coherence sometimes force an infelicitous reading of modal utterances, even when pragmatic account would predict plausible alternative interpretations.²⁶ Another, perhaps even clearer, example of this is the following one from Yalcin (2007):

(25) If it is not raining and it might be raining, then I'm uninformed about the current weather.

As Yalcin notes, (25) is infelicitous. We are left with the impression that we are asked to consider an inconsistent hypothetical scenario. This is puzzling from the standpoint of a contextualist who maintains that modals are simply quantifiers quantifying over contextually provided sets of worlds, where the relevant worlds are determined according to the standard pragmatic meta-semantics. The problem is that considerations of charity and plausibility do not predict that (25) should be infelicitous. There are available bodies of information that could serve as restriction on the domain of quantification for a modal 'might' that would deliver a perfectly consistent interpretation. For

²⁵Note that if the pronoun in (24) is accented, one understands it as referring to Thatcher. However, in that case the discourse is no longer understood as organized by Parallel. Rather, we understand it as explaining how Bush follows Thatcher's opinions. I shall return to this point in § 1.5.

²⁶For detailed discussion of this point and further examples, see Chapter 4.

instance, the modal could be quantifying over the speaker's body of information, thus yielding a perfectly consistent interpretation, namely, that if it is raining, and for all the speaker knows it is not, then the speaker is uninformed about the weather. But the inconsistent reading is predicted if we assume that mechanisms of discourse coherence make certain possibilities prominent as restrictors for modals as a matter of grammar. We naturally understand the second conjunct in the antecedent as elaborating on the first one. Elaboration makes the non-raining scenario described by the first conjunct prominent, and the modal in the second conjunct selects it as its restrictor. Thus, we get the reading that if it is raining and it might not be, *given that it is*, the speaker is uninformed about the weather. Of course, the scenario described by the antecedent, on this reading, is inconsistent, so infelicity is predicted. Again, the fact that this reading is forced even though there are other possible, and vastly more plausible, interpretations, suggests that the effect of Elaboration is a matter of a linguistic convention.

Another piece of evidence for thinking that the effects of mechanisms of coherence are a matter of a linguistic convention comes from cross-linguistic variations that these mechanisms and these effects exhibit. To illustrate, while one might expect that understanding the second sentence in (13) as elaborating on the first one, and hence, interpreting (13) as being organized by Elaboration relation, is a matter of common sense reasoning—after all, it only makes sense that the scenario in which the addressee is being eaten provides a further elaboration of the one in which a wolf walks in—Asher and McCready (2007) point out that its Japanese translation of (13) is infelicitous unless it contains an overt discourse marker that signals the appropriate relation. Similarly, in the case of pronouns, languages that use different means of signaling prominence of candidate referents—for instance, word order or topic markers—are often more constrained in the kinds of shifts they associate with coherence.²⁷

The same kinds of considerations extend to other kinds of examples we have considered so far. For instance, take (23). Clearly, resolving 'local' in such a way so as to mean *local-to-Paris* would make the most sense in (23), especially assuming that all the parties to the conversation know well the relevant facts about Paris and Belgrade. Yet, it's surprisingly difficult to get this reading. In particular, so long as we understand the final sentence in (23) as a part of the chain

²⁷For a more detailed discussion of the cross-linguistic data suggesting that the effects of coherence on the interpretation of pronouns and modals are linguistically encoded, see Stojnić, Stone, and Lepore (2014) and Chapter 4, respectively.

of elaborations about the Belgrade trip, ‘local’ is interpreted as *local-to-Belgrade*. This is again surprising if Elaboration merely pragmatically suggests a certain interpretation. And the same point is illustrated by the following example as well:

(26) Some dogs chased a few cats. Most cats ran away. A few caught one.

Given the available information contributed by the first two sentences in (26), it would clearly make the most sense that ‘a few’ in the last sentence quantifies over the set of dogs that did the chasing. Yet, this is not the interpretation we find. Rather, we expect that ‘a few’ quantifies over the set of cats, resulting in an intuitively odd interpretation.

The same goes for incomplete definite descriptions. To illustrate, consider the following example:

(27) I left the living room and walked into the bedroom. The bedroom had no windows. The window was open.

In (27), one cannot understand ‘the window’ as being the window in the living room, even though the living room has been mentioned, and the bedroom ruled out. This is because we understand both the description of the window being open, and the bedroom being windowless as providing background circumstances to the event of entering the bedroom, and are thus restricted to the situation describing the bedroom. Since the descriptions are incompatible—there can be no window in a windowless room—(27) is infelicitous.

All these examples illustrate the same point—not only do the mechanisms of discourse coherence affect the interpretation of context-sensitive items, but their effect is not reducible to common-sense inference that draws holistically on the epistemic cues reflecting speaker intentions. In particular, these effects seem to draw on the knowledge of linguistic conventions that underscore the dynamics of prominence. These conventions fix interpretations in ways difficult for general world-knowledge or common-sense reasoning about speaker’s intentions to override.

1.5 Discourse Structure, Information Structure and Grammar

We have seen in the previous sections that mechanisms of discourse coherence affect the resolution of context-sensitive items in a systematic and rule-governed way. Of course, I do not mean to suggest that discourse coherence is the only relevant linguistic factor in context-sensitivity resolution. Other linguistic mechanisms can affect the interpretation in analogous ways. In particular, context-sensitivity resolution is also affected by mechanisms of information structure. Information structure marks the dimension of meaning signaling how the information conveyed with an utterance is related to the prior information, how it is organized with respect to categories such as focus/background, topic/comment, how it contrasts with relevant alternatives, and which move it makes in the dynamics of communication.²⁸ In English, information structure is typically signaled through intonation.²⁹

That information structure impacts the interpretation of context-sensitive items has been noted in the literature Evans (1980); Lappin (1993a,b, 1996). To illustrate, consider the following example, from Evans (1980):

- (28) Everyone has finally realized that Oscar is incompetent. Even HE has finally realized that Oscar is incompetent.

The observation is that, though normally ‘he’ could not be interpreted as co-referring with Oscar, due to the constraints on binding,³⁰ when there is contrastive stress on the pronoun ‘he’, one nevertheless understands ‘he’ as picking out Oscar. In a similar fashion, (as Kehler himself notes), one can improve (4), by putting contrastive stress on the pronoun *her*:

- (29) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships HER.

With this modification, we can understand the pronoun as referring to Thatcher. The focal stress

²⁸For more on information structure, see, e.g. Halliday (1967); Vallduví (1992, 1993); Ginzburg (2012); Roberts (1996a).

²⁹This fact is, again, language specific. Many languages use word order, or overt morphology, to signal a particular information structure. This suggests that the signals of information structure, too, are a part of the conventional linguistic repertoire of a given language. See Lepore and Stone (2015) for discussion of this point.

³⁰In particular, this would require a violation of *Principle C* of the Government and Binding Theory (Chomsky, 1981).

achieves a sort of implicit demonstration, making Thatcher the most prominent referent at this point in a discourse, again. How does this bear on the argument of the previous sections, that coherence relations determine the resolution of pronouns? We have seen in the discussion of (4) that Parallel relation requires that ‘her’ is resolved to the object of the previous clause—Reagan. Does the effect of focus we see in (29) violate this rule? The answer is negative. It is important to note that, while putting the focal stress on ‘her’ improves the overall felicity of the discourse, it also changes its overall interpretation. In particular, one no longer understands the discourse in (29) as comparing Thatcher’s and Bush’s respective attitudes (towards Reagan), but rather as explaining Bush’s respect for Thatcher’s opinions. Importantly, there is no reading according to which the discourse is organized by Parallel, yet the pronoun refers to Thatcher.³¹

This observation is further supported by examples like the following one from Kehler (2002):

- (30) The lawyer defended Bill against the accusations, and HE did too.

As Kehler notes, even with the contrastive stress on the pronoun, (30) does not have an interpretation according to which ‘he’ is understood to refer to Bill. The problem is that it is hard to understand (30) as organized by anything other than Parallel, and thus, the pronoun ‘he’, is expected to co-vary with the NP in the subject position in the first clause. Since, even with the contrastive stress, the Parallel reading is still the only one available, the infelicity of (30) is predicted: Parallel requires that ‘he’ be the lawyer, whereas the contrastive stress requires that it be someone else—Bill. In effect, it is as though the speaker is pointing in two different directions along with a single use of a pronoun.

Another important aspect of information structure has to do with signaling which part of an utterance connects to the prior contributions, and which is the new distinctive contribution that the utterance is making—that is signaling the distinction between the theme, or topic, and rheme, or comment. One way to think about this distinction assumes that the theme is signaling which topic, or question at issue, the utterance is answering, while the rheme provides an answer to it (Roberts, 1996a; Ginzburg, 2012). In English, topic/comment signaling also exploits prosodic accents, as illustrated by the following example:

³¹See Stojnić, Stone, and Lepore (2014) for further discussion of this point.

(31) John likes MARY.

(32) JOHN likes Mary.

(31) is fine in a context in which we are wondering whom John likes, say, Mary or Sue, but not in a context in which we are wondering who likes Mary, say, Bill or John; the opposite is true of (32).³²

This dimension of information structure likewise has an impact on the interpretation of context-sensitive items.³³ To illustrate, consider the following:

(33) John only likes MARY.

(34) John only LIKES Mary.

Even though *prima facie* (33) and (34) are just one and the same string of words, they convey two different things: (33) is true just in case John likes Mary and nobody else (out of the alternative individuals available in the context), whereas (34) is true just in case John likes Mary, but does not have any other attitude (out of contextually relevant ones) towards her. In other words, the focal stress signals which topic (33) and (34) address: whether the one regarding the possible objects of John's liking, or the one regarding the possible attitudes John might bear towards Mary. In turn, the topic constrains the domain of quantification of 'only', thus affecting the truth-conditions of (33) and (34).³⁴

These kinds of examples only illustrate that, though discourse structure affects the resolution of context-sensitive items, it is not the only dimension of grammar that carries this effect. The

³²See Roberts (1996a). Also, see Rooth (1992), for more on sentential focus.

³³See, for instance, Rooth (1992); Roberts (1996b, 2004); Beaver and Clark (2008); Roberts (2011); Schaffer and Szabó (2013); Schoubye and Stokke (2015) for more on this point.

³⁴Some authors have suggested that the apparent contextual variability of knowledge ascriptions seems to be affected by the topic or question under discussion, as well (Dretske, 1981, 1991; Schaffer and Szabó, 2013). They observe the contrast between examples such as (i-a) and (i-b), from Schaffer and Szabó (2013):

- (i) Claire has stolen the diamonds. Ann finds Claire's fingerprints all over the safe. So Ann says to Ben:
 - a. I know that CLAIRE stole the diamonds.
 - b. I know that Claire stole THE DIAMONDS.

Given the scenario described by (i), (i-a) seems true, but (i-b) false. But (i-a) and (i-b) appear to be one and same string of words; their difference is simply in prosody. In particular, the focal stress in (i-a) signals that the relevant topic—the question under discussion—is *who stole the diamonds*, whereas in (i-b), it is *what was stolen*.

examples above show that mechanisms of information structure contribute to resolution in similar ways. A natural question then arises about the relation between these two aspects of grammar.³⁵ Fully addressing this question would take us too far afield. For our purposes, it is important to recognize that while these two aspects of meaning both affect the resolution of context-sensitive items, and while they work in complementary ways, they are nevertheless distinct. In particular, it is often left implicit which question is currently under discussion at a particular point in a discourse, and even when there is an explicit question introduced, the speaker can choose to address it directly or indirectly, by addressing a more specific subquestion. We have seen that in English, prosody is exploited to signal the question that the utterance is addressing, but this is certainly not the only cue used to this effect. In a similar fashion, coherence relations among successive utterances can guide which topic a new utterance is addressing: for instance, cause-effect relations can organize the contribution to address why an event occurred (Explanation), or what resulted from it (Result), occasion relations can be used to address what happened during or after an event (Elaboration, Narration), while resemblance relations can be used to address how events compare to other events (Contrast or Parallel).³⁶

Clearly then, prosody, through signaling the topic, can serve as a cue as to which relation is operative at a particular point in a discourse. However, there are other cues that seem to conventionally signal that a particular relation is operative, and recognizing that a relation is operative can in turn help to recognize which question the speaker will chose to address next with an utterance. To see what is at stake, consider first the following (Stojnić, Stone, and Lepore, forthcoming):

(35) I knocked the glass off. It broke.

(36) The glass broke. I knocked it off.

While (35) and (36) are both an instance of a cause-effect relation, and indeed could both be describing the same exact events in the world, the former is understood as organized by Result relation and the latter by Explanation. And while one might think that we understand the second sentence in (36) as explaining the event described in the first because this interpretation “makes the most

³⁵See e.g. Roberts (2004).

³⁶See Lepore and Stone (2015); Roberts (2004); van Kuppevelt (1995) for a further discussion of this point.

sense”, as argued in Stojnić, Stone, and Lepore (forthcoming), the particular organization seems to involve an element of conventionality. To appreciate this, consider the following example:

(37) John shot Bill. He was dead.

Though certainly death can result from a shooting event, it is hard to understand the second sentence in (37) as describing the result of the event of shooting. Rather, we understand (37) to convey that Bill was already dead when shot by John. We typically understand stative descriptions as connected by Background relation to descriptions of ongoing events.³⁷ Clearly, this is not just a matter of common-sense or world-knowledge—after all, such considerations would favor the reading of (37) according to which the second sentence is reporting a result of the event described by the first one. Thus, both information structure, and discourse coherence, though working in complementary ways, are dependent on a variety of linguistic cues, which in turn cross-cut the world-knowledge, and common-sense and plausibility considerations.

Finally, I close by addressing a possible objection to the meta-semantic account I defended. The view I defended maintains that mechanisms of discourse coherence affect the resolution of context-sensitive items by making certain values of contextually relevant parameters prominent as a matter of linguistic convention. However, an intentionalist, or a pragmatist might then object that, even if coherence relations conventionally promote certain values of contextual parameters, it is ultimately the speaker’s intention, and/or general world-knowledge, which determines which coherence relation is organizing a discourse, and so, we still need to rely on intentions for meaning determination. In other words, one might argue that even if the metasemantics of, say, a pronoun, is a matter of convention, the metasemantics of coherence relations might still involve speaker intentions, and other contextual cues.

To respond to this kind of a worry, it is important to note that I am not suggesting that coherence relations themselves are context-sensitive items, that need to be resolved in a context, even though discourses may be *ambiguous* with respect to the coherence relation they harbor. To see what is at stake recall (3), repeated below as (38):

³⁷See Asher and Lascarides (2003); Kehler (2002).

(38) Phil tickled Stanley, and Liz poked him.

What I suggest is that discourses like (38), are *ambiguous* between (at least) the logical form harboring Parallel, and the one harboring Result. Thus, determining a coherence relation amounts to choosing between distinct fully specified logical forms of a discourse.³⁸ And while general reasoning about interlocutors intentions might well be potentially used as a tool of disambiguation, speakers often resolve ambiguities by relying on shallow linguistic cues that favor an interpretation within a set of alternative interpretations delivered by the grammar. This is reflected in examples like (37). The fact that discourse structure can often be ambiguous is hardly a vindication of an intentionalist meta-semantics.³⁹

1.6 Conclusion

The often assumed intentionalist, pragmatic account of the resolution of context-sensitivity is *prima facie* very plausible. After all, speakers typically speak with an intention to convey certain thoughts to their audience, and it is moreover natural to think that in using context-sensitive language, they normally use it with an intention to convey the meaning that would help express their thought. Yet, while all this is true, it does not amount to an intentionalist meta-semantics. While speakers typically intend to use context-sensitive language to convey the intended content, this is compatible with the rules of language, rather than an appropriate intention, or a set of general contextual cues, being what ultimately determines the meaning of such expressions.

Indeed, I have argued that context-sensitivity is resolved by appeal to linguistic mechanisms that govern the dynamics of prominence of particular values of contextual parameters associated with context-sensitive expressions. In particular, we have seen that mechanisms of discourse structure and information structure govern the dynamics of prominence in context, and do so as a matter of linguistic convention. The resulting conventionalist meta-semantics does not deny that speakers normally have intentions to convey a particular meaning with a particular use of a context-sensitive expression. However, speakers are constrained by the rules of language in what they can intend

³⁸See Stojnić, Stone, and Lepore (2014).

³⁹For a more detailed discussion, see Stojnić, Stone, and Lepore (forthcoming).

to convey with the use of a context-sensitive expression.⁴⁰ That they should be so constrained is telling: it is because the rules of language determine a particular resolution of a context-sensitive expression, that we cannot use it as flexibly as one would expect on a traditional view.

⁴⁰Of course, the intentionalists do not have to deny this particular point. For instance, King (2014a,b) makes and endorses this point. But, insofar as we are as thoroughly constrained by linguistic rules with respect to what we can intend, and insofar as the rules are as pervasive as I have suggested they are, it seems that the presence of an appropriate intention will neither be a necessary, nor a sufficient condition for meaning-determination.

Chapter 2

Discourse and Logical Form: Pronouns, Attention and Coherence¹

2.1 Introduction

The meaning of the pronoun ‘he’ does not seem ambiguous in the way that, say, the noun ‘bank’ is. Yet what we express in uttering ‘He is happy,’ pointing at Bill, differs from what we express when uttering it, pointing at a different person, Sam. ‘he’ in the first case refers to Bill, in the second to Sam. And further with an utterance of ‘A man always drives the car he owns,’ or ‘A man came in. He sat down,’ without an accompanying pointing gesture, ‘he’ refers not at all; its interpretation, instead, co-varies with possible instances in the range of the quantifier ‘A man’. Those uses in which the pronoun is interpreted referentially are usually called ‘demonstrative,’ and those in which it is interpreted non-referentially are usually called ‘bound’. We immediately face two questions: what is the semantic value of a pronoun in a context, given that it can have both bound and demonstrative uses? And, since this value can vary with context, how does context fix it? The first question concerns the semantics of pronouns, and the second their meta-semantics.

The received view answers the first question by positing an ambiguity: bound uses of ‘he’ function as bound variables; demonstrative ones, by contrast, are referential. The received view’s answer to the second question is that the semantic value of a bound use co-varies with possible instances in the range of a binding expression, and for demonstrative uses, linguistic meaning constrains, but does not fully determine, its semantic value in context. For instance, in uttering ‘He is happy,’ pointing at Bill, its linguistic meaning is supplemented by a pointing gesture in fixing the referent of ‘he’, and, in general, a surrounding contextual cue, either an accompanying demonstration, or an appropriate and adequately transparent speaker intention, or even both, is required to fix the referent of

¹This chapter is a paper jointly co-authored with Matthew Stone and Ernie Lepore, (Stojnić, Stone, and Lepore, 2014).

a demonstrative pronoun.² Such supplementation is understood to be extra-linguistic or pragmatic, and, most certainly, not linguistically encoded.

This paper challenges the received view on both counts: we reject the ambiguity thesis and any meta-semantics that invokes extra-linguistic determinants of semantic values in context. We argue that the semantic value of a demonstrative pronoun, like that of a bound variable, depends on, and can co-vary with, antecedent expressions; and further, the linguistic meaning of a pronoun, on both uses, fully determines its semantic value in a context. In particular, ‘he’ is like the first person pronoun ‘I’ since for both context automatically fixes a semantic value on any occasion of use. Just as ‘I,’ given its character, refers to the speaker of a context, so too, contra the received view, ‘he’, given its character, determines its semantic value as a function of a fixed feature of context. On first pass, its semantic value is the male at the “center of attention” in a context. (We will precisify “center of attention” below.) In contrast to the received view, what lies at the center of attention is *not* governed by extra-linguistic mechanisms, speaker intentions or demonstrations, but entirely by linguistic rules. More precisely, linguistic mechanisms govern the dynamics of prominence in a discourse, rendering certain semantic values preferred for pronoun resolution in a context. The meaning of a pronoun fully determines its referent in a context—it picks out whatever linguistic rules determine to be “at the center of attention” at that point in a discourse. In what follows, we spell out these rules and how they affect pronoun resolution.

As a clarification, note that our claim that linguistic conventions play a role in pronoun resolution is not altogether novel. Most theories of pronoun resolution maintain linguistic conventions play *some* role in fixing semantic value. Many believe conventions play a role in constraining what speakers can even be taken to intend—e.g., in a normal context, one cannot reasonably intend to refer to a woman using ‘he’. We believe, however, such constraints undermine the main ingredient of these accounts, namely, that speaker intentions, or additional epistemic cues, play a crucial role in fixing the semantic value of a pronoun. As we have argued elsewhere, the meaning of ‘he’ is determined in conformity with linguistic conventions *regardless* of the presence or absence of a

²Proponents of the received view are many. Notable ones are Kaplan (1989b,a); King (2014a,b); Neale (2004, 1990); Reimer (1992). There are important differences between these different versions, but they do not matter for our present purposes. Our focus is primarily here on our positive view. For our criticisms of such views, see Stojnić, Stone, and Lepore (2013).

speaker intention, or an epistemic cue.³

Briefly, to see what we mean, suppose while Sue is making herself highly salient by jumping up and down, you have been asked, ‘Who came to the party last night,’ to which you respond, ‘Mary came to the party. She had fun.’ In uttering ‘she had fun’, you make no overt gesture towards Sue, not even a glance towards her, but nevertheless you intended to refer to Sue. It is our intuition that what you have said is *Mary* came to the party and had fun, despite your intention or Sue’s salience. In short, contrary to the received view, an intention is neither necessary nor sufficient to fix a referent on an occasion of use. To say in such cases that linguistic conventions constrain what can reasonably be intended is to agree with us that these conventions fix a referent on an occasion of use.

We also note that various authors argue that linguistic rules do play a role in resolving bound uses of pronouns, and more generally, so-called anaphoric uses, where pronoun resolution is dependent on a prior linguistic expression.⁴ Our view goes beyond, as well as diverges from, these authors in several respects.⁵ For one, we argue for a unique linguistic meaning (a character) for *all* uses of pronouns—demonstrative, bound and anaphoric alike—and that, for *all* these uses, the interpretation of a pronoun is fully determined by linguistic rules. Secondly, we elucidate a wide range of linguistic mechanisms that affect the dynamics of prominence in a discourse, which have either gone unnoticed in the current debate, or otherwise are confused with pragmatic effects on interpretation.

³We have provided extensive criticism of such accounts elsewhere (Stojnić, Stone, and Lepore, 2013).

⁴Our distinction between bound and demonstrative uses of pronouns does not perfectly coincide with one between demonstrative and anaphoric uses of pronouns, where an anaphoric use is any use the interpretation of which depends on an expression introduced earlier in the discourse. Namely, some anaphoric pronouns are clearly referential. To wit, consider: ‘Bill came in. He sat down.’ Since we will argue that all uses of pronouns have the same semantics, and are resolved by appeal to the same linguistic mechanism, all demonstrative, bound and anaphoric uses are, according to us, alike.

⁵Theorists who stress linguistic constraints on the interpretation of an anaphoric pronoun to various degrees include Chomsky’s Binding Theory (cf. Chomsky (1981) as well as various forms of dynamic semantics (e.g., Kamp (1981) and Heim (1982))). Almost everyone agrees there are linguistic constraints on the resolution of anaphoric pronouns. Fiengo and May (1994, 2006), for instance, hold that any two co-indexed expression–tokens are tokens of the same expression type, and are coreferential as a matter of grammar, and point to a number of linguistic constraints on co-indexed, including, but not limited to, those of Binding Theory. Though we agree with these authors that anaphoric dependencies are a matter of linguistic conventions, we both go beyond, and diverge from them, in several respects. First, we argue *all* pronoun resolution—whether demonstrative, bound or anaphoric—is determined by a linguistic rules. Secondly, the theories described allow for a single pronoun to be represented by different expressions in the logical form (corresponding to different indices). This suggests an ambiguity in the word ‘he’ like, e.g., the word ‘bank’. We, by contrast, argue that ‘he’, has a unique, unambiguous linguistic meaning, and that all its occurrences are represented by a token of the same expression-type in logical form. Thirdly, and most importantly, we characterize a wider spectrum of hitherto unappreciated linguistic mechanisms that govern pronoun resolution.

Given the mechanisms that maintain the dynamics of prominence in a discourse, any occurrence of a pronoun picks out its semantic value automatically, according to its character—selecting whatever is most prominent, or “at the center of attention”.

Finally, we want to register that our view has far-reaching philosophical consequences beyond the semantics and meta-semantics of pronouns. Philosophers regularly invoke context-sensitivity in analyzing philosophically interesting terms, such as ‘know’, ‘good’, ‘might’ and ‘all’, where their semantic values are alleged to be determined by contextual parameters such as the epistemic standards, standards of precision, value commitments, implicit restrictions on quantifier domains, etc. In building theories about how the values of these parameters are fixed in context, philosophers are invariably guided by the received view of the meta-semantics of pronouns: namely, that the values of these parameters are determined jointly by linguistic meaning and non-linguistic cues.⁶ Often, in fact, both the linguistic analyses of these terms, and the philosophical arguments exploiting these analyses, depend on this model of context dependence. If we are right, however, the standard view is misguided even for pronouns. This suggests we should be wary of meta-semantical theories about philosophically interesting expressions built on this model. In fact, our view recommends a project of its extension to other cases; if so, philosophers will have to rethink the role of context-sensitivity in their arguments.⁷

In what follows, we develop an approach to the semantics and meta-semantics of pronouns we call the *Attention–Coherence Approach*; it contains two components—a ranking of candidate interpretations according to relative prominence (attention), and implicit mechanisms affecting this ranking (coherence). The role of both components will be made precise in what follows. We begin in the next section to develop the attention component; the theory will not receive its full shape until after we present the coherence component in later sections.

⁶That the meta-semantics of context-sensitive expressions generally should be modeled on the received view of the meta-semantics of demonstrative pronouns is often either implicitly assumed, or explicitly endorsed. See. e.g. Dowell (2011); Cohen (1998); King (2014a,b); King and Stanley (2005); Neale (2004, 1990).

⁷See Chapter 1.

2.2 Attention

On a rough first pass, our theory is that pronouns, as a matter of linguistic meaning, (*character*, if you like), always pick out whatever is ‘in the *center* of attention’ in a discourse. To capture this, we borrow resources from *Centering theory*,⁸ according to which, candidate resolutions of a pronoun in a discourse are ranked according to their relative prominence: those higher in the ranking are preferred as interpretations of pronouns over those lower in the ranking. Call this ranking the *attentional state* of an ongoing discourse.⁹ As a discourse progresses, utterances bring new candidates into focus and adjust the prominence of old ones, thereby changing the attentional state. We shall argue that pronouns, as a matter of meaning, always pick out the top-ranked candidate in the current attentional state, and further, that the attentional state itself is maintained through linguistic rules.

To motivate the dynamics of an attentional state, we begin with (1):

- (1) A man walked in. He sat down.

One reading is that ‘He’ is interpreted as co-varying with ‘A man’. On this construal, (1) is true just in case some man walked in and sat down. On a second reading, the speaker, perhaps uttering ‘He’ pointing at a small child, Billy, is interpreted as referring to Billy. The discourse is true just in case a man came in and *Billy* sat down.

Speakers can use (1) in different ways and hearers can point to reasons why certain interpretations are recovered, and others not. Pronoun resolution is guided by an implicit organization that knits together information in discourse. On its anaphoric reading, the discourse begins with a description involving ‘A man’ and proceeds directly to develop a narrative: accordingly, ‘He’ is interpreted as dependent on ‘A man’. On its deictic reading, however, the speaker marks a transition from talk of past events to the present situation with a demonstrative gesture: on this organization, ‘He’ is interpreted as dependent on the individual demonstrated, and not on ‘A man’. We shall argue that the difference between these interpretations of (1) is a difference in the dynamics of attention,

⁸See Bittner (2014); Grosz, Joshi, and Weinstein (1995); Sidner (1983).

⁹Centering theory is traditionally understood as pragmatic; its preference for higher ranked entities is one of many non-linguistic cues that can potentially give evidence about the intended referent of a demonstrative pronoun. Our view, by contrast, is semantic: the ranking is maintained and updated by linguistic rules. The character of a pronoun, given this ranking, determines its semantic value.

guided by linguistic mechanisms. It is because the description ‘A man’ makes its potential witnesses prominent, and the discourse proceeds to develop a narrative about whomever is a witness for ‘A man’, thus keeping it in the center of attention, that ‘He’ is understood as co-varying with potential witnesses in the domain of the existential ‘A man’. Similarly, it is because the pointing gesture shifts the attention from potential witnesses for ‘A man’ to the child pointed at, Billy, that ‘He’ on the deictic reading is understood to refer to Billy. These changes in attention, governing the prominence of potential resolutions for a pronoun, we shall argue, are guided by linguistic mechanisms. To make our proposal clear and precise, we shall endorse a particular way of formalizing our key ideas. But our philosophical point—that pronominal interpretation is fixed through linguistic rules—is independent of our formalization, and could be implemented in other frameworks. Accordingly, our preferred formalization is not essential to our argument.¹⁰

Since on one interpretation of (1), ‘He’ is understood as ‘bound’ by the indefinite ‘A man’, we adopt a model for bound variables that allows us to capture relevant interpretive dependencies across discourses. One such model cleaves closely to the familiar Tarskian machinery for variable binding, according to which the truth of a formula is specified relative to an assignment of values to variables. A variable, x_i , is interpreted by retrieving the i^{th} element of a current assignment. Bound variables have dependent interpretations because quantifiers vary the assignments in force within their scope. One shortcoming is that interpretation so-construed is limited to the syntactic scope of a quantifier. But quantifiers can introduce dependencies that persist across subsequent sentences, as in the anaphoric reading of (1). We need a framework for extending the familiar notions of a variable and binding to the kind of dependent interpretations in (1). A fix is provided by *dynamic semantics*.¹¹

In dynamic semantics, just as in a Tarskian framework, variable binding creates dependent interpretations by varying assignment functions. However, unlike the Tarskian approach, dynamic semantics assigns truth conditions to formulas relative to a *pair* of assignment functions, not to a

¹⁰For alternative frameworks, see e.g., Elbourne (2008); Fiengo and May (1994, 2006).

¹¹See Dekker (2011) for an accessible overview of dynamic semantics.

single one. The first assignment accounts for interpretive dependencies potentially already available in the discourse for interpreting a formula; the second, however, accounts for interpretive dependencies the formula itself makes available for interpreting subsequent discourse. A formula is interpreted as an instruction to update the available interpretive dependencies: it augments the potential interpretive dependencies from prior discourse, encoded in the input assignment, with the additional interpretive dependencies a formula introduces for interpreting subsequent discourse, so as to create the output assignment. Just as in Tarskian semantics, in dynamic semantics, we interpret a bound variable x_i by retrieving the i^{th} element of the input assignment, but now we can define the existential quantifier in a way that allows for alternative *output* assignments, so that the interpretation of subsequent variables can continue to take a bound interpretation throughout the subsequent discourse. Overall, a formula is true on a given input assignment if an output assignment makes it true.¹²

Due to the existential quantifier in (1), the formula modeling the first sentence in (1) is true on a given input assignment just in case there is an output assignment according to which some unspecified man—a witness for the existential—walked in.¹³ The second sentence is modeled with a formula that takes this output assignment as a new input; ‘he’ can then be interpreted as a bound variable that takes as its value whatever the witness for ‘a man’ is in this new input assignment, resulting in the anaphoric reading of (1). The discourse is true on the initial input assignment just in case a man walked in and sat down.

But how does the interpretation of ‘he’ in (1) get bound by the existential quantifier ‘A man’? Both in Tarskian and dynamic semantics alike, variables behave like temporary names. A quantifier like $\exists x_i$ re-defines what x_i names within its scope; we can then interpret occurrences of x_i as local names for potential witnesses of the quantifier. So, if ‘A man’ is translated with ‘ $\exists x_i$ ’, then ‘he’

¹²In this framework, we interpret formulas via updates that relate an input assignment g to an output one h . The simplest update contributes information fixed by a condition C , written as ‘ $[C]$ ’ and interpreted as a partial identity relation between assignments. If g is an input assignment, and h an output one, $\llbracket [C] \rrbracket(g, h)$ iff $g = h$ and C holds on the interpretation of variables given by g . In standard fashion, a sequence of updates H and K is represented by a single one $H;K$ that performs H followed by K . An update is true for an input assignment iff it is related to some output one by the update relation. In addition to the simplest updates, we need updates that affect the prominence ranking of the candidate referents. We introduce these presently. As noted earlier, we exploit this particular formal system just for the sake of concreteness and clarity; our philosophical ideas can be implemented in other frameworks. For detailed exposition of the formal system see the Appendix (§ A.1).

¹³Any output assignment that makes it true that there’s a man who walked in will do. This is the sense in which the witness for the existential is an “unspecified” man—any one will do.

gets translated with an ‘ x_i ’, thus looking for the i^{th} element of the input assignment; this element, given the prior discourse, will be required to be a man who walked in. By contrast, since we want to model prominence in a discourse, we will organize interpretive dependencies *not* with names, but by *prominence*. To achieve this, we treat assignments as stacks in the sense of theoretical computer science. Each assignment specifies a sequence of possible individuals ordered by prominence. The most prominent individual is in initial position—the top of the stack—and ones in subsequent positions are deeper in the stack, thereby, receding in prominence. Quantifiers introduce new possibilities for dependent interpretations in subsequent discourse by pushing values onto the stack: inserting them at a specified position, for example, at the top, thereby, decreasing the prominence of other candidates in the discourse by one position. In this way, they vary output assignment functions. We treat a variable not as a temporary name, but as a marker of prominence; we can define a variable that picks out the top of the current stack, and thus, co-varies in its interpretation with whichever quantifier most recently pushed a new value there. We use ‘@’ as a variable, interpreted relative to an assignment g , as specifying the top-ranked element of g . (The mnemonic is ‘@’ is at the center of attention.)

The meaning of ‘@’ is a first approximation to the meaning of an English pronoun. To illustrate, let ‘ $[man(@)]$ ’ be a condition that requires that whatever is at the center of attention be a man. And let ‘ $\langle\alpha\rangle$ ’ indicate a dynamic existential quantifier, which changes the input assignment by introducing a new unspecified top-ranked individual at the center of attention, with all other candidates demoted one position in the ordering.¹⁴ This update is a first approximation to the meaning of the English indefinite article. It affects the prominence ranking, changing what is at the center of attention. (1), on its ‘anaphoric’ interpretation, is represented as (2):

$$(2) \quad \langle\alpha\rangle; [man(@)]; [walk.in(@)]; [sit.down(@)]$$

(2) begins with an existential quantifier, corresponding to the indefinite article, that introduces its witness as the new top-ranked resolution for subsequent variables; the witness is constrained to be a man, to have walked in, and—once picked up by the anaphoric pronoun—to have sat down. Thus,

¹⁴This is a non-deterministic update in the jargon of dynamic semantics. $\langle\alpha\rangle$ relates an input assignment function g to an output one h iff h potentially differs with g in the top position, and for every subsequent position i , $h_i = g_{i-1}$.

the formula is true just in case some man walked in and sat down.¹⁵

By contrast, consider a deictic reading of (1), uttered pointing at a child, as in (3):

(3) A man walked in. He [pointing to a child, Bill] sat down.

We know its indefinite NP puts its witness at the top of the stack, while the pronoun, as a matter of meaning, resolves to the top-ranked candidate. Yet, ‘He’ in (3) doesn’t resolve to the witness for ‘A man’. This is because of a further *shift* in attention—one triggered by a pointing gesture. This shift updates the attentional state so that the entity indicated by the pointing gesture becomes top-ranked.

We introduce a family of updates, written ‘ $\langle \pi c \rangle$ ’, where ‘ π ’ corresponds to the act of pointing and ‘ c ’ to some individual being pointed at.¹⁶ This update stores c as the top-ranked entity, and (as always) pushes all others down a position in the ordering. It thereby represents the effect of the pointing that accompanies the use of ‘he’. (3) is represented formally as (4):

(4) $\langle \alpha \rangle; [man(@)]; [walk.in(@)]; \langle \pi b \rangle; [sit.down(@)]$

In (4), we formalize ‘He sat down’ with ‘ $[sit.down(@)]$ ’, just as in (2). But in its context in (4), the condition is true just in case Bill sat down. The intervening expression ‘ $\langle \pi b \rangle$ ’, corresponding to the act of pointing at Bill, updates the attentional state. As a result, whomever the speaker has pointed at is at the top of the assignment when ‘ $[sit.down(@)]$ ’ is interpreted. In other words, a pointing gesture updates the prominence ranking, making Bill the center of attention.

As in (4), the shift of attention provoked by a pointing gesture is reflected in logical form. However, traditionally, the effect of a pointing gesture accompanying the use of a demonstrative pronoun has been treated as a result of a pragmatic process in which the hearer comes to recognize the speaker’s referential intention by exploiting the epistemic cue provided by the speaker’s pointing gesture. The most plausible way to make sense of the speaker having chosen to accompany her use of the pronoun with a pointing gesture is that she intends to refer to some object in the indicated direction. Crucially, the pointing gesture does *not* render the demonstrated object prominent as a

¹⁵The formula is true for an input assignment g iff it is true for some output assignment. Since (2) has no free occurrences of a variable, its truth does not depend on an initial input assignment at all.

¹⁶More generally, ‘ $\langle \alpha \rangle$ ’ corresponds to the contribution of indefinites, and ‘ $\langle \pi c \rangle$ ’ to definites, where ‘ c ’ is a denoting expression. We do not offer here an account of definites other than pronouns.

matter of a convention. It serves as an epistemic cue that helps manifest the speaker's intention because it makes sense to point at an intended referent.

We reject this view. The change of attention contributed by the pointing gesture in (3) is and *should be* reflected in logical form, we maintain, because the sorts of gestures that secure pronoun resolution in these cases do so as a matter of grammar. Our claim is that rules of language dictate that a demonstrative act of pointing introduces a new candidate referent at the top of the stack.¹⁷ We submit several reasons for treating deictic gestures as governed by grammar.

First, as argued by Kendon (2004), non-verbal means of indicating an entity are governed by rules sensitive to form, meaning, and the relationship with ongoing speech. English speakers seem to count deixis as well-formed only when the pointing action is synchronized appropriately with the prosody of the accompanying utterance. They often repair utterances when their performance fails to align speech and gesture in time, as one would expect if the requirement for synchronicity were dictated by an underlying convention.¹⁸

In addition, Kendon observes that although English speakers can use a range of hand shapes when they indicate an object, their particular choice affects the semantic contribution they make. A gesture with the index finger and thumb extended, and the other fingers curled closed, uses the direction of the index finger to single out an object as an individual distinct from its alternatives. By contrast, a gesture with a flat hand open toward the audience, its four fingers extended in a tight line, uses the direction of the fingers to exhibit an object as a representative of a broader class. Moreover, there are many possible gestures that are *not* typically used as deictic gestures by English speakers though they could have been. English speakers use the thumb in the “thumbs up” hand shape, with the thumb extended from a tight fist, to demonstrate *movement* in the direction that the thumb points, but normally *not* to demonstrate *an object* located in that direction, though nothing in the gesture itself precludes it from playing this function. The arbitrary fact that some spatially directed actions are taken to indicate objects—while analogous ones are not—reveals convention at work in English speakers' demonstrations. Numerous languages, unlike English, allow speakers to indicate objects by deictic gestures of their lips, and not an extended index finger. Cuna, a language from Panama,

¹⁷For a grammar integrating gesture and speech, see Alahverdzhieva and Lascarides (2011).

¹⁸This is similar to how prosodic focus is grammatically constrained to appear at a particular position with a particular contour to contribute a semantic contribution.

is one of many examples.¹⁹ Such conventionality and cross-linguistic variability is a hallmark of linguistic meaning.²⁰

Further, in typical cases, acts signaling shifts in attention are *indispensable* for the appropriate interpretation. In (3), explicit signaling of a shift of attention is necessary to establish a deictic reading, even to a particularly salient individual in the situation of utterance; otherwise, grammar seems to commit the speaker to an anaphoric reading of the pronoun. Even with Bill jumping up and down in front of the interlocutors, attracting their psychological attention, unless an overt signal establishes Bill as the referent of ‘he’, the audience will prefer the default anaphoric reading. Indeed, some means of explicit signaling of a shift of attention is required even when a linguistic antecedent is *unavailable*, such as when a pronoun is used deictically in an utterance that initiates a conversation. If deictic gestures were epistemic cues that suggest real-world salience, this would be quite mysterious.²¹

Moreover, the linguistic contribution of a pointing gesture clearly affects the inference patterns (1) licenses. The inference from (1), on its anaphoric interpretation, where ‘He’ is understood as co-varying with ‘A man,’ to ‘Some man sat down’ is valid; but, on its deictic reading, it is not. Given that, as we have argued, the contribution of a pointing gesture is underwritten by linguistic conventions, and that this contribution affects which inference patterns (1) can license, we submit that the contribution of the pointing gesture should be reflected in the logical form of (1). Indeed, (2) entails that some man came in and sat down, as a matter of logical form; but (4), on the other hand, does not.²²

Finally, we deny that pragmatic mechanisms determine the *demonstratum* of a demonstrative gesture. We represent the pointing in (3) as a *pointing at Bill*, not as pointing *simpliciter*; with an

¹⁹See Wilkins (2003).

²⁰We understand convention *a la* Lewis (1969), though we need not subscribe to all of his conditions to register our point.

²¹Grammar specifies a diverse set of resources for raising entities to prominence (e.g., deictic gestures and indefinite NPs). Our formalism is expressive enough to make such resources available even for utterances that are not following up an ongoing discourse. Our prediction is therefore not that deixis without demonstration is impossible at the beginning of a discourse, but that it succeeds only for utterances that recognizably accomplish acts that independently require construing the referent as the center of attention in the current state of the ongoing discourse. We provide a more detailed account of deixis without overt accompanying demonstrative gestures in Stojnić, Stone, and Lepore (2013).

²²An inference from a formula K to H is valid iff for any assignments g , and h , if K relates g to h , then there is an assignment i such that H relates H to i . Formal definitions are in the Appendix (§ A.1).

intended demonstratum to be determined by context. In other words, we deny pointing is semantically interpreted as having a uniform character that, given a context, determines a referent. Pointing at Bill is not interpreted as *pointing at x*, where *x* is contextually determined to be Bill. Rather, pointing gestures are ambiguous between multiple possible forms, e.g., *pointing at Bill*, *pointing at his shirt*, *pointing at his button*. Acts of pointing are, on our account, ambiguous. The interpretive work in (3) comes in settling whether the *form* of the gesture is pointing at Bill or pointing at something else. But *disambiguation* is pre-semantics, in Kaplan's 1989a sense: it involves the interpretive work needed to settle *the linguistic form* of an utterance. This should not be confused with what is involved in the semantic interpretation of a given linguistic form. Though, according to us, epistemic cues play no role in semantic interpretation, they *can* allow observers to recognize *which form* of a pointing gesture is in play. This is the role they'd play in helping to disambiguate *any* ambiguity, e.g., in disambiguating a use of 'bank'. And, as with other ambiguities, conventions governing demonstrative actions constrain possible disambiguations. As we have seen, a flat hand shape with the fingers towards the audience allows for a certain range of interpretations, but not others; similarly, for an extended index finger.

To sum up: we have argued there are reasons to treat the effect of an attention shift triggered by an indefinite NP or an act of demonstration as rule governed and as represented in logical form. The recognition of these systematic effects on the prominence ranking of candidate referents in a context allows us to assign one meaning to each pronoun that fully determines its resolution in a given context—roughly, pronouns pick out the most prominent candidate, the one at the center of attention. Refinements, however, are obviously required; we turn to them directly.

2.3 Refinements

That 'he', as a matter of meaning, takes on the most prominent candidate interpretation cannot be the whole story. 'He' in (5) cannot be interpreted as co-varying with 'A girl'. It must be resolved to a *male*.

(5) A girl came in. He sat down.

'She' in (6) cannot be interpreted as co-varying with the witness for 'Some girls'. 'she' must have

a *singular* interpretation.

(6) Some girls came in. She sat down.

‘we’ and ‘they’ in (7) cannot be interpreted as co-varying since ‘they’ requires a third-person interpretation *disjoint* from speaker and addressee, and ‘we’ requires a first-person pronoun interpretation that includes the speaker.²³

(7) We came in. They sat down.

Pronouns come with person, gender and number requirements that must be satisfied in pronoun resolution. These requirements are *linguistic*: it matters whether a speaker utters ‘he’, ‘she’ or ‘they’ because they differ in meaning.²⁴

Further, in (8), ‘him’ cannot semantically co-vary with Paul, even though Paul would normally be an eligible referent for ‘him’.

(8) Paul met him.

The explanation is that reference is constrained by syntactic principles—in this case, Principle B of Chomsky’s Binding Theory, which requires non-reflexive pronouns to be free in their governing category. Roughly, this means the non-reflexive pronoun must not be bound by an expression in the same clause, that is, that it cannot share a clause with its own antecedent.²⁵ So, in (8), where ‘him’

²³Sometimes a speaker can use ‘he’ to refer to himself or ‘they’ to pick out a group including himself. Witness a Perry-type scenario, where the speaker seeing, but not recognizing, himself in a mirror says, “He is the messy shopper” (Perry, 1979). Such cases are arguably marked in the sense that the exploitation of an extant rule makes for the surprising effect. Others suggest that the disjointedness condition is pragmatically implicated, not grammatically dictated (see Fiengo and May (2006)). By contrast, we treat number requirement as a part of the character of a pronoun. We agree that sometimes the speaker can manage to ‘speaker-refer’, in the sense of Kripke (1977), to himself with ‘he’, or to a group including himself with ‘they’. Also, sometimes an audience can understand that the speaker meant to pick out himself with ‘he’, or some group including himself with ‘they’, in the same sense as we can sometimes understand a speaker who mistakenly attempted to pick out a male with ‘her’. This is so even though the character of a pronoun encodes number, gender and person features. However, we could modify our account, if one insisted on excluding these features from the meaning of a pronoun. Whichever property *P* specifies the character of ‘he’, we maintain its referent is the top-ranked entity satisfying that property.

²⁴Again, this is true regardless of whether a speaker can manage to ‘speaker-refer’ in the sense of Kripke, to a woman with the pronoun ‘he’ (or a definite ‘that man’).

²⁵More precisely, the antecedent of a non-reflexive anaphoric pronoun must not be local, or c-command the pronoun. For precise definitions, see Chomsky (1981).

and ‘Paul’ are clausemates, ‘him’ cannot refer to Paul on pain of violating Principle B.

With this in mind, we say a pronoun denotes the value stored at the highest-ranked position on the stack that respects grammatical features and applicable syntactic principles. To interpret a pronoun, we consider candidates in order of prominence within the current attentional state of the discourse, until we find one that satisfies the operative linguistic constraints.²⁶ 2.3.1 integrates the constraints of semantics and syntax for the pronoun ‘he’.

Definition 2.3.1.

When interpreted relative to an assignment g modeling the available dependent interpretations on an occasion of use, ‘he’ denotes g ’s highest-ranked entity that is singular, masculine, and disjoint from the speaker and addressee of the utterance, and that yields an interpretation where the occurrence of the pronoun is free in its governing category.²⁷

To incorporate 2.3.1 formally, let ‘he’ be a predicate representing the constraints associated with the English third person singular male pronoun and ‘@he’ be an individual expression denoting the highest-ranked entity from the current assignment that satisfies the property denoted by ‘he’. (2)—(3) become (9)–(10).

(9) $\langle \alpha \rangle; [man(@)]; [walk.in(@)]; [sit.down(@he)]$

(10) $\langle \alpha \rangle; [man(@)]; [walk.in(@)]; \langle \pi b \rangle; [sit.down(@he)]$

However, even with these qualifications, (11) and (12) point to deficiencies.

(11) A man met Sam. He greeted him.

(12) John was disappointed with Tim.

a. He fired him.

²⁶The proposal to resolve the pronoun to the top-ranked element on the stack involves *search* rather than a simultaneous imposition of potentially competing constraints; the pronoun picks out the value stored at the highest ranked position on the stack *that satisfies the constraints*.

²⁷Roberts (2002) proposes similar rules for demonstrative pronouns, but with crucial differences. Instead of a linguistically governed notion of attentional prominence ranking, she uses a liberal notion of contextual, “real-world” salience. Moreover, she treats semantic constraints on pronoun resolution as presuppositions of a pronoun, but we understand them as part of character.

- b. He did sloppy work.

In (11), English speakers would normally resolve ‘He’ to the man introduced in the first sentence and ‘him’ to Sam.²⁸ Although the syntactic constraint on binding explains why English speakers do not treat the pronouns in (11) as co-referential (violation of Principle B), (11) it won’t explain why the first pronoun is resolved to the man introduced in the first sentence, and not to Sam. The explanation can’t be that only ‘A man’ and not ‘Sam’ updates the top of the stack to allow for a new dependent interpretation; a name can also affect attention focus.²⁹ So, why is the former, and not the latter, the highest-ranked candidate for ‘He’? Similarly, given what we’ve said so far, (12) is puzzling. In (12-a), ‘He’ is naturally resolved to John, but in (12-b), to Tim. Clearly, both referents can’t be top-ranked simultaneously. What makes John prominent in (12-a) and Tim in (12-b)? We take up these challenges in turn.

For (11), an explanation is available. In English, attention ranking mirrors the grammatical role in which noun phrases are realized; the noun phrase in subject position takes precedence over the one in object position.³⁰ This is why, with (11), ‘He’ is resolved to the candidate referent introduced by the previous subject, ‘A man’, rather than one introduced by the previous object ‘Sam’. Crucially, this preference for referents introduced by noun phrases in subject position is a grammatical feature of English, and not universally shared across languages.³¹ But what about ‘him’? The dependent interpretation linked to ‘A man’ remains most prominent when we turn to interpreting ‘him’; the ranking remains constant throughout the second sentence of (11). But because ‘him’ cannot be so resolved, on pain of violating Principle B, it must resolve to the highest-ranked candidate lower in the ranking that satisfies the relevant constraints, i.e., Sam.

A streamlined way of representing this aspect of English formally is to link grammatical roles to specific positions on the stack of candidate interpretations: the subject corresponds to position 0, the

²⁸Assume (11) is uttered without accompanying demonstrative gestures. That the NP in subject position is preferred as an antecedent for the subsequent anaphor over the one in object position is well-documented by a number of corpus and psycholinguistic studies. See Kameyama (1996); Walker, Masayo, and Cote (1994); Bittner (2014).

²⁹Witness: “Bill came in. He sat down”.

³⁰This observation is well documented. See e.g. Kameyama (1996); Kehler (2002); Bittner (2014).

³¹Languages with a flexible syntax can exploit word order to indicate prominence. Some grammaticize prominence with morphemes like topic markers that crosscut word order and grammatical role. See Kameyama (1996); Walker, Masayo, and Cote (1994); Bittner (2014).

direct object to 1, the indirect object to 2. We use expressions of forms ‘ $\langle \alpha n \rangle$ ’ and ‘ $\langle \pi n c \rangle$ ’ to encode updates that push referents to position n in the stack (arbitrarily limiting ourselves to $n = 0, 1$ or 2 and assuming the value of n follows compositionally from the grammatical status of the expression being represented), pushing all other one position down. On this strategy, the first sentences of (1) and (11) are formalized as (13) and (14) respectively.

$$(13) \quad \langle \alpha 0 \rangle; [man(x0)]; [walk.in(x0)]$$

$$(14) \quad \langle \alpha 0 \rangle; [man(x0)]; \langle \pi 1 s \rangle; [met(x0, x1)]; [greeted(x0, x1)]$$

Given these LFs, in formalizing successive sentences across discourse, we should encode the pronoun as assigned a grammatical role in its own sentence, and represent this by an explicit further update to the attentional state of the discourse. (To represent the effect of Principle B in resolving the pronoun, we introduce ‘ $@he^{x0}$ ’, an expression that denotes the most prominent referent whose value satisfies the pronoun other than $x0$ —meaning the top-ranked male other than the subject.) We get (15) and (16):³²

$$(15) \quad \langle \alpha 0 \rangle; [man(x0)]; [walk.in(x0)]; \langle \pi 0 @he \rangle; [sit.down(x0)]$$

$$(16) \quad \langle \alpha 0 \rangle; [man(x0)]; \langle \pi 1 s \rangle; [met(x0, x1)]; \\ \langle \pi 0 @he \rangle; \langle \pi c 1 @he^{x0} \rangle; [greeted(x0, x1)]$$

This approach achieves the most uniform possible syntax–semantics interface. In particular, our formal representation consists entirely of updates readable directly off the lexical items that comprise the utterance, along with their grammatical roles. What about the second challenge created by (12)? While (12-a) is explained in parallel fashion to (11), by appealing to subject preference for ‘he’ taking John, and Principle B for ‘him’ taking Tim, (12-b) seems to run counter, for, though there is a prominent, accessible antecedent in subject position (‘John’) satisfying the character of the pronoun, the preferred resolution of ‘he’ is to the referent introduced by the antecedent in object

³²Since the pronoun is in subject position, the update—‘ $\langle \pi 0 @ \rangle$ ’—that corresponds to the pronoun promotes the candidate referent associated with ‘he’ as the top-ranked referent. It does not introduce a new one; it rather re–stores (i.e., pushes a copy of) an old one to the designated position on the list of prominent candidates. If the entity is already at the top-ranked position, as in (15)–(16), the update associated with the pronoun will not change possible interpretations of subsequent ones.

position ('Tim'). This is where coherence enters.

2.4 Coherence

Coherence theory starts from an obvious observation: a discourse is more than a sequence of grammatical sentences. Successive contributions must be linked together by a recognizable flow of interpretive relationships. We see the requirement of coherence in (17) and (18), from (Hobbs, 1979).

(17) John took the train from Paris to Istanbul. He has family there.

(18) John took the train from Paris to Istanbul. He likes spinach.

(17) doesn't just reveal two facts about John: it suggests the reason he went to Istanbul was to visit his family. A coherence relation of Explanation links the second sentence to the first.³³ Coherence theory recommends we represent this connection explicitly to capture the correct interpretation of (17). Conversely, the requirement that a discourse must be coherent is strikingly evident in the interpretive effort (18) elicits. Given apparently unrelated facts about John in (18), we search for a connection. Is Istanbul known for its spinach? Is the train? Clearly, interlocutors must use the common ground to disambiguate between discourses harboring different coherence relations, just as they must use the common ground to resolve other ambiguities. But just as clearly, a failure to acknowledge any of these ingredients of interpretation constitutes a failure to understand the discourse.³⁴

Kehler advances the view that coherence relations cluster into (at least) three qualitatively different sorts, reflecting alternative strategies for organizing discourse,³⁵ as illustrated in (19)–(21).

³³See Hobbs (1979); Asher and Lascarides (2003); Kehler (2002).

³⁴Representing coherence relations explicitly does not mean an explanatory relation is a part of the truth-conditional content of (17). Though we will argue that connections signaled by coherence relations, like Explanation in (17), are not conveyed as a conversational implicature, this is not at odds with this content being not at-issue, in the sense of Tonhauser et al. (2013); conventional encoding of non-asserted content is not unusual (semantic presuppositions, conventional implicatures, expressive content). For simplicity, our formalism includes only one dimension of semantic content, in the sense of Potts (2005), but we could easily modify our framework to capture conventionalized not at-issue content. However, this task is for another paper.

³⁵Kehler's typology is useful in providing an intuitive picture of the kinds of interpretive connections we find in discourse. Other typologies might be better suited for other purposes. Cf. Mann and Thompson (1988); Knott (1996).

- (19) Max spilt a bucket of water. He tripped on his shoelace.
- (20) Max spilt a bucket of water. He spilt it all over the rug.
- (21) Max spilt a bucket of water. John dropped a jar of cookies.³⁶

(19) illustrates the kind of an explanatory discourse we considered in (17). For Kehler, Explanation is an instance of a broader class of cause–effect (or event–result) relations that speakers can utilize to organize discourse. (20), meanwhile, gives an extended description of unfolding events—thus, a Narrative connection, which for Kehler epitomizes a broader class of Contiguity relations. (21) exemplifies what Kehler calls Resemblance relations, organizing a discourse to draw comparisons and contrasts. In (21), there’s a Parallel between Max’s and John’s respective accidents. Different coherence relations are alike in signaling relationships among propositions in discourse. However, as we shall see, these relations also shape how other material in a discourse is interpreted. This is particularly important for context–dependent elements. The best way to capture these interpretive effects formally, we argue, is to represent coherence relations explicitly in logical form.

Coherence theorists view *identifying coherence relations* and *resolving semantic ambiguities* as mutually constraining. In (12-b), only resolving ‘he’ to Tim allows for a plausible explanation of John’s disappointment. In (17), only resolving ‘there’ to Istanbul allows for a plausible explanation of John’s trip.³⁷ In (19)–(21), meanwhile, we infer a temporal relation between the spill described initially and the tripping, spilling or dropping described next that matches the inferred coherence relation. Reference and coherence relations fit together in such cases.

A number of studies have experimentally confirmed the interdependence of resolving pronouns and establishing coherence.³⁸ An illustrative example is (22), from Smyth (1994):

³⁶See Kehler (2002). (19) and (20) are adapted by Kehler from Asher and Lascarides (2003).

³⁷Note that changing (18) to ‘John took a train from Paris to Istanbul. He likes spinach there,’ still only allows resolving ‘there’ to Istanbul for a plausible explanation of John’s trip. And since Explanation is naturally signaled in such examples, as in (17) and (18), the pronoun is understood to refer to Istanbul. This further supports our claim that the choice of relation governs the resolution of the pronoun. (Of course, so modified the discourse is no longer as incoherent as the one sequencing random information in (18), because it is clear how the second sentence provides an explanation of the first, as is required by Explanation. But that is beside our point.)

³⁸See, for instance, Keiser (2009), and references therein.

(22) Phil tickled Stanley, and Liz poked him.³⁹

Speakers tend to interpret (22) in either of two ways. One way assumes Liz's action was *prompted* by Phil's. Liz is perhaps reacting with disapproval to what Phil has done. This cause–effect interpretation comes with the understanding that ‘him’ must refer to Phil. The second reading assumes Liz's action was *similar* to Phil's in certain respects. This parallel interpretation comes with the understanding that ‘him’ must refer to Stanley. Crucially, the choice of a coherence relation (Result or Parallel) and the pronoun resolution (to Phil or Stanley) go in tandem.⁴⁰ Pragmatic theories of reference resolution—even on standard coherence approaches—take this as evidence of an inferential relationship between a speaker's intention in organizing the discourse and her referential intentions. According to us, these studies confirm a tighter connection between coherence relations and pronoun resolution than Coherence theorists have been inclined to posit.

The contrast between these interpretations leads us to conclude it's a mistake to treat (22) as harboring separate ambiguities audiences must resolve in turn—one about discourse coherence, another about pronoun resolution. The examples suggest that, once a coherence relation is established, a certain pronoun resolution is automatically set up. We hypothesize that pronoun resolution is *settled* by whichever coherence relations organize a discourse.

We argue coherence relations are another sort of mechanism that effects changes in the attentional state of a discourse. More precisely, we formally represent coherence relations in the logical form of a discourse, but propose further that this requires representing not only inferential connections, but also shifts in attention associated with coherence relations. Put simply, coherence

³⁹We assume no demonstrative gestures are accompanying (22).

⁴⁰When (22) is embedded within a larger discourse, we might be able to get other resolutions, but this is because embeddings can license different coherence relations. Similarly, prosodic marking can affect pronoun resolution. This is expected, since prosodic marking can affect prominence just as demonstrative gestures can. Likewise, filling in a context might change the interpretive dependencies that affect pronoun resolution by virtue of changing the coherence relations organizing a discourse, or by giving rise to other linguistic constraints on interpretation. One might imagine a scenario in which it is a part of the common ground that Liz always copies Phil, and tries to mimic whatever he does. One could then say, “Phil tickled Stanley, and as a result, she poked him”, where ‘him’ is interpreted as Phil. Two things to note: first, we should not assume coherence relations are mutually exclusive. In many cases, more than one is needed to capture the structure of a discourse (see Asher and Lascarides (2003)). This might give rise to a more complex pattern of promoting referents to prominence. And, second, the presence of an explicit descriptive signal such as “as a result” need not automatically mean a particular relation—Result—is organizing the discourse (see Webber et al. (2003)). Coherence relations are a matter of linguistic dependencies and not just any way of describing how events are related gives rise to a particular coherence relation. Similarly, not just any way of describing cause-effect relations counts as Explanation. “as a result” might be targeting a background presupposition about Liz's known pattern of behavior. There is reason to think presuppositions place further constraints on possible antecedents (see Hobbs (1979)).

establishes an interpretive relation, and moreover, (as we shall argue) as a matter of *linguistic* contribution, they promote certain entities to prominence. We depart from standard Coherence theorists inasmuch as we maintain that coherence relations come with grammatically encoded shifts in attention, much like NPs or pointing gestures do. It is only after these shifts are acknowledged as a part of a linguistic contribution that we can represent the intuitively correct interpretations of pronouns, while giving pronouns the uniform and unambiguous meanings we have already proposed. We first explain how this account captures data from (12-a) and (12-b), and then, present several reasons in defense of the view that the attention–shifts associated with coherence relations are grammatically encoded, not pragmatically implied. According to us, the difference between resolutions in (12-a) and (12-b) is fixed by the coherence relation that figures in their respective representations; these come with different attention shifting updates, which affects, semantically, the resolution of subsequent pronouns. In (12-a), the coherence relation is Narration: the content of the second sentence *follows up* on the content of the first, providing an extended description of unfolding events. In (12-b), the coherence relation is Explanation: the content of the second sentence is taken to explain the content of the first. According to Coherence theory, one or the other of these relations surfaces in formal representations of (12-a) and (12-b). Our proposal is that, *as a matter of language*, these distinct coherence relations affect the attentional state of the discourse, promoting one or the other of the candidate referents to the top–ranked position. When a formal representation features Narration, the attention–shifting operation raises John (the subject) to prominence; and when it features Explanation, the attention–shifting operation raises Tim (the direct object) to prominence.⁴¹ Accordingly, we represent (12-a) as (23), and (12-b) as (23) and (23):

$$\begin{aligned}
 (23) \quad & \langle \pi_0 j \rangle; \langle \pi_1 t \rangle; [was.disappointed.with(x_0, x_1)]; \\
 & [Narration(x_0)]; \langle \pi_0 x_0 \rangle; \\
 & \langle \pi_0 @ he \rangle; \langle \pi_1 @ hex_0 \rangle; [fired(x_0, x_1)]
 \end{aligned}$$

$$\begin{aligned}
 (24) \quad & \langle \pi_0 j \rangle; \langle \pi_1 t \rangle; [was.disappointed.with(x_0, x_1)]; \\
 & [Explanation(x_0, x_1)]; \langle \pi_0 x_1 \rangle;
 \end{aligned}$$

⁴¹ As we shall see below, not all explanatory coherence relations raise the direct object to prominence. Coherence relations encode linguistic interdependencies in discourse, not mere common–sense dependencies between events in the world. Thus, which entity is raised to prominence by which coherence relation is a matter of empirical inquiry, not *a priori* judgment, or reasoning about common–sense relations between events in the world.

$$\langle \pi 0 @ he \rangle; \langle \alpha 1 \rangle; [work(x1)]; [sloppy(x1)]; [did(x0, x1)]$$

(23) introduces John into subject position, and Tim into direct object position, requiring that John was disappointed by Tim. The second sentence continues a narrative about John; we represent the contribution of Narration as: $\langle [Narration(x0)]; \langle \pi 0 x0 \rangle \rangle$, an update that requires that the discourse continues the narrative about $x0$, and correspondingly, renders $x0$ prominent in the attentional state of the discourse.⁴² ‘He’ picks out the most prominent candidate appropriate for it, which is John, while ‘him’ picks out the currently most prominent candidate other than John, which is Tim. Further, it is required the subject (John) fired the object (Tim).

The first line in (24) is the same as in (23). The crucial difference comes next. This discourse is modeled as harboring Explanation, which we formalize using the update $\langle [Explanation(x0, x1)]; \langle \pi 0 x1 \rangle \rangle$. This update requires that the two bits of discourse stand in an explanatory relation, and promotes the entity in object position to prominence. ‘he’ continues to pick out the top-ranked interpretation, but due to the update associated with Explanation, this is Tim. The formula then proceeds with an additional dynamic existential quantifier, and further conditions that ensure its witness is sloppy work done by the subject Tim.

Crucially, differences in pronoun resolution *follow from* specifications of the coherence relations that organize the discourse and update the attentional state, and *not* from an open-ended process of pronoun resolution. It might seem surprising to describe attention-shifts associated with coherence as grammatically encoded. By contrast, you might hold that attention shifts contributed by coherence relations only reflect speaker intentions and the hearer’s common-sense inference. This view is, in fact, standard among Coherence theorists. So construed, an attention shift that guarantees the correct interpretation of the subsequent pronoun is not a linguistic effect of coherence, as we urge, but rather of pragmatic reasoning that occurs once a hearer has established that a particular coherence relation is structuring the discourse. A hearer reasons that in (12-a), since Narration holds between the two sentences, the speaker must be intending to promote a certain referent, in this case,

⁴²This representation of Narration suffices to capture the effects on attention that concern us. For other purposes, we might want to refine the representation of Narration to account for spatiotemporal and causal links we find in narrative discourse. Hobbs (1979, 1990) formalizes these inferential connections in terms of a relationship between eventualities described in successive sentences. Asher and Lascarides (2003) model it as a relationship between dynamic propositions expressed by successive sentences.

John, to the center of attention. On this view, intention recognition affects the re-ranking of the list of prominent referents.

It certainly *makes sense* for attention to shift in the ways (23) and (24) suggest. It would be perverse to reverse these preferences, so that when we came to consider a narrative about John who is disappointed in Tim, the referent of ‘he’ was Tim, or when we came to explain what about Tim made John disappointed in him, suddenly the referent for ‘he’ was John. To organize discourse in such a confusing manner would make it much harder to communicate our ideas concisely. However, ask not *where* coherence relations shift attention, but rather *in what circumstances* they do so. This question gives us reason for maintaining that attention-shifting operations are grammatically encoded. To see why, first observe that speakers and hearers take a restricted set of cues into account in instances of pronoun resolution. They privilege linguistic ones, over the broader constraints of background knowledge and rational inference that they might potentially consider. For a perfect example of this regularity, consider (25), from Kehler (2002):

(25) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

Kehler reports (25) is generally judged to be infelicitous by subjects. What explains this? By virtue of following ‘admires’ with ‘absolutely worships’—a stronger term in an obvious scalar relationship—the speaker provides clear evidence (25) is organized around a contrast between Margaret Thatcher’s and George W. Bush’s comparable attitudes. Coherence theory predicts this parallel should make Reagan (the object of the first clause) the preferred referent for the pronoun in object position in the second clause.⁴³ And, indeed, reading (25), it seems as if the speaker has erred, inadvertently referring to Reagan as ‘her’. Of course, since Thatcher has been evoked in the previous sentence, in subject position, and is a well-known object of Bush’s admiration, you’d expect it would be rather easy to refer to Thatcher with ‘her’ were the effect of coherence on prominence merely a by-product of the general pragmatic, common-sense reasoning that interpreters use to recognize a plausible interpretation; yet, this is not what we find.

⁴³Kehler is not committed to this preference being a result of a grammaticized contribution of coherence. He uses (25) for different purposes, but it perfectly illustrates our point.

This point is analogous to the conventionality of demonstration. Thatcher may attract our visual attention with what she's doing, but if the speaker is pointing elsewhere while saying 'her'—or continuing an ongoing discourse about someone else—we do not take 'her' to refer to Thatcher. Likewise, the Parallel relation encoded in (25) accomplishes a kind of inferred demonstration, indicating Reagan in a way that's difficult for a common-sense inference to override.⁴⁴

Further support for our thesis comes from the variation we find across languages. Many have explicit operations for shifting attention, such as grammatical topic marking, or a distinction between topic and non-topic pronouns. Some are more constrained than English in the sorts of shifts they permit to take place implicitly. A rough inspection of Serbian, a language that allows (sometimes requires) pronouns to be “dropped”, i.e. remain unpronounced in certain grammatical positions, suggests it is a language of this kind. So, consider two possible translations of (12-a):

(26) Džon je bio razočaran Timom.

John–NOM is–PRS–3ms be–PPA–3ms disappointed–ADJ–3ms Tim–INS.

a. Otpustio ga je.

Fired–PPA–3ms him is–PRS–3ms.

b. On ga je otpustio.

He him is–PRS–3ms fired–PPA–3ms.

When (12-a) is translated into Serbian with its third-person male singular pronoun “dropped” (as in (26-a)), the interpretation is that John fired Tim. Yet when the pronoun is overt as in (26-b), the interpretation is Tim fired John. It seems that the overt pronoun signals a change in prominence the covert counterpart does not. English, lacking this explicit means of signaling a shift in attention, is more flexible with implicit shifts in attention; thus, we witness an ambiguity in (12-a)–(12-b).

The Greenlandic language Kalaallisut provides another clear case. Bittner (2014), contrasts (27-a) in English with (27-b) in Kalaallisut.

⁴⁴Focal stress on 'her' makes (25) felicitous, with the referent of 'her' being Thatcher (as Kehler reports). But this stress carries interpretive requirements, changing the point of the discourse: it explains how Bush follows Thatcher's opinions, not how conservative politicians feel about Reagan. There is no reading where the relation is Parallel and the pronoun resolves to Thatcher. But this should be achievable if the attention-shifting effect of Parallel were a pragmatic, default one.

(27) *Traveling to Denmark*

- a. When I came to Denmark, I bought my ticket six months in advance.
- b. Danmarkimut tikikkama,
 Danmarki-mut tikit-ga-ma,
 Denmark-sg.DAT come-FCT-1s
 When I came to Denmark,
 qaammatit arvinillit siuqqullugu billitsisivunga.
 qaammat-t arvinilli-t siuqqut-llu-gu billitsi-si-pu-nga.
 month-pl six-pl v.ahead-ELA-3s ticket-get-IND.IV-1s
 I got a ticket (for some other event) six months ahead (of that event).

Following Moens, and Steedman and Webber, we take sentences with subordinate clauses, such as ‘when’ clauses, to be mini-discourses, with coherent interpretive connections between the clauses that mirror interpretive connections found between successive sentences in a discourse, or between conjuncts.⁴⁵ For (27-a), English speakers find a natural interpretation where buying the ticket early is a description that elaborates how the speaker came to Denmark. For its translation (27-b), however, the analogous interpretation is unavailable to Kalaallisut speakers. The Kalaallisut sentence requires an interpretation where the main clause describes what happens after the speaker came to Denmark—which is compatible with the narrative interpretation, but not elaboration. We witness this discrepancy because Kalaallisut marks the resolution of temporal anaphora, so that an event verb forces the temporal progression, making the time “right after the event described” prominent for temporal anaphora. Since such a temporal progression, dictated by an eventive verb, is consistent with the narrative interpretation, we find this interpretation available. However, though this language has grammatical means for signaling shifts in attention for temporal anaphora, it, unlike English, seems not to associate coherence relations (viz., Elaboration) with attention shifting updates. In English, Elaboration contributes an attention shifting update that makes the time prior to coming to Denmark prominent for temporal anaphora. In Kalaallisut, Elaboration lacks this attention shifting update, and so, absent other grammatically marked ways of rendering the time before the trip to

⁴⁵See Moens and Steedman (1988); Webber (1988); Webber et al. (2003).

Denmark prominent, the only way to resolve temporal anaphora in (27-b) is for the trip to Denmark to precede buying the ticket. This interpretation is consistent with Narrative, not Elaboration, and so, there's no Elaboration interpretation of (27-b). The discrepancy between English and Kalaallisut would be puzzling were the shifts induced by Elaboration in English a result of pragmatic reasoning. If it were so, we would expect to find the same range of interpretations in Kalaallisut. However, we do not. Only if languages interpret analogous expressions differently as signals of transitions in discourse can we accommodate the differences we find. So, we infer that each coherence relation in logical form carries a rule-governed contribution to attention in discourse. Given the state of attention, pronoun resolution follows.

Our view is consistent with general reasoning being crucial to interpretation. When the grammar delivers multiple candidate readings, hearers need to choose one that makes the most sense on a given occasion. (22) is ambiguous between a discourse containing Result and one featuring Parallel. Some general reasoning might be invoked in disambiguating between these, much as it might be involved in figuring out whether a speaker means a financial institution or a river bank, with a use of 'bank', or which quantifier scope is intended with a use of 'Every boy kissed a girl'. To interpret, a hearer must first settle disambiguations. This may involve assessing the plausibility of inferential links conveyed by coherence relations. It may involve evaluating whether a reading engenders a plausible resolution of pronouns. Such reasoning constitutes an important principle of disambiguation, but it cannot contribute content to logical form. It serves to privilege a logical form of discourse, among available ones grammar delivers. But once a coherence relation is established, pronoun resolution is determined by grammar, not general reasoning. In sum, strong evidence favors treating attention-shifting effects of coherence relations as governed by linguistic rules, not as a byproduct of pragmatic reasoning. These attention-shifting updates change the attentional state of a discourse, thus, setting the parameters of the context that determine the resolution of a pronoun. A pronoun, in turn, as a function of these parameters, automatically selects a referent, according to linguistic meaning. Thus, not only is the resolution of a pronoun determined by its linguistic meaning as a function of context, but moreover, relevant features of context that fix this reference are themselves orchestrated by rules of language. So, pronominal resolution is linguistically determined—through and through. Linguistic mechanisms of attention-shifting updates and rules governing pronominal usage *together determine* pronominal resolution on any occasion of use.

2.5 Conclusion

We defended a joint attention–coherence account for pronoun resolution that assigns one linguistic meaning to each pronoun. Pronouns are variables with dependent interpretations, but are interpreted relative to a prominence ranking. Each resolution is restricted by additional constraints, including person, gender and number features, and independently motivated syntactic constraints. They trigger a search for a matching interpretation. Think of the character of a pronoun as incorporating these constraints; on our account, it determines the resolution of a pronoun automatically as function of context. To make this work, context must be appropriately set up. This is achieved by the attentional state of a discourse, the result of a series of attention–shifting updates, which intuitively re–rank candidate resolutions for pronouns; formally, they are pushing new entities onto a stack of values for variables, demoting others. These updates are contributed by the diverse mechanisms for structuring discourse and shaping pronominal interpretation, including evoking discourse entities in specific grammatical roles, demonstrating entities with non–verbal (yet linguistic) actions, and signaling the direction of discourse through various interpretive connections between clauses. Although these mechanisms are heterogeneous, their contribution is governed by rules, not pragmatic reasoning, and as such, should be formally represented in logical form.

Our account is provisional in several respects. We said little about the modal profile of context–sensitive utterances.⁴⁶ We remained silent about many ambiguities associated with discourses containing pronouns—both within English and across languages.⁴⁷ And we have not provided a (exhaustive) list of coherence relations. To see what’s at stake, consider (28) from Winograd (1972):

- (28) The city council denied the demonstrators a permit.
- a. They feared violence.
 - b. They advocated violence.

Both (28-a) and (28-b) exhibit an explanatory relation, yet the occurrences of ‘they’ are resolved differently. Our suggestion is that, though in (28-a) and (28-b) the content of the second sentence

⁴⁶Still, our formal semantics has resources to capture basic facts about these modal profiles.

⁴⁷For example, we have not offered a treatment of dependent clauses.

is taken to explain the content of the first, these explanatory relations are qualitatively different. In (28-b), the council's decision about the demonstrators can be explained on the basis of the former's beliefs about the latter: the relevant explanation being that because the demonstrators are potentially violent, or believed to be so by the city council, the council denied them a permit. Meanwhile, in (28-a), the council's decision about the demonstrators can be explained based on (other aspects of) the council's attitudes: it is because the council feared violence that they decided to deny the demonstrators a permit. These qualitative differences suggest distinct coherence relations are at play. So, it's no surprise they come with different attention-shifting operations—when the explanation goes via subject, the subject is promoted to prominence; when it goes via object, the pattern is reversed. To flesh out this suggestion, we need to substantiate systematic differences between these kinds of explanations. Only further empirical research can guide us closer to having a firm grasp on all of these issues, and we submit our proposal here to the scrutiny of this future research.

We remain excited about the philosophical ramifications of the tools developed here. Philosophers often use linguistic examples to argue for context dependence. The kind of context dependence they have in mind, implicitly, is where they see context-dependent elements as freely selecting one interpretation from an open-ended array of candidates by unspecified, broadly pragmatic and open-ended mechanisms. Often, both the linguistic analyses and ensuing philosophical arguments depend on this model of context dependence. But if we are right, context is not as powerful as philosophers have presumed. In interpreting even a straightforward case like (1), our discourse ultimately lacks a dependency on *non-linguistic* context. This suggests a project of extending our account to other cases of apparent contextual variability. If we can succeed in capturing the apparent contextual variability with uniform meanings and constrained variation, philosophers will have to give up many customary appeals to context sensitivity.

Chapter 3

One's *Modus Ponens*: Modality, Coherence and Logic¹

3.1 Introduction

Recently, there has been a shift away from traditional truth conditional accounts of meaning towards non-truth-conditional ones, e.g., expressivism, relativism and certain forms of dynamic semantics. Fueling this trend are some puzzling behavior of modal discourse. One particularly surprising manifestation of such behavior is the alleged failure of some of the most entrenched classical rules of inference; viz., *modus ponens* (MP) and *modus tollens* (MT). Thus, several authors have independently touted counterexamples to MP and MT.² Because each challenge arises in the presence of modal language (typically involving embedded modals or conditionals), these critics believe they have uncovered a tension between the behavior of modal vocabulary and classical logic, with the moral being to revise the semantics for modality in order to invalidate certain classical patterns. These revisionary, non-truth-conditional accounts tout the alleged tension between the behavior of modal vocabulary in natural language and classical logic, as data in support of their departure from tradition, since the revisionary semantics invalidate some of these patterns.³ I, instead, offer a

¹The chapter is structured in two parts, followed by a formal appendix. In Part I (§ 3.2), I develop my account of modality. Part II (§ 3.3) describes the formal tools that allow us to build the semantics described in Part I. A formal appendix, (§ B.1) provides the complete formalization, along with a proof that the semantics preserves classical logic. Part II and the appendix should thus be seen as complementing Part I, by providing a formal implementation of the account developed in Part I. A version of this chapter is forthcoming in *Philosophy and Phenomenological Research*.

²See in particular McGee (1985) and Willer (2014) for counterexamples to MP, and Yalcin (2012b) and Veltman (1985) for counterexamples to MT. See also Kolodny and MacFarlane (2010) who reject MP as a reaction to certain puzzling behavior of deontic modals, and Cantwell (2008) who argues that MP, MT and reasoning by cases all fail when they involve modals in the consequents of conditionals. The revisionary accounts are typically non-truth-conditional, insofar as they deny that an utterance expresses a proposition that is true or false depending on the way the world is. In other words, these accounts deny utterances express propositional content. An exception to this is McGee (1985), who rejects MP, and thus endorses a non-classical account, but one that is nevertheless not non-truth-conditional. In this respect, it differs significantly from these other accounts, which invalidate the relevant patterns of inference in part by rejecting truth-conditionality of modal vocabulary.

³This is the course explicitly taken by Yalcin (2012b), who appeals to his counterexample as a basis for a rejection of MT. Some relativists also reject MP and/or MT based on the puzzling behavior of modal language (Kolodny and

semantics for modality with the resources to accommodate the puzzling data while preserving classical logic, thus affirming the tradition that modals express ordinary truth-conditional content. My account shows that the real lesson of the apparent counterexamples is not the one the critics draw, but rather one they missed: namely, that there are linguistic mechanisms, reflected in the logical form of a discourse, that affect the interpretation of modal language in a context in a systematic and precise way, which have to be captured by any adequate semantic account of the interaction between discourse context and modal vocabulary. I develop and defend a theory that captures this interaction: according to my account, context affects the interpretation of modals through mechanisms that are independently motivated, and are required to explain the effects of context on the interpretation of context-sensitive expressions quite generally, even in the most basic case—the case of pronouns. In particular, I shall argue, the relevant mechanisms are the ones that govern how individual utterances are organized to form a coherent discourse. While the impacts of these mechanisms explain (and predict) the appearance of counterexamples, the underlying logic, as I shall prove, is classical.⁴

In short: I defend a theory that captures how context affects the interpretation of modals, and does so in a way that reconciles classical logic with the semantics of modal language. My strategy will be to focus on one particular case, due to Yalcin, but the account I develop naturally accommodates other known cases as well.⁵ The paper divides into two parts. In part § 3.2, I argue that a systematic impact of discourse context on the interpretation of modals explains the apparent counterexample, through general mechanisms, while preserving the validity of MT, and I sketch an account of modality and context-change that captures this interplay between discourse context, and modal expressions. In part § 3.3, I demonstrate that we can formalize this systematic impact, in a way that explains the counterexample, and preserves classical logic.

MacFarlane, 2010). And some dynamic semanticists, too, invalidate MT, where the diagnosis for rejection lies once again in the behavior of modal language (Gillies, 2010, 2004). I should note that, though the example I will focus on involves epistemic modality, the problem arises for other flavors of modality, as well; in particular, deontic modals, too, exhibit the behavior that *prima facie* violates classical patterns of inference. (See e.g. Kolodny and MacFarlane (2010).) The theory I will develop is not limited to epistemic modality, and my treatment of the apparent failure of classical patterns of inference naturally extends to other examples involving other types of modality, as well.

⁴To be precise, the underlying logic is a classical modal logic—a simple extension of classical logic with a modal necessity operator (in particular, system S_4 .) This logic preserves a classical inference system, and is sound and complete. Importantly, unlike another famous attempt at a defense of classical logic, pioneered by H. P. Grice (1989), I shall not be arguing that the truth-conditions of the English indicative conditional are those of the material conditional, but rather those of the strict conditional.

⁵E.g., the apparent counterexamples due to McGee (1985) and Kolodny and MacFarlane (2010).

3.2 Part I

3.2.1 The Challenge to *Modus Tollens*

The pattern $[If\ p, q; \neg q \therefore \neg p]$ is known as *modus tollens*. Yalcin argues for its invalidity as follows:

Take an urn with a 100 marbles. 10 of them are big and blue, 30 big and red, 50 small and blue, and 10 are small and red. One marble is randomly selected and hidden (you do not know which). Given this setup, (1) and (2) are licensed:

- (1) If the marble is big, then it is likely red.
- (2) The marble is not likely red.

But, surely, (3) does not follow:

- (3) So, the marble is not big.

Since the inference from (1) and (2) to (3) looks to be an application of MT, we would seem to have a counterexample.⁶

Possible reactions to the apparent counterexample are:

- i embrace the invalidity of MT, or
- ii deny that (1)–(3) is a genuine instance of MT.

(i) requires modifying our semantics for modals in such a way so as to invalidate MT. This is Yalcin's option, one that embraces a "revisionary" semantics for modal language that denies that modal and conditional sentences express ordinary truth-conditional content.⁷ It is important to note

⁶It is easy to construct other counterexamples along similar lines. See Veltman (1985) for an earlier counterexample to MT with right-nested conditionals. I shall focus on Yalcin, but my considerations extend to Veltman.

⁷See e.g. Yalcin (2012b), Yalcin (2007) and Moss (2015) for an expressivist version of revisionary semantics, see e.g. Kolodny and MacFarlane (2010) for a relativist version, and see Gillies (2010, 2004) for dynamic semantics version. These semantics are 'revisionary' precisely insofar as they invalidate classical patterns of inference, and deny that modal discourse expresses standard truth-conditions. It is worth noting that, although Yalcin revises the standard compositional semantics for modals and conditionals, he defends expressivism as a *pragmatic* thesis. In particular, he does not hold that the semantic content of a modal utterance is the informational content the utterance expresses; in fact, he denies that modal utterances express any informational content. The semantics I will develop and defend will require a substantial departure

that these revisionary frameworks do not deny that, intuitively, the big premise in (1)–(3) is (in some sense) about conditional probability, i.e., that it concerns the probability of the marble being red, conditional on it being big, while the small premise is, intuitively, ‘unrestricted’ in this sense. Everyone in the debate concedes this. What the revisionists deny, however, is that this intuition can be adequately captured by saying that the *truth-conditional content* of the big premise *describes* the conditional probability, while the *truth-conditional content* of the small premise *describes* the unrestricted one. They all argue, in one way or another, that there is no plausible and systematic way of deriving the right truth-conditions given a context of utterance.⁸

Note that, though it might seem that the revisionary tendencies are localized to particular counterexamples to a particular inference pattern—MT, this reaction opens the floodgates: it’s easy to devise similar counterexamples to numerous other inference patterns that are classically valid. The question then becomes which of the deductive rules of inference we should reject. To illustrate, consider the same scenario as in Yalcin’s original counterexample: there’s an urn with 100 marbles, etc. But then:

- (4) Suppose that the marble is big.
- (5) Then it is likely red.
- (6) But the marble is not likely red.
- (7) So, the marble is not big.

(4)–(7) is also horrible. If (1)–(3) provides grounds for rejecting MT, then by parity (4)–(7) provides grounds for rejecting MP or *reductio*. We might then reject both *modus ponens* and *modus tollens* (as in fact Kolodny and MacFarlane (2010) do). Or we might reject *reductio*. Should we reject all of these rules? Of course, we can derive MT through MP and *reductio*, on the assumption of the monotonicity of logical consequence. More generally, we know that rules of inference are holistic in this sense. But that is part of the problem—how do we isolate the culprit(s)? How do we choose?

from the traditional semantics for modal discourse (as in Kratzer (1977), and Kratzer (1983)), but one that allows us to vindicate the idea that modal utterances express truth-conditional content. Thus, my semantics is conservative, insofar as it preserves classical patterns of inference, and the classical truth-conditions for modal vocabulary.

⁸Again, assuming the standard notion of truth-conditions.

None of this is, strictly speaking, an argument against option (i), since, the proponents of the revisionary views already accept that our modal vocabulary is at odds with classical logic.⁹ But it does show that there is a lot more at stake than the loss of MT. Moreover, as we shall see shortly, contrary to what the proponents of the revisionary views assume, the problem is *not* tied specifically to modal vocabulary, but can be replicated with the most basic case of context-sensitive expressions—demonstrative pronouns. This suggests that the real problem may lie elsewhere. A theory of context-sensitivity resolution that preserves the aforementioned inference rules, while explaining away the appearance of counterexamples has a lot going for it. Since I am precisely interested in developing such a theory, I shall reject option (i).

One way to pursue option (ii) that has received some attention in the literature is to claim that epistemic modals—in particular, the modal operator ‘likely’—takes obligatory wide-scope over the conditional, so that (1)–(3) is not a genuine instance of MT.¹⁰

If this were the only way to implement option (ii), we would face a hard choice indeed, for obligatory wide-scoping faces well-known problems. For one, it generates intuitively incorrect predictions. For example, on the (supposedly obligatory) wide-scope reading of the modal ‘likely’, we get the intuitively wrong interpretation for (8).¹¹

- (8) If Bill comes to the party, then John will come and it is likely that Margaret will come, too.

(8) does not have a reading according to which it is likely that if Bill comes to the party, John and Margaret will come, too.

We can, alternatively, defend option (ii) by appeal to the context-sensitivity of modal operators.

⁹A number of authors have suggested that the failure of MT is best explained by adopting the informational approach to logic and consequence relation. We can then provide a proof system for such logic, and study its relations to classical logic. See e.g. Yalcin (2007, 2012b); Bledin (2014, forthcoming). See also Veltman (1985); Gillies (2004); Willer (2014) for a related approach to logical consequence. These authors point out that various classically valid inference patterns—MT and *reductio* for instance—are valid within a restricted fragment that does not contain modal expressions. However, here, we are precisely interested in whether the presence of modal expressions gives rise to failures of classically valid patterns, and whether the best semantics for modality violates classical logic. In other words, the interesting question is whether MT or *reductio* are valid in a fragment containing a modal expression. I shall argue, contrary to what these authors maintain, that the answer to this question is positive.

¹⁰This proposal has been discussed by e.g. Yalcin (2012b), Dorr and Hawthorne (2014), and Kratzer (1983), though all of these authors reject it as unsatisfactory.

¹¹Yalcin (2012b) provides further discussion of this problem.

It is not particularly controversial that the interpretation of modal operators depends on context (although, exactly how is a matter of great controversy). A familiar view is that modals are quantifiers over possible worlds, but just which worlds depends on the context (Kratzer, 1977, 1981). We can exploit this to argue that the problematic counterexample can be explained away by maintaining that the modal ‘likely’ in (1) contributes a different semantic content than the one in (2), due to contextual effects on the interpretation of the two occurrences of the modal; and so, (2) and the consequent of (1) fail to contradict each other. Accordingly, (1)–(3) is not really an instance of MT.

This strategy captures the intuition that the consequent of (1) talks about a restricted (conditional) probability, while (2) talks about an unrestricted one.¹² The challenge is to explain exactly why and how the context secures different (and intuitively correct) interpretations for the two occurrences of the modal. To do so in a non-*ad hoc* way is notoriously difficult.¹³

To give more bite to the *ad hoc*-ness charge, note that it becomes even more pressing once we acknowledge that contextual effects are not freely available with many other uncontroversially context-sensitive expressions. To illustrate, consider the following example:

(9) If John ate the food from the fridge, then the fridge is empty.

(10) But the fridge is not empty.

(11) So, John didn’t eat the food from the fridge.

One cannot freely shift the context so that ‘empty’ in (10) means *in the state of a vacuum*, and thus avoid the MT reading. So, why assume that context can freely shift between (1) and (2), but not between (9) and (10)?¹⁴

This apparent asymmetry raises a worry about the relation between language and logic on

¹²However, as Yalcin (2012b) notes, one should not be too quick to conclude, given this intuition, that the consequent of (1) and (2) express different contents, since there are non-truth-conditional, non-contextualist accounts available, Yalcin’s included, that capture this intuition without positing a difference in content.

¹³See e.g. King (2014a), and Stanley (2000).

¹⁴Of course, one could try to modify (9)–(11) to achieve the context-shifting effect by, e.g. putting a focal stress on ‘empty’. But focal stress is one of the linguistic mechanisms that systematically affect the interpretation of context-sensitive expressions, in a way that is predicted by the theory I shall defend. (More on this below.) The point is just that one has to additionally signal when there is a contextual shift via a mechanism such as, e.g., focal stress. The need for some such mechanism precisely shows that context *cannot* shift freely. So, the challenge of spelling out the mechanism(s) that would yield the interpretation of (1)–(3) that the proponent of option (ii) defends becomes only more pressing.

the contextualist accounts. These accounts typically leave the resolution of context-sensitivity to broadly open-ended pragmatic mechanisms (e.g. speaker's intentions, or non-linguistic contextual cues).¹⁵ But, if logical forms are partly fixed by broadly unsystematic, open-ended mechanisms, then validity becomes partly a matter of such mechanisms, as well. As a result, on such accounts, there is no systematic, rule-governed way of determining whether (1)–(3), or (9)–(11), is expressing a valid argument, or not. For sure, given a set of fully-specified logical forms, one can determine a subset of valid ones; but given a surface form, like (1)–(3), one derives the fully-specified logical form only through open-ended, defeasible pragmatic processes. The rules of language alone do not determine whether a pattern like (1)–(3) is valid or not. Thus, logical inference becomes dependent on non-linguistic and psychological factors, like epistemic cues and communicative intentions. The link between grammar and logic is thus only indirect, mediated by pragmatic principles. This is worrisome if we cannot provide a systematic, yet non-*ad hoc*, account of when and how such mechanisms affect the resolution of context-sensitive expressions. And the problem is that the pragmatic mechanisms are too flexible to provide such systematic constraints: recall, if we want to claim that (9)–(11) is an instance of MT, while (1)–(3) is not, we need a principled story of why context affects the resolution in a particular way in one case, but not in the other. And we also need a story about why, by contrast with (9)–(11), we do not find contexts in which (1)–(3) expresses a valid logical form. It's hard to see what such a story would be, if it is to rely on speaker's intentions and general, non-linguistic contextual cues.

By contrast, I offer a systematic formal account of the effects of context on the interpretation of modals that is not *ad hoc*. I shall argue in what follows that we have *independent* evidence for context-change in (1)–(3). Moreover, I shall argue that we have evidence for the kind of context-change that is governed by mechanisms that affect the interpretation of context-sensitive vocabulary quite generally, and is resolved in a systematic, rule-governed way.¹⁶ Once we see the import of these mechanisms, we will see that the puzzling behavior of modals is not in tension with classical

¹⁵This kind of an approach was initially advocated by Kaplan (1989a), as an account of the resolution of demonstrative pronouns. The idea is roughly that, with the possible exception of the so-called *pure indexicals*—expressions like 'I', whose linguistic meaning fully determines the interpretation given a context—context-sensitive expressions require pragmatic supplementation in order to be interpreted in a context.

¹⁶As we shall see, moreover, these mechanisms do not merely affect the resolution of the context-sensitive items, but are also reflected in the logical form of an argument like (1)–(3). See also Chapter 4 for further discussion.

logic. In particular, we will be able to devise a precise theory of context that predicts the appearance of examples in (1)–(3), while at the same time, provably preserves classical logic.

3.2.2 Modals and Pronouns

Consider the following example:

- (12) If Jane is out, then she is having fun.
- (13) She (pointing at Mary) is not having fun.
- (14) So, Jane is not out.

(12)–(14) is obviously bad. But no one would invoke this argument to present a counterexample to MT. Obviously, (13) does not *really* contradict the consequent of (12). The occurrences of ‘she’ in (12) and in (13) resolve to different referents. This explanation works because it identifies non-*ad hoc* context-shifting in the inference. We can point to clear reasons why the reference is resolved in a particular way in the relevant cases. In (12), the pronoun ‘she’ occurs in an elaboration of a hypothetical scenario about Jane, and thus, is resolved to Jane. In (13), the pronoun occurs in tandem with a demonstrative gesture—a pointing at Mary—and thus is resolved to Mary. The *elaboration* in (12) and the pointing gesture in (13) render a certain referent the most prominent for the subsequent anaphora; the subsequent pronoun then, as a matter of its meaning—the character, in the sense of Kaplan (1989b)—picks out the most prominent referent.¹⁷ As I have argued elsewhere (Stojnić, Stone, and Lepore, 2013, 2014), and will amplify on in the next section, we can capture these observations by rendering explicit the mechanisms that systematically govern the resolution of

¹⁷Note that the example in (12)–(14) is not merely trading on the difference between deictic and anaphoric (uses of) pronouns. Similar examples can easily be constructed where all the pronouns are interpreted anaphorically:

- (i) Mary is upset because Jane is much luckier than she is.
 - a. If Jane buys a ticket, she always wins.
 - b. But, *she* does not always win.
 - c. #So, Jane didn’t buy a ticket.

No one would treat (i) as a serious threat to MT. We can point to reasons why a certain interpretation is naturally retrieved in the context: while ‘she’ in (i-a) is uttered in the course of an elaboration of a hypothetical scenario about Jane, and is thus resolved to Jane, the contrastive focus on the occurrence of ‘she’ in (i-b) requires that it refers to the contextually most prominent referent other than Jane—and this is Mary.

the pronoun. We can pinpoint the mechanisms that, for our purposes, play the key role in governing the resolution of context-sensitivity by building on the resources of so called Discourse Coherence Theories. As we shall see in what follows, these mechanisms are reflected in the logical form of the discourse like (12)—(14), which will be particularly important insofar as we want to treat validity as a matter of logical form.¹⁸

3.2.3 Coherence

The key insight behind Coherence Theory is the simple but often neglected observation that a discourse is more than a random sequence of sentences. To flesh this out, we begin with an illustrative example from Hobbs (1979):

(15) John took the train from Paris to Istanbul. He has family there.

(16) John took the train from Paris to Istanbul. He likes spinach.

There is a stark contrast between (15) and (16). While (15) is a perfectly felicitous piece of discourse, (16) (out of the blue) is odd. What explains this contrast? Note that (15) does not merely report two random, unrelated facts about John. It signals that John took the train from Paris to Istanbul *because* he has family there; we understand the second sentence as providing an explanation of the events described in the first. Recognizing this explanatory connection between the two bits of discourse in (15) is part of understanding the contribution in (15)—unless we recognize it, we have simply failed to fully understand the discourse. Unless we understand how the two bits of discourse are related, we cannot fully understand the speaker's contribution. By contrast, due to difficulties in establishing such a connection, (16) seems off. We are left wondering: is Istanbul famous for its spinach? Does spinach cause a fear of flying? That such an interpretive effort is in play in an attempt to understand (16) suggests the requirement of a readily recoverable implicit organization of the discourse that renders it coherent.

¹⁸See Hobbs (1979), Kehler (2002), and Asher and Lascarides (2003). I do not mean to suggest that mechanisms of discourse coherence are the only mechanisms affecting the resolution of either pronouns or modals. Other linguistic mechanisms, such as prosody, for example, likewise can affect the prominence in a discourse. See Stojnić, Stone, and Lepore (2013, 2014) for an account that incorporates the effects of various different linguistic mechanisms on the resolution of demonstrative pronouns.

Drawing on these observations, Coherence Theorists postulate an implicit organization of discourse that establishes inferential connections—*coherence relations*—among utterances (Hobbs, 1979; Kehler, 2002; Asher and Lascarides, 2003). This implicit organization arises from the communicative strategies that interlocutors exploit to convey and organize their ideas through an ongoing discourse. As demonstrated by the contrast between (15) and (16), successive contributions to a discourse must be linked by a recognizable network of interpretive relationships. The speaker must signal how she structures her contributions according to shared standards and conventions.¹⁹ So, for instance, (15) is understood as connected by the coherence relation of *Explanation*.²⁰ Failure to establish such a connection in (16) makes it seem off.²¹

Crucial for us is that the task of establishing discourse coherence and resolving semantic ambiguities turn out to be mutually correlated processes. In particular, as has been confirmed by a number of empirical studies, pronoun resolution co-varies with the choice of coherence relation.²² Here is an illustration:

(17) Phil tickled Stanley, and Liz poked him. (Smyth, 1994)

There are (at least) two ways we can understand (17). The second clause could be taken to describe the result of the event described by the first: Phil tickled Stanley, and *so*, Liz poked him (i.e. “Liz is avenging Stanley”); or, one could understand the two clauses as comparing and contrasting two parallel events: Phil tickled Stanley, and Liz poked him *as well* (i.e. “What happened to poor Stanley?”). Crucially, if the discourse is understood as connected by the Result relation, the pronoun refers to Phil; if it is organized around the Parallel relation, then ‘he’ is Stanley. The choice of a coherence relation guides the choice of pronoun resolution.

¹⁹This is not to say that a discourse cannot be ambiguous with respect to coherence relations it harbors. We see one such example of ambiguity in (17) below. Part of the interpretive effort in understanding a discourse is in resolving such ambiguities.

²⁰Coherence Theorists typically capture these observations by representing inferential relations—coherence relations—in the logical form of a discourse. Cf. Asher and Lascarides (2003). There is both good syntactic and good semantic evidence for representing coherence relations in the logical form of a discourse. See, in particular, Asher and Lascarides (2003) and Webber et al. (2003) for further discussion.

²¹Of course, we could make (16) felicitous, were we to provide a sufficiently rich context, which would allow us to establish the relevant relation—for example, if it were a part of the common ground that the best spinach is grown in Istanbul. This is just as expected.

²²See e.g. Wolf, Gibson, and Desmet (2004), Kehler et al. (2008), and Keiser (2009).

Moreover, the data suggest that the mutual constraints between these two tasks are both systematic and robust. That is, given the choice of a coherence relation, the interlocutors are radically constrained in the possible interpretation of a pronoun. To illustrate, consider the following example from Kehler (2002):

- (18) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

Kehler reports that (18) is generally judged infelicitous by his subjects. The subjects expect the pronoun in the second sentence to resolve to Reagan, and intuitively feel the speaker has erred in uttering ‘her’ instead of ‘him’. This is explained, again, by the interaction between the task of establishing coherence and that of resolving a pronoun. The discourse follows ‘admires’ by ‘absolutely worships’—a stronger term in a scalar relationship—thus signaling that the discourse is organized by the coherence relation of Parallel. Parallel requires that the occurrence of a pronoun in the object position be resolved to Reagan (the object of the first clause). Given the gender mismatch, the utterance is judged infelicitous. This is surprising given the available referent for ‘her’ in the first conjunct, one that is a well-known subject of Bush’s admiration. If the correlation were really a matter of mere general pragmatic defeasible reasoning, the perceived infelicity of (18) would be mysterious. The infelicity is, by contrast, expected if the effect of Parallel is a matter of an underlying convention—the convention determines that the referent has to be Reagan, and that is why we are stuck with infelicity, even in the presence of a nearby plausible interpretation.²³

I take these (and other²⁴) observations to show that coherence relations render certain referents

²³Though Coherence Theorists agree that, due to the semantic effects of coherence relations on the truth-conditions of a discourse, as well as due to syntactic constraints on discourse structure, coherence relations need to be represented in the logical form of a discourse (cf. Asher and Lascarides (2003); Webber et al. (2003)), they typically take the process of pronoun resolution and establishing coherence to be merely mutually constraining. They understand the preference for a particular coherence relation to be an input to general holistic process of reasoning that attempts to find the overall most plausible interpretation of the discourse. By contrast, I have argued elsewhere that we should understand these dependencies not merely as mutual constraints between two related tasks, but rather as linguistic effects of coherence relations on pronoun resolution; coherence relations as a matter of linguistic convention make particular entities preferred candidates for subsequent anaphora. (For detailed defense of this view, see Stojnić, Stone, and Lepore (2013, 2014).) I advance considerations below that the effects of discourse coherence on modals is likewise conventionalized. For a more detailed defense of this position see Chapters 1, 2 and 4.

²⁴It’s worth noting, moreover, that languages differ with respect to the effects of coherence relations on the interpretation of pronouns, which suggests that, indeed, the effect is a matter of linguistic convention. See Stojnić, Stone, and Lepore (2014).

prominent for subsequent anaphora resolution.²⁵ Here's one way to capture the constraints that the choice of coherence relation places on pronoun resolution. Suppose we rank candidate referents for anaphora in a discourse according to their relative prominence.²⁶ The import of discourse coherence is to affect this ranking, by making certain referents prominent for subsequent anaphora. The pronoun, then, according to its meaning, refers to the most prominent referent that satisfies the grammatical constraints encoded by the linguistic meaning of the pronoun (e.g. 'he' refers, roughly, to the top-ranked, third-person, singular, male candidate referent).²⁷ All of this supports our observations about (12)–(14). The consequent in (12) stands in an Elaboration relation with the antecedent; it is *because* 'she' in (12) occurs in an elaboration of the hypothetical scenario described by the antecedent that the pronoun refers to Jane. And it is *because* in (13) the pronoun occurs in a tandem with a pointing gesture, that it refers to the individual pointed at, namely, Mary. As a result, (12)–(14) is not an instance of MT.²⁸

So far, I have demonstrated that even in the case of pronouns, if we fail to appreciate how they are resolved within a discourse, we would be misled to interpret examples such as (12)–(14) as “counterexamples”. Moreover, I have argued that pronoun resolution is responsive to discourse structuring mechanisms, in particular, mechanisms of discourse coherence. Next, I shall argue that modals are analogous to pronouns in two crucial respects. First, like pronouns, their interpretation is an anaphoric process, by which I mean that they require a contextually available antecedent which

²⁵The notion of prominence within a discourse that I'm relying on here should not be confused with the intuitive notion of real-world salience. That a referent is salient is not sufficient (or necessary) to make it prominent in my sense. Consider an utterance of “Jim came in. He sat down,” while Bill is jumping up and down, making himself the most salient entity in the given situation. Unless the speaker points or somehow demonstrates that Bill is the referent, the occurrence of the pronoun 'he' will be interpreted as Jim. This is not to say that prominence has nothing to do with the entity being salient in the pre-theoretic sense; rather, it is just to say that interlocutors tend to rely on a relatively narrow set of intra-discourse cues that language provides for guiding their attention towards particular referents in a discourse. See Stojnić, Stone, and Lepore (2014).

²⁶The idea of a ranking of referential candidates in a discourse comes from Centering Theory, which hypothesizes that the referential candidates in a discourse are ranked according to relative prominence, those ranked higher being preferred over those ranked lower as interpretations for subsequent anaphora. See Sidner (1983), Grosz, Joshi, and Weinstein (1995), and Bittner (2014).

²⁷For a detailed development of this view, see Stojnić, Stone, and Lepore (2014).

²⁸Note that according to this explanation, there are changes in context—making certain referents prominent as antecedents for subsequent pronouns at particular points in discourse—which explain why the pronouns in (12)–(14) are resolved the way they are. But these contextual changes are induced by linguistic mechanisms, reflected in the logical form of a discourse; in particular, they are governed by mechanisms of discourse coherence.

is either linguistically introduced or available from a non-linguistic context.²⁹ And second, like the resolution of pronouns, the interpretation of modals is guided by the mechanisms that structure the information in discourse, in particular, mechanisms of discourse coherence. These mechanisms govern the interpretation of both types of expressions in a systematic, and rule-governed way. Once we incorporate the effects of these mechanisms on the resolution of context-sensitivity, we will see that we can devise a semantics for modals, the underlying logic of which remains classical.

3.2.4 Modals as Pronouns

That modals exhibit anaphoric-like behavior has been observed in the literature on modal subordination, as in (19), where a modal is interpreted relative to some other modal expression introduced earlier in a discourse (i.e. relative to a linguistically introduced antecedent):³⁰

- (19) A wolf might walk in. It would eat you first. (Roberts, 1989)

Perhaps the clearest argument for the full analogy between modals and pronouns with respect to the range of interpretive effects they permit has been offered by Stone (1997, 1999).³¹ First, he observes that both can depend for their interpretation on an antecedent introduced linguistically (either by means of indefinite or definite reference) earlier in the discourse, as illustrated by (19). Just as ‘he’ in (20) is naturally understood to refer to the man introduced in the first sentence,³² so ‘would’ in (19) is naturally understood restrictedly, as describing the hypothetical scenario introduced by the modal ‘might’ in the first clause.

- (20) John owns a donkey. He beats it. (Based on Partee (1984).)

Second, like pronouns, modals allow for reference to specific entities from a non-linguistic context.

²⁹Some theorists like to reserve the terms “anaphora” and “anaphoric” for cases where an item is bound by a linguistically introduced antecedent. The reader should bear in mind that I use these terms throughout in this broader sense specified in the text.

³⁰See Roberts (1989). How exactly to account for modal subordination has been a matter of much debate, one that has generated a vast literature. The account defended in this paper shows how to model modal subordination.

³¹Stone’s arguments for the parallel between modals and pronouns are analogous to Partee’s famous arguments for the parallel between pronouns and tense. See Partee (1973, 1984).

³²I assume that there are no accompanying pointing gestures in (20).

In particular, just as (21), can be uttered out of the blue to refer to some significant woman available in the discourse context, so too (22), uttered out of the blue, can be understood to describe the hypothetical scenario that is salient in the discourse context (the scenario in which the speaker buys the stereo):

(21) (Referring to a certain significant female) She left me. (Partee, 1973)

(22) (Looking at a high-end stereo in an electronics store) My neighbors would kill me.
(Stone, 1997)

Third, both types of expression allow for bound readings, where intuitively their semantic interpretation co-varies with the domain of some higher binding operator, as in (23) and (24):

(23) Every woman believes that she is happy.³³ (Partee, 1984)

(24) If a concert goer arrives late, he or she will not be permitted into the auditorium.³⁴
(Stone, 1997)

Finally, both types of expression allow for so called “donkey anaphora” readings, as witnessed by (25) and (26); crucially, just like ‘it’ in (25) co-varies with the indefinite NP ‘a donkey’ (without being within its syntactic scope), so in (26) the modal in the consequent ‘will’ co-varies with the sub-constituent of the antecedent clause ‘if the enemy captures it’.

(25) If a man owns a donkey, he beats it. (Partee, 1984)

(26) If a submarine cannot self-destruct if an enemy captures it, the enemy will learn its secrets. (Stone, 1997)

These data strongly suggest an analogy between modals and pronouns with respect to the kind of interpretive dependencies they allow; both search for an antecedent either previously linguistically

³³Similar examples are found in e.g. Partee (1973) and May (1977).

³⁴I cite the original example that Stone provides, but the effect is easily replicated without ‘will’: “If a concert goer arrives late, he or she might not be permitted into the auditorium.” The same goes for Stone’s other examples involving ‘will’ (in particular, example (26)).

introduced, as in (19)–(20), and (23)–(26) or provided by the context, as in (21)–(22).³⁵

Observe that conditionals, too, display this type of anaphoric behavior. For example, in (27), the second conditional depends on a scenario introduced by the first one; it is not simply evaluated against all epistemically accessible worlds, available in the context discourse initially. And the modal in the consequent of the second conditional is thus interpreted only relative to the hypothetical scenario introduced by the antecedent of the second conditional *relative to the hypothetical scenario described by the first conditional*.³⁶

- (27) If a wolf walks in, we will use the tranquilizer gun. If we manage to shoot it, we will be safe.

One way we can think about the observed anaphoricity (in the sense of “anaphoricity” described above) of modals and conditionals is as follows. It is customary to treat modals as quantifier expressions, quantifying over possible worlds.³⁷ We know that quantifiers require a contextually supplied domain restriction—“Everyone had fun today” does not mean that everyone *in the universe* had fun today.³⁸ The same goes for modals—like other quantifiers, they also require that their domain of quantification be further contextually restricted.³⁹ “A wolf might walk in” does not convey the meaning that in at least one world *out of all possible worlds* a wolf walks in. Rather it conveys a more restricted meaning—at least one world out of the relevant, epistemically accessible worlds is such that in it a wolf walks in. Since in the case of the modals the domain of quantification is just a set of worlds, the restrictor on the domain will likewise be a set of worlds (a possibility, for short). I suggest that modal anaphora resolution is a matter of retrieving the possibility that serves as the

³⁵Further support for the analogy between modals and pronouns with respect to their interpretive range is provided by examples from languages such as Warlpiri or American Sign Language (ASL) that allow for a single anaphoric expression to be ambiguous between pronominal and modal interpretation. See Bittner (2001) for data on Warlpiri, and Schlenker (2013) for data on ASL.

³⁶Just as, as we shall see, the anaphoricity of modals plays a key role in explaining the counterexamples like the one in (1)–(3), so this anaphoricity of conditionals plays a key role in explaining the counterexamples to MP and MT involving right-nested conditionals, as McGee’s (1985) original counterexample to MP, or Veltman’s (1985) counterexample to MT.

³⁷Cf. Kratzer (1977, 1981, 2012).

³⁸For more on quantifier domain restriction, see von Stechow (1994), and Stanley and Szabo (2000).

³⁹Typically, in an ordinary discourse, the restrictor will eliminate at least some worlds from the domain of quantification, on pain of redundancy. In principle, however, the restrictor need not eliminate anything. This can happen, for instance, if the restrictor is a necessary proposition.

restrictor on the domain of quantification of a modal operator. That is, modals (and conditionals) are anaphoric insofar as they require an anaphorically retrieved restrictor: just as with the antecedent of a pronoun, the restrictor can either be explicitly linguistically introduced in the discourse (e.g. by some modal utterance prior in the discourse), as in (19), (24), (26) and (27), or otherwise made prominent in the context, as in (22).⁴⁰

How this anaphoric dependency is resolved—i.e. how the restrictor is retrieved in a context—brings us to our second analogy between pronouns and modals. Crucially, modals exhibit the same kind of responsiveness to mechanisms of discourse coherence as pronouns. The idea is that mechanisms of discourse coherence affect the prominence of possibilities that are candidates for the restrictor of a subsequent modal, just as they can affect the prominence of referents that are candidates for subsequent pronominal anaphora; they render certain possibilities prominent to serve as a restrictor of a subsequent modal. We see this already in (19), where the second sentence elaborates on the hypothetical possibility described by the first. It is because this Elaboration relation holds that the modal ‘would’ in the second sentence is understood as restricted by the possibility described by ‘might’ in the first— all of the hypothetical scenarios *out of those epistemically accessible ones in which the wolf walks in* are such that the addressee is eaten first. (Similarly, it is because there is an Elaboration relation between the antecedent and the consequent in (24) that the modal ‘will not’ in the consequent is understood as restricted by the possibility described by the antecedent.) The Elaboration relation is what makes this scenario prominent for the subsequent modal to pick up on.

The import of coherence is crucial in making a certain possibility prominent as a restrictor of a subsequent modal. Note that mere sequencing of modals is not sufficient for the correct interpretation. That is, it is not always the case that the modal will be restricted by a possibility introduced by the immediately preceding modal (when there is an immediately preceding modal). The fact that one modal follows another does not suffice to establish that the hypothetical scenario described by the second modal further elaborates the one introduced by the first. We easily see this with examples

⁴⁰Does this mean that quantifiers are also anaphoric in my sense? The answer is yes. Like modals they require a restrictor, provided either by the non-linguistic context, or the prior discourse. Moreover, though the details exceed the scope of the present paper, there are good reasons to hold that the way in which the restrictor of a quantifier is retrieved in a context is analogous to the way in which the restrictor of a modal and an antecedent of a pronoun is, i.e. that quantifier domain restriction is sensitive to discourse structuring mechanisms. (See Chapter 1.)

like (28).⁴¹

(28) If a wolf walks in, it would eat you. But one probably won't walk in.

As before, the consequent of the first sentence in (28) elaborates on the information provided by the antecedent. The Elaboration relation between the antecedent and the consequent makes the hypothetical scenario introduced by the antecedent the most prominent one, and as a result, 'would' in the consequent further describes this scenario. Crucially, however, the modal 'won't' in the second sentence does not further elaborate upon scenario described by the two modals in the first sentence. The two sentences stand in a relation of Contrast, signaled by the discourse marker 'but', and are understood as contrasting two hypothetical scenarios—one in which a wolf walks in, and one in which one does not.

Intuitively, the Contrast relation requires that the first and the second sentence provide contrasting information about some body of information. A bit more precisely, the two bits of discourse provide contrasting information about some body of information regarding some common topic—in our example the topic is what is possible regarding a wolf's entrance.⁴² The first sentence already sets the stage in determining the body of information the contrast has to be about—(assuming that the conditional is uttered discourse initially) it is interpreted relative to the set of epistemically accessible worlds determined by the context discourse initially, describing what might be the case if a wolf walked in, *given this overall body of knowledge*. The second sentence, then, has to provide a contrasting bit of information, regarding the possibility of a wolf's entrance, about this body of information available discourse initially. The Contrast relation makes this body of information prominent, and consequently, the modal in the second sentence selects it as its restrictor—*given this*

⁴¹For similar examples, see Asher and McCready (2007).

⁴²The relevant topic is typically signaled by the cues in the information structure—the way the information is packaged together. One way of signaling this, in English, is by exploiting prosodic accents. For example, compare the following two utterances:

(i) John likes MARY.

(ii) JOHN likes Mary.

(31) is fine in the context in which we are wondering whom John likes, say, Mary or Sue, but not in the context in which we are wondering who likes Mary, say, Bill or John; the opposite is true of (32). For more on sentential focus, see Rooth (1992), and for more on information structure, see e.g. Roberts (1996a).

overall body of knowledge, a wolf probably won't come in. This is the intuitively correct interpretation: *given the overall body of knowledge, if a wolf walks in*, it would eat the addressee, but *given the same body of knowledge*, one probably won't walk in.

What this sort of example establishes is that, just as with pronouns, the impact of discourse coherence on making a certain possibility prominent as a restrictor for a subsequent modal is crucial. As the contrast between Elaboration and Contrast shows, precise discourse mechanisms govern the prominence of possibilities in a discourse. We can capture this idea in a way similar to the way in which we captured the effects of coherence on the prominence of candidate referents for pronoun resolution. Here is a first pass proposal: let the context represent a ranking of sets of worlds (possibilities, for short) that are candidates for domain restrictors of modals in a discourse according to their relative prominence—the most prominent being the top-ranked one. A modal simply retrieves the most prominent epistemically live possibility as the restrictor for its domain of quantification.⁴³ The prominence ranking of candidate possibilities, in turn, is affected by a range of linguistic mechanisms, most notably, those of discourse coherence; coherence relations make certain possibilities prominent for subsequent modal anaphora.⁴⁴

I shall assume that at the beginning of a conversation the top-ranked possibility is just a set of epistemically possible worlds—a set of worlds epistemically accessible from the actual world. The intuitive idea behind this assumption is the familiar idea that the ultimate goal of a conversation is to narrow down possible ways the actual world could be.⁴⁵ However, as the discourse progresses,

⁴³I will use the term “epistemically live possibility” to denote the possibility that is not ruled out given the relevant body of knowledge, i.e., which is accessible given the relevant epistemic accessibility relation. For a more precise notion of epistemic accessibility, see § 3.3.1 and § B.1.2.

⁴⁴Though limited space prevents me from going into details here, I have argued elsewhere that there are good reasons to treat these effects of coherence on the resolution of modal anaphora as linguistically encoded, rather than as byproduct of general reasoning. (See Chapter 4.) For instance, as Asher and McCready point out, the direct translation of (19) in Japanese is infelicitous, unless there is an overt discourse marker signaling an Elaboration relation (Asher and McCready, 2006.) This would be surprising if the elaboration reading, where the second modal is understood as elaborating on the scenario described by the first one, were merely a byproduct of general reasoning. Moreover, as I argue in Chapter 4, the mechanisms of discourse coherence sometimes force inconsistent interpretations of modal discourse, even when there are possible plausible alternative interpretations that could have been retrieved instead. That the inconsistent readings are retrieved even in the face of alternative common-sense interpretations is easily explained if the mechanisms of discourse coherence conventionally mandate such interpretations; however, this would be mysterious if the effects of these mechanisms were byproducts of pragmatic reasoning, since in such cases, we have overwhelming common-sense reasons to disprefer the inconsistent interpretations, and so override the effects of discourse coherence. These considerations, among others, suggest that the effects of the mechanisms of discourse coherence are a part of our linguistic repertoire, and should hence be reflected in the logical form of the discourse.

⁴⁵Cf. Stalnaker (1978). More precisely, since typically there is uncertainty about which world is the actual one (given that our knowledge is limited), the initial set of epistemically live possibilities will be a set of worlds W that contains, for

the prominence ranking changes; this change in ranking is precisely what we want to model, since this is precisely the change relevant for the interpretation of modals. The changes arise as an effect of introducing novel possibilities (e.g. through utterances containing modals or conditionals), and through discourse structuring mechanisms that change and affect the prominence ranking, in the ways described earlier.

Connecting this with our previous observations, the anaphoricity of modals is captured by requiring that the restriction on the domain of quantification be retrieved in a way similar to how the antecedent of a pronoun is—either provided by the context, or explicitly by the prior discourse. The way anaphora is resolved, in both cases, is determined by discourse structuring mechanisms, in particular, mechanisms of discourse coherence. Just as the mechanisms of discourse coherence affect the prominence of candidate referents for the resolution of a pronoun, so too they affect the prominence of a candidate possibility for a restrictor on a subsequent modal.

To sum up: modals are like pronouns in two crucial respects: (a) they are anaphoric expressions, and (b) the resolution of modal anaphora is responsive to the same mechanisms that pronoun resolution is responsive to. We now have all the ingredients we need to tackle the original counterexample.⁴⁶

each world w that is a candidate for the actual world, the set of worlds accessible from w . Of course, if we are concerned with the interpretation of discourse initial uses of modals, it will matter a great deal whose body of knowledge is relevant for determining the set of epistemically accessible worlds. Since I shall not deal with this issue in the present paper, I shall simply assume that the relevant body of epistemically accessible worlds will be contextually provided discourse initially. However, *pace* Kolodny and MacFarlane (2010), Yalcin (2007), and Moss (2015), the resources of my account can show how this body of information is selected in a context, even discourse initially. (See Chapter 4.) Relatedly, as I argue at length in Chapter 4, the resources of my account naturally explain the patterns of intra- and inter-contextual (dis)agreement involving modal language, that have been argued to favor the revisionary theories precisely on grounds of the alleged failure of context to determine an adequate body of information in such cases. (For the revisionary arguments based on (dis)agreement patterns, see e.g. Egan, Hawthorne, and Weatherson (2005); MacFarlane (2014). For alternative contextualist responses to these arguments see, e.g., Dowell (2011) and von Fintel and Gillies (2009). These contextualist accounts (much like Kratzer's account) assume general pragmatic mechanisms of context-sensitivity resolution that are too unconstrained to systematically account for the apparent counterexamples to MT, as well as the puzzling behavior of modals in various embedding environments, including, but not limited to antecedents of conditionals. See Chapter 4. for a detailed discussion.)

⁴⁶Note that, according to the theory developed here, modals express truth-conditional content, and the explanation of the counterexample exploits the difference in the truth-conditional content between the consequent of the big premise, and the small premise in (1)–(3). Several authors have challenged the view on which modal vocabulary expresses truth-conditional content on grounds that are not directly related to the failure of classical inferences. (See e.g. Egan, Hawthorne, and Weatherson (2005), Yalcin (2007), and Moss (2015).) Though the theory developed here, given its systematic account of context-change, has means for accounting for those challenges as well, addressing them here is beyond the scope of the present paper. (I address these issues in detail in Chapter 4.) I shall just note that we can, for example, easily explain the problematic behavior of epistemic modals embedded in antecedents of conditionals—one of the key data points used by Yalcin to argue against the truth-conditional accounts (Yalcin, 2007, 2011). Yalcin notes the discrepancy between the following two conditionals:

3.2.5 Yalcin's Counterexample Explained Away

We can now explain what is going on with (1)–(3) as follows: in (1), the consequent of the conditional is understood as elaborating on the possibility introduced by its antecedent (which is just the set of epistemically accessible worlds in which the antecedent holds), and thus, the Elaboration relation renders this possibility the most prominent one for subsequent anaphora. Consequently, the modal ‘likely’ in the consequent, which is searching for the most prominent epistemically accessible possibility, selects this possibility as the restrictor for its domain of quantification. The consequent is thus understood as further describing the possibility introduced in the antecedent, providing the intuitively correct restricted reading—the marble is likely red, *given that it is big*. In turn, we naturally understand (2) as being linked to (1) by the relation of Contrast. This is seen even more clearly if we insert an explicit discourse marker ‘but’ in (2): ‘But the marble is not likely red.’ Note that some such way of signaling contrast is required, for the discourse consisting of (1) followed by (2) to be felicitous.

As always, the question of which relation holds (and between which relata) is a matter of disambiguation in a discourse, much as in the case of (17). There are often discourse-internal, linguistic cues that signal a particular relation (e.g. a discourse marker ‘but’), but context can play a role in disambiguation as well. For example, here the initial context sets up a topic—the color of the marble, depending on a certain assumption about its size. (1) and (2) are then understood as providing

-
- (i) If it is raining and it might not be raining, ...
 - (ii) If it is raining and I/we don't know it, ...

The first one is odd, while the second perfectly fine, pointing to an apparent problem for the standard truth-conditional accounts of conditionals, that interpret modals as quantifying over some salient body of knowledge. Yet, the natural way to understand “it might not be raining” in the conditional above is as elaborating on the hypothetical raining scenario. This affects the resolution of modal anaphora—‘might’ is understood as quantifying over all the relevant epistemically accessible worlds in which it is raining, and so it is no surprise that the conditional winds up being bad. Note that my account does not predict that the conditional with the reversed order of conjuncts in the antecedent, i.e. of the form: “If it might p, and not p, then...,” will automatically be bad. This is a desired result since, as Dorr and Hawthorne (2014) note, switching the order of the conjuncts in some cases (in particular, in Yalcin's original example) makes the conditional felicitous. Yet, my account does not predict that all such cases will automatically be felicitous either. This is because the badness of such a conditional will depend on which coherence relations, and other interpretive dependencies, can be established between the two conjuncts in the antecedent, and between the conditional and the rest of the discourse in which it is embedded, for these factors can all affect the resolution of modal anaphora in a particular case. This is, again, a desired result. And the point holds more generally—embedding any sentence (including the original Yalcin's conditional) in a broader discourse might give rise to various interpretive dependencies that might yield a different interpretation than the one we get when interpreting the same sentence-type out of the blue.

contrasting bits of information about some body of information regarding this topic. As before, the first sentence already sets the stage in determining the body of information the contrast has to be about—since (1) is uttered discourse initially, it is interpreted relative to the overall body of knowledge available discourse initially. Thus, one understands (1) and (2) as providing two contrasting bits of information about *the overall body of knowledge or information available discourse initially*, regarding the likelihood of the marble being red, depending on a certain assumption about its size. Namely, given the overall available information discourse initially, *if the marble is big*, then it's likely red, but *additionally*, given the overall available information, the marble is not likely red (given no particular assumption about its size).

A bit more precisely, the effect of Contrast in (1) and (2) is the same as in (28). (1) and (2) are understood as contrasting two different bits of information about some initial overall body of information, regarding the likelihood of the redness of the marble, given some or no assumption about its size. The first sentence already sets the stage in determining the body of information that the contrast has to be about—the conditional is interpreted relative to the set of epistemically accessible worlds determined by the context discourse initially,⁴⁷ describing what the likelihood of the redness of the marble is, *given this body of knowledge*, provided that it's big. The second sentence thus has to provide a contrasting bit of information about this same body of information (i.e. the set of epistemically accessible worlds, discourse initially) regarding the likelihood of the marble being red. Thus, this body of information is made prominent by Contrast. Consequently, the modal in (2), selects this body of information as its restrictor, conveying that the marble is not likely red *given this overall body of knowledge* (given no particular assumption about its size). Thus, we see that the two occurrences of 'likely' in (1) and (2) are interpreted differently, much like the two occurrences of 'she' are in (12) and (13). And thus, (1)–(3) is no more an instance of MT than (12)–(14) is.⁴⁸

⁴⁷Provided that the conditional is uttered out of the blue, which, by assumption, it is.

⁴⁸Note that, while my account maintains that there is a change in context that transpires in (1)–(3), affecting the interpretation of modal expressions, the account does not maintain that such a change is somehow 'illicit'. In an important sense, there is no way of holding the context fixed throughout the course of (1)–(3). This is because (1)–(3) harbors linguistic elements (in particular, modals, antecedents of conditionals, and coherence relations) part of the meaning of which is to change the context in a way that affects the truth-conditions (by introducing certain possibilities, making them prominent, and demoting others).

3.2.6 Conclusion to Part I

In § 3.2, I argued that the lesson of the alleged counterexamples to classically valid patterns of inference is *not* that we need a revisionary, non-truth-conditional semantics for modal vocabulary that invalidates these patterns, but rather that we need a different, more constrained picture of how context affects interpretation. In particular, I argued that modals are like pronouns not merely in that they are anaphoric expressions, but in that they are sensitive to some of the same interpretive mechanisms of anaphora resolution as pronouns are—in particular, the mechanisms of discourse structure. The “counterexamples” arise due to a failure to appreciate the import of these mechanisms on the interpretation of modals. But, a failure to appreciate how context changes with the evolution of a discourse would lead us to mistakenly treat even the examples like (12)–(14) as counterexamples. That the effect of context-change on modals is more subtle and less well explored only makes the failure to appreciate it that much more dangerous.⁴⁹

The mechanisms of context-sensitivity that I have argued are relevant for the interpretation of modals, though pervasive and nuanced, are not unruly. We can devise a *formal* theory of context-change that models the import of these mechanisms on the interpretation of modals in a rule governed and systematic fashion, and indeed, this is what I do in § 3.3, and in the Appendix (§ B.1). Thus, the formalism developed there should be understood as an existence proof in support of the main argument of § 3.2. The theory accommodates the anaphoricity of modals, while preserving their standard truth-conditions, and maintaining the validity of the classical patterns of inference. Though the resulting semantics departs from the standard semantics for modal discourse, it does so in a way that allows us to capture the interplay between discourse structuring mechanisms and the interpretation of modal anaphora. It is precisely this departure that preserves classical logic, and permits modal utterances to express standard truth-conditional content.

It is important to note that the account developed here opens a promising line for further development. In particular, philosophers often identify context-sensitivity of various philosophically interesting terms—e.g., of implicit restriction on quantifier domains, knowledge ascriptions, vague

⁴⁹As I argue in more detail in Chapter 4, the revisionary accounts in fact are not well suited to explain the full range of data concerning modal anaphora. For instance, they do not capture the modal anaphora data described in § 3.2.4. My account, in turn is precisely designed to account for these data. Thus, apart from preserving truth-conditions and classical inference patterns, the account is well motivated on independent grounds.

predicates, normative terms—and use this context-sensitivity to motivate broad philosophical conclusions. But in doing so, they typically assume a model of context-sensitivity that is resolved by freely selecting one candidate resolution out of an open-ended list of potential ones, through general pragmatic mechanisms (e.g. speaker’s intentions, and non-linguistic contextual cues). This predicts a level of flexibility that often fails to be born out in practice, and this flexibility in turn shapes the philosophical arguments that appeal to such context-sensitivity. Though arguing for this in full generality obviously exceeds the scope of the present paper, going beyond modals and pronouns, the kinds of tools developed here open a path for exploring the potential systematic constraints on other context-sensitive expressions. If the approach advocated here can be extended to capture the contextual-sensitivity of these other kinds of expressions, then that would show that contexts are much less powerful, and the resolution of context-sensitivity a much more constrained process than what philosophers typically assume in their arguments.⁵⁰

3.3 Part II: A New Semantics for Modality

In what follows, we have two tasks: first, to modify the standard truth-conditional account, developed in Kratzer (1977, 1981, 1983, 2012), in order to accommodate the anaphoricity of modals and conditionals, and second, to develop an account of prominence and context-change that explains how modal anaphora is resolved in a discourse. I will undertake the first task in § 3.3.1, and the second task in § 3.3.2. As I shall demonstrate, the formal semantics that emerges assigns standard truth-conditions to modals and conditionals, and as I prove, in the Appendix (§ B.1), preserves classical logic.

3.3.1 Anaphoricity and Truth-Conditions

As mentioned earlier, the classic account treats modals as quantifying over a contextually supplied modal base, i.e. a set of contextually relevant worlds that comprise the domain of quantification. A modal then requires that a particular relation holds between the proposition expressed by the

⁵⁰Indeed, in Chapter 1. I argued that the kind of account developed here extends to numerous other cases of context-sensitivity, and that the kind of mechanisms I described here are suited to explain context-sensitivity quite generally.

clause embedded under a modal (which is typically called *the prejacent*⁵¹) and the modal base.⁵² For example “It might rain” requires that there is some possible world in the contextually supplied modal base in which it rains. Let ‘ q ’ denote the proposition expressed by the prejacent ϕ of ‘*might* ϕ ’ in a given context and ‘ R ’ an accessibility relation, that specifies the epistemically accessible worlds, i.e. the modal base.⁵³ Then the truth-conditions expressed by ‘*might* ϕ ’, relative to a context of utterance, are:

Definition 3.3.1.

$$\{w \mid \exists w' : wRw' \ \& \ w' \in q\}$$

That is, ‘*might* ϕ ’ is true at a world w (relative to a context c) just in case there is some world among the worlds epistemically accessible from w in which the proposition, q , expressed by the prejacent in c , holds. Anaphoric dependency is easily factored into the standard truth-conditions of modals explicitly as follows. (I shall use ‘ $M(p, q)$ ’ for the truth-condition expressed by an utterance of ‘*might* ϕ ’, where q is the proposition expressed by the prejacent ϕ of the utterance of ‘*might* ϕ ’, and p the proposition corresponding to an anaphorically retrieved restrictor. I omit the details about how context determines truth-conditions (and, in particular, the restrictor p) here. This will be the topic of § 3.3.2.)

Definition 3.3.2.

$$M(p, q) = \{w \mid \exists w' : wRw' \ \& \ w' \in p \ \& \ w' \in q\}$$

This gives us the resources to define other modal expressions: as is standard, ‘*must*’ is the universal dual of ‘*might*’.⁵⁴ For probability modals, such as ‘*likely*’, we need a probability measure

⁵¹So, in ‘*might* ϕ ’, ϕ is the prejacent.

⁵²This captures the main insights from Kratzer’s account, though I suppress formal details for simplicity. In particular, I suppress an ordering source parameter, which provides an ordering of worlds in a modal base according to some contextually provided standard, and the formal machinery that serves to derive a modal base parameter. (Cf. Kratzer (1977, 1981, 1983, 2012).) We could easily factor these elements back in. As Stone points out, modal base and ordering source parameters, as specified in Kratzer’s account, cannot accommodate the anaphoricity of modals, since both are determined in complex ways, and neither provides a semantic parameter that can be contributed by prior discourse, so a modification of Kratzer’s account is needed regardless (Stone, 1997).

⁵³The relation R plays the role of a Kratzerian modal base. On Kratzer’s account the modal base is contextually determined in complex ways. For our purposes, we can simplify even further, and let R be supplied by the model, since we are not dealing with discourse-initial uses of modals. See Chapter 4 for details on how R is determined in context, discourse initially.

⁵⁴For the definition of truth-conditions of ‘*must* ϕ ’, see § B.1.2.

over the accessible worlds. Let \mathcal{P} be a probability measure over the set of worlds in the universe W , that maps each subset of W to $[0, 1]$, satisfying the following constraints:

- $\mathcal{P}(W) = 1$
- $\mathcal{P}(p \cup q) = \mathcal{P}(p) + \mathcal{P}(q)$, when p and q are disjoint subsets of W .⁵⁵

Then, where q is the proposition expressed by the prejacent ϕ of an utterance of ‘*likely* ϕ ’, and p the proposition corresponding to the anaphorically retrieved restrictor, the truth-conditions, $P(p, q)$, expressed by the utterance of ‘*likely* ϕ ’ are as follows:

Definition 3.3.3.

$$P(p, q) := \{w \mid \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p \ \& \ w' \in q\}) / \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p\}) > .5\}^{56}$$

As expected, the truth-conditions expressed by an utterance of ‘*likely* ϕ ’ are the set of worlds such that, for each w in the set, the ratio of the probability that an R -accessible world from w be a p and q world to the probability that the R -accessible world from w be a p -world is greater than .5; i.e. an utterance of ‘*likely* ϕ ’ is true in w just in case given our modal base, the conditional probability of the prejacent, q , given the p -restricted modal base is greater than .5.⁵⁷

Finally, we need to specify the truth-conditional contribution of a conditional. We can easily factor in the anaphoric potential into the truth-conditions of a conditional, just as we did with modals, while otherwise preserving the standard truth-conditions. Where p , as before, is the anaphorically retrieved restrictor with respect to which the conditional is uttered, q corresponds to the proposition

⁵⁵I shall assume that the probability measure is supplied by the model. Alternatively, it could be provided by context. The choice is inessential for the context-sensitivity I’m aiming to model. For simplicity, I assume that W is finite. I also assume that \mathcal{P} is a regular probability measure, i.e. assigning non-zero probability to all non-empty sets of worlds. Insofar as $\mathcal{P}(p)$ represents the prior, if $\mathcal{P}(p) = 0$, p wouldn’t really be a possibility.

⁵⁶It is typically assumed that the restrictor on the domain of quantification, p , is not empty. But if we wanted to allow for the possibility in which p is empty, we could modify our definition in a following way: $P(p, q) := \{w \mid \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p \ \& \ w' \in q\}) > 1/2 \ \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p\})\}$. Also, one might wonder whether we should always impose a threshold of .5, or perhaps the threshold might vary with the context. For simplicity I choose the former option. The choice is inessential for our purposes.

⁵⁷I depart from Kratzer (1991) in suggesting this quantitative characterization of the truth-conditional contribution of probability operators, rather than a qualitative one cashed out in terms of relative likelihood. There are well known problems with a purely qualitative account. For a detailed discussion, see Yalcin (2010); for a discussion of the prospects of basing the quantitative notion of probability on a qualitative one, see Kratzer (2012, ch 2.) and Holliday and Icard (2013). This particular choice I make here, though independently motivated, is inessential to the overall point of the paper. What matters for us is that the anaphoric potential of a probability modal is correctly captured. We can build the anaphoric dependency into the truth-conditions in the way suggested, regardless of what we think the correct account of the truth-conditions is. The point holds for the truth-conditions of other modals as well.

expressed by the antecedent, and r to the one expressed by the consequent, an utterance of a conditional expresses truth-conditions corresponding to a set of worlds such that for each w in the set, all the worlds w' , R -accessible from w , that are p and q worlds, are r worlds as well; i.e. an utterance of a conditional is true in w if and only if all the p and q worlds in the modal base are r worlds as well.⁵⁸

Definition 3.3.4.

$Cond(p, q, r) := \{w \mid \forall w' : wRw', \text{ if } w' \in p \ \& \ w' \in q, \text{ then } w' \in r\}$

This in essence preserves the standard truth-conditions associated with a conditional, but factors in the fact that a conditional itself is always evaluated against some prominent body of information, that need not always correspond to the complete unrestricted set of epistemically live worlds.

This concludes my characterization of the truth-conditional contribution of modals and conditionals. According to the proposed account, the utterances containing modals and conditionals express the same truth-conditions one would expect given the canonical account,⁵⁹ except that the anaphorically retrieved restrictor is factored in separately. This allows us to flexibly track the way it is retrieved in a context. It is important to note that, provided that all anaphoric restrictors are resolved to the same set of worlds (e.g. to the set of all possible worlds), we get exactly the truth-conditions for modals we would expect in standard propositional modal logic (our conditional is a standard strict-conditional).⁶⁰ However, we have yet to specify the effects of context and context-change on the determination of truth-conditions of a given utterance. As I have argued, the most important impact of context for us will be in the resolution of anaphoric dependencies of modals and conditionals, i.e. in the role of context in determining the restrictor of the modal base of modals (and conditionals). In § 3.3.2, I lay out a formal theory of context-change, which will allow us to

⁵⁸I ignore the ordering source, but we could easily factor this parameter in, and then state the truth-conditions by making a conditional true (in a world w and at a context c) just in case all the p and q worlds that are closest to w in the modal base are r worlds as well. Apart from factoring in the anaphora (and modulo the abstraction of the ordering source parameter), the truth-conditions I propose here match the ones developed in Kratzer (1983). Kratzer's account makes a conditional true in a world and at a context just in case all the (closest) antecedent worlds within a modal base are consequent worlds as well.

⁵⁹Cf. Kratzer (1977), Kratzer (1981), Kratzer (1983) and Kratzer (2012)

⁶⁰More precisely, on the assumption that R is reflexive and transitive, we would get the system $S4$. For a proof, see section § B.1.7. This is a common and natural assumption. Reflexivity, e.g., ensures that *must* p entails p , and also that p entails *might* p , and transitivity ensures, e.g., that *might*(*might* p) entails *might* p , again, provided that all anaphoric restrictors are resolved to the same set of worlds (e.g. to the set of all possible worlds).

tackle the alleged counterexample we began with.

3.3.2 Truth-conditions and the Dynamics of Context-change

In order to capture the idea that the restrictor of a modal expression is the most prominent possibility, we need to have a way of modeling the prominence of possibilities in a discourse context.⁶¹ As suggested in § 3.2.4, we can let a context represent a ranking of possibilities—sets of worlds, or propositions—according to relative prominence, the top-ranked possibility being the most prominent one. One way to think about such a context is as a sort of a conversational record, in the sense of Lewis (1979), an abstract “scoreboard” that tracks the moves and contributions interlocutors make in the flow of a discourse, and that comprises information relevant for interpretation, such as who’s speaking, what the conversation is about, etc. Crucially, for us, the conversational record tracks propositions put into play in the course of a conversation as well as their relative prominence. Since this is the only aspect of the scoreboard that will matter for modal anaphora, we can abstract away from all other aspects that might be otherwise needed. In this spirit, we can see my context as modeling one aspect of a Lewisian scoreboard—the relative prominence of possibilities within a discourse.

We will let a context represent the ranking of possibilities according to their relative prominence. To model a context, we will exploit the idea of an assignment function, modeled as a stack, in the sense of theoretical computer science; an assignment function, or a stack, specifies an ordered sequence of worlds.⁶² Let the context be a set of such stacks.⁶³ So, a context is just a set of sequences of worlds. For a given context G , let ‘ \mathbf{w}_i ’ be a variable that stores a world at the i^{th} position of every stack in G .⁶⁴ Then, relative to G , ‘ \mathbf{w}_i ’ stores a set of worlds that collects the

⁶¹I sketch the key bits of the formal account here, but for the fully precise formal definitions and details, consult § B.1.

⁶²Formally, a stack is an assignment function, mapping a finite convex subset of \mathbb{N} to a set of worlds together with an undefined value \perp .

⁶³This formalization, unlike the one exploiting a single stack ranking sets of worlds, allows us to keep track of not only sets of worlds, but also relations between individual members of different sets. (Cf. Bittner (2014), Brasoveanu (2010), and van den Berg (1996).) Though this is not strictly speaking crucial for our present purposes, it does become crucial once we want to develop an integrated semantics for anaphora both in the modal, and in the pronominal domain.

⁶⁴In describing context change I am adopting the following strategy. I define a dynamic language that models the dynamics of prominence, and then provide a translation of a fragment of English into this language. The dynamic language has atomic expressions (propositional constants (p, q, r), and variables (\mathbf{w}_i for $i \in \mathbb{N}$), conditions (propositional expressions comprising set of atoms closed under \wedge and \neg), and update expressions, which we will define and describe shortly. For details see § B.1 (in particular, § B.1.1, § B.1.4 and § B.1.7).

worlds at the i^{th} position of every stack in G . This, in particular, will allow us to keep track of the propositions that have been introduced and promoted (as well as demoted), during the course of a discourse. I shall assume that every stack in G begins with the 0^{th} position, as the top-ranked position on the stack, and that for each position n , the position $n + 1$ is one position lower in the ranking. Let ' G_i ' denote the set of worlds that collects the worlds at the i^{th} position of every stack in G .⁶⁵ Then, G_0 is the top-ranked proposition in G , and for each n , G_{n+1} is a proposition one step lower in the ranking. This allows us to keep track of the relative prominence of possibilities for modal anaphora.

As noted in § 3.2.4, I shall assume that at the beginning of a conversation, the top-ranked possibility is just a set of epistemically accessible worlds. In turn, as the discourse progresses utterances affect the prominence ranking of candidate possibilities. As we have already seen in an informal way in § 3.2, utterances can promote novel possibilities and re-rank the ones introduced prior in the discourse. The simplest way to model this is to represent utterances as updates to the context that change the relative prominence of candidate possibilities. Formally, an update is represented as a relation between an input context and an output context where the update is true (relative to an input context and a world of evaluation w) just in case it relates an input context to some non-empty output context (relative to w).⁶⁶

Thus, utterances have a two-fold contribution—they express truth-conditions (as per § 3.3.1), but they also contribute updates that change the context, by updating the ranking of possibilities. Both aspects of the interpretation are crucial. The updates associated with utterances capture the way in which these utterances change the context; in turn, such a dynamically maintained context determines the relevant truth-conditional content of a given utterance. Thus, we effectively model the two-way interaction between a context and an utterance: on the one hand, an utterance changes the context, and on the other, such a dynamically evolving context determines its truth-conditional content.

⁶⁵Thus, G_i is the proposition stored at the i^{th} position in G .

⁶⁶This is the standard notion of truth exploited in dynamic semantics. (See Dekker (2011).) Recall that on my account utterances also express ordinary truth-conditional content. The dynamic notion of truth defined here is exploited to capture the logic of context-change. We will see that once we have the logic of context-change in place, the ordinary truth-conditional content (and the ordinary notion of truth) falls out of it straightforwardly.

Before laying down the updates associated with modals and conditionals, we need some preliminaries. We said that utterances contribute updates to a context by promoting and re-ranking propositions, i.e. possibilities, but they also express propositions. I am going to separate the role that propositions play in anaphora, from the one they play in semantic composition. To illustrate why we need this separation, take the example: “A wolf might eat Harvey”. First, we need to compose the proposition expressed by the prejacent, *that a wolf eats Harvey*, with the modal ‘*might*’. But if this proposition automatically became the top-ranked possibility in the context, since the modal ‘*might*’ selects the top-ranked epistemically live proposition as its restrictor, so long as there’s some epistemically accessible world in which a wolf eats Harvey, we would predict that the restrictor for the modal in “A wolf might eat Harvey” is the proposition comprising all the epistemically live worlds in which a wolf eats Harvey. But, obviously, this proposition should not automatically be made a restrictor, as witnessed by examples like: “A wolf might walk in. It might eat Harvey.”

There are several ways to get around this problem, but one elegant way is to store separately the propositions that enter into semantic composition, and the ones that are candidate restrictors for subsequent modals. That insures that the truth-conditional contribution does not automatically interfere with prominence ranking.⁶⁷

To this end, I reserve a designated position on each stack in a context G that does not affect the ranking. Let ‘*comp*’ (for compositional) denote the proposition that comprises all the worlds stored at this position in every stack in G .⁶⁸ We exploit ‘*comp*’ to store, relative to G , each bit of propositional content that enters into semantic composition.⁶⁹ Here is how we do that. We will treat expressions that contain no proper propositional subparts as atomic formulae in our system. When ϕ is an atomic formula (and so, does not contain modals, or conditionals), its interpretation will just be the simplest update, defined below in 3.3.5, which stores the proposition expressed by that formula in the input context G as a new value of ‘*comp*’. This update relates the input context G to an output context H , (relative to a world of evaluation w) just in case H differs from G in at most

⁶⁷This is not to say that one and the same proposition cannot play a role both in semantic composition and also be prominent for anaphora resolution. This is just to say that the mechanisms by which a proposition completes these two roles are best kept separate.

⁶⁸So, now, formally, a stack is just a function from a finite convex subset of \mathbb{N} plus *comp* to a set of worlds together with an undefined value, \perp .

⁶⁹Similarly as before, ‘ G_{comp} ’ denotes the value of ‘*comp*’ in G .

the value of ‘*comp*’, and the value of ‘*comp*’ is the proposition that ‘*p*’ expresses in G and w . That is, where G is the input context, H the output context, w the world of evaluation, and $G \sim_n H$ just in case G differs from H in (at most) the n^{th} position, we can define this update as follows:

Definition 3.3.5.

$$\llbracket \langle \text{comp} := \phi \rangle \rrbracket(w, G, H) \text{ iff } G \sim_{\text{comp}} H \ \& \ H_{\text{comp}} = \llbracket \phi \rrbracket^{G, w}.$$

When ϕ is non-atomic, it will be interpreted as a more complex update. However, crucially, its truth-conditional contribution determined by the update will still be stored as a value of ‘*comp*’ of the output context, in a way to be specified presently.⁷⁰

Observations in § 3.2.3 and § 3.2.4 show that modals require an anaphorically retrieved restrictor, but also introduce possibilities that can subsequently be picked up by other modals. I suggested that the anaphorically retrieved restrictor will be the top-ranked epistemically live possibility. Let ‘@ E ’ denote the top-ranked epistemically live possibility (the top-ranked possibility that is a subset of epistemically accessible worlds) in a given context.⁷¹ Intuitively, we want the update associated with an utterance of ‘*might* ϕ ’ to have the following effect: first, (as with all other updates) its truth-conditions (as defined in 3.3.2) are stored as the value of ‘*comp*’ of the output context. Second, it introduces a possibility comprising the top-ranked epistemically live worlds in which the prejacent holds. Here is how we can informally describe the effect on context that an update associated with ‘*might* ϕ ’ carries out. Where G is an input context, and H an output context, the update relates an input context G to an output context H just in case there are intermediate contexts G' and G'' and:

- a. G' is a result of updating G with the update associated with prejacent of the modal (recall that since updates associated with both atomic and non-atomic formulae will store their truth-conditional content as the value of ‘*comp*’, G'_{comp} will just be the truth-conditional contribution of the prejacent in G , i.e. $\llbracket \phi \rrbracket^{G, w}$),
- b. G'' is just like G' except that it stores the top-ranked epistemically live possibility in which

⁷⁰In particular, 3.3.5–3.3.12 are clauses of a single recursive definition of update expressions in our dynamic language. See the Appendix (§ B.1) for details.

⁷¹Formally, we include in the basic vocabulary of our dynamic language unary predicates and a unary operator ‘@’. Where ‘ E ’ is a unary predicate and ‘@’ a unary operator, ‘@ E ’ is an atom. ‘@ E ’ is interpreted as taking a property denoted by ‘ E ’ and delivering the top-ranked proposition satisfying it, denoted by ‘@ E ’.

the prejacent holds (which is just $(G'_{comp} \cap \llbracket @E \rrbracket^{G,w})$) as a novel top-ranked possibility, and pushes all other values one position down, and

- c. the final output context H is just like G'' except that it stores the truth-conditional content expressed by the utterance as a new value of ‘*comp*’, which by 3.3.2, is the set that comprises all the worlds w such that there’s some world epistemically accessible from w in which both the restrictor (the top-ranked epistemically live possibility in the input context G) and the prejacent hold (i.e. $M(\llbracket @E \rrbracket^{G,w}, G'_{comp})$).

Putting all this together, we can now define the update associated with ‘*might* ϕ ’. Let us define a relation between contexts, \approx_n , where for any two contexts G and H , this relation holds just in case H is obtained from G by storing a novel value for n and pushing all other values one position down in the ranking; more precisely, $G \approx_n H$ just in case H is identical to G up until n , it differs from G in (at most) the n^{th} position, and for all m , such that $m \geq n$, $G_m = H_{m+1}$.⁷² Then, where K is an update associated with the prejacent ϕ , we define the update associated with an utterance of ‘*might* ϕ ’, that carries out the steps in (a.)–(c.), in the following way:

Definition 3.3.6.

$\llbracket \text{MIGHT}(@E, K) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ G'' \sim_{comp} H \ \& \ H_{comp} = M(\llbracket @E \rrbracket^{G,w}, G'_{comp})$.

Let us see how this works through an example.

(29) A wolf might walk in.

Given 3.3.5 and 3.3.6, we can represent (29) as follows. Let ‘ p ’ stand for the prejacent (‘a wolf walks in’); since atomic, by 3.3.5, it’s just interpreted as $\langle comp := p \rangle$. Thus, we get:

- $\llbracket \text{MIGHT}(@E, \langle comp := p \rangle) \rrbracket(w, G, H)$ which by 3.3.6 holds just in case there is a G' and G'' such that: $\llbracket \langle comp := p \rangle \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ G'' \sim_{comp} H \ \& \ H_{comp} = M(\llbracket @E \rrbracket^{G,w}, G'_{comp})$.

⁷²Even more precisely: $G \approx_n H$ iff $\{g_{0,n} + g_{n+1} \dots | g \in H\} = G$ and $G_{comp} = H_{comp}$. See section § B.1.3 of the Appendix for details.

- By 3.3.5, $\llbracket \langle comp := p \rangle \rrbracket(w, G, G')$ holds just in case $G \sim_{comp} G'$ & $G'_{comp} = \llbracket p \rrbracket^{G,w}$. That is, it holds just in case G'_{comp} is the proposition expressed by p in G and w , namely, the proposition *that a wolf walks in*.
- Moreover, the possibility corresponding to a set of top-ranked epistemically live worlds in which a wolf walks in is introduced as a novel top-ranked possibility, pushing all other possibilities further down in the ranking (thus we get G'').
- Finally, the proposition expressed by the modal utterance is stored as the value of ‘*comp*’ in the final output context H , which is otherwise just like G'' ; by 3.3.2, this proposition stored as the value of ‘*comp*’ in H corresponds to a set of worlds R -related to some world in which both the anaphorically retrieved restrictor and the prejacent of the modal hold (i.e. $M(\llbracket @E \rrbracket^{G,w}, G'_{comp})$). The anaphorically retrieved restrictor is the top-ranked epistemically live possibility in the input context G , (i.e. $\llbracket @E \rrbracket^{G,w}$) which, assuming that (29) is uttered out of the blue, just is the set of epistemically accessible worlds discourse initially. Thus, the proposition expressed by (29), the value of *comp* in the output context H , is the proposition that some of the epistemically accessible worlds discourse initially are such that in them a wolf might walk in.

Putting all this together, we get that (29) (a) expresses the proposition that it is compatible with what is known discourse initially that a wolf walks in, and (b) introduces a possibility comprising the top-ranked epistemically live worlds in which a wolf walks in, making it available for subsequent modal anaphora. This is just the desired result.

It is now also easy to see what the updates associated with utterances containing ‘*must*’ and ‘*likely*’ should look like. They will exactly parallel the update associated with utterances containing ‘*might*’, the only difference being in the truth-conditions. So, the update associated with ‘*likely*’ proceeds in exactly the same steps (a.)–(c.). The only difference will transpire in step (c.), where now the truth-conditional content stored as a value of ‘*comp*’ of the final output context, (H), will be the truth-conditional content expressed by an utterance containing ‘*likely*’, which, by 3.3.3, is $P(\llbracket @E \rrbracket^{G,w}, G'_{comp})$, i.e. a proposition that requires that given our modal base, the conditional probability of the prejacent, G'_{comp} , given the appropriately restricted modal base, $\llbracket @E \rrbracket^{G,w}$, is greater

than .5. So, the update associated with ‘*likely*’ can be defined as follows:⁷³

Definition 3.3.7.

$\llbracket \text{LIKELY}(@E, K) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ G'' \sim_{comp} H \ \& \ H_{comp} = P(\llbracket @E \rrbracket^{G,w}, G'_{comp})$

Finally, we need to specify the update associated with a conditional. The update will proceed in a similar fashion as before, with one difference: now we have to first process the update associated with the antecedent, and then with the consequent. That is, where G is an input context, the update first stores the proposition expressed by the antecedent (in G , relative to the world of evaluation w), as the value of ‘*comp*’ and introduces the top-ranked epistemically live worlds in which the antecedent holds as the top-ranked possibility (pushing all other possibilities one position down). Relative to thus obtained intermediate context (G''), it stores the proposition expressed by the consequent (in G'' , relative to w), as the value of ‘*comp*’ and introduces the top-ranked epistemically live worlds in which the consequent holds pushing all other possibilities one position down in the ranking, resulting in the intermediate context (G'''). Lastly, the final output context (H) differs from the intermediate context G''' only insofar as it stores the propositional contribution of the conditional as the value of ‘*comp*’: as per 3.3.4, it stores the proposition true at a world w just in case all the epistemically accessible worlds from w in which the anaphorically retrieved restrictor and the antecedent hold are such that the consequent holds in them as well.⁷⁴ More precisely, the update relates an input context G and an output H just in case there are some contexts G' , G'' , G''' and G'''' , and:

- a'. G' is a result of updating G with the update associated with the antecedent (thus storing the truth-condition expressed by the antecedent in G , as the value of ‘*comp*’ in G'),
- b'. G'' is just like G' except that it stores the top-ranked epistemically live possibility in which the antecedent holds, as a novel top-ranked possibility (G''_0), and pushes all other values one position down,

⁷³For the definition of an update associated with ‘*must*’, see § B.1.4.

⁷⁴As before, I abstract away from the contribution of the ordering source.

- c'. G''' is the result of updating G'' with the update associated with the consequent (thus storing the truth-condition expressed by the consequent in G'' , as the value of ‘comp’ in G'''),
- d'. G'''' is just like G''' except that it stores the top-ranked epistemically live possibility in which the consequent holds, as a novel top-ranked possibility (G_0''''), and pushes all other values one position down, and finally,
- e'. the final output context H is just like G'''' except that it stores the truth-conditional content expressed by the whole utterance of the conditional as a new value of ‘comp’, which by 3.3.4, just is $Cond(\llbracket @E \rrbracket^{G,w}, G'_{comp}, G'''_{comp})$, i.e. a proposition that requires that in all the epistemically accessible worlds in which the antecedent (G'_{comp}) and the restrictor for the conditional ($\llbracket @E \rrbracket^{G,w}$) hold, the consequent (G'''_{comp}) holds as well.

Putting all this together, where K_1 and K_2 represent updates associated with the antecedent and the consequent, respectively, we define the update associated with the conditional as follows:

Definition 3.3.8.

$\llbracket \text{IF}(@E, K_1, K_2) \rrbracket(w, G, H)$ iff there is a G', G'', G''' and G'''' such that
 $\llbracket K_1 \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ \llbracket K_2 \rrbracket(w, G'', G''') \ \& \ G''' \approx_0 G'''' \ \& \ G_0'''' = G'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \ \& \ G''''_{comp} \sim H \ \& \ H_{comp} = Cond(\llbracket @E \rrbracket^{G,w}, G'_{comp}, G'''_{comp})$

This completes the specifications of updates associated with modals and conditionals. As I have shown, we capture both aspects of their behavior—namely, we characterize the truth-conditions expressed by an utterance containing a modal or a conditional, and the way in which such an utterance changes the context in which it is uttered. The updates associated with utterances change the context, by updating the prominence ranking of possibilities that are candidate restrictors for subsequent modals and conditionals; the choice of a restrictor in turn affects the truth-conditions of an utterance containing a modal or a conditional. More generally, the updates associated with utterances affect the way in which the context evolves as the discourse progresses; the context in turn determines the truth-conditions expressed by the utterances. Both aspects of interpretation are crucial, and they are interrelated—unless we captured the change in the context prompted by an update associated with an utterance containing a modal or a conditional, we would not be able to

tell how the modal can make a possibility available for subsequent anaphora; and unless we calculated in the anaphoric dependency of utterances containing modals or conditionals, we would not be able to correctly predict which proposition a given utterance containing a modal or a conditional expresses, since the anaphorically retrieved restrictor crucially factors into its truth-conditions.

As we have seen, the updates in 3.3.5–3.3.8 all store their corresponding utterances’ truth-conditional content as the value of ‘*comp*’ of the output context. But, we also want to characterize what it takes to *assert* this content. Minimally, an assertion of an utterance requires that the proposition expressed holds at the world of evaluation. Plausibly, it also makes the possibility associated with the asserted content prominent. We can capture this by ensuring that an assertion promotes the set of top-ranked epistemically live worlds in which the asserted content holds as a novel top-ranked possibility,⁷⁵ and requires that the actual world be within that set. We can introduce a simple assertion update that achieves this effect:

Definition 3.3.9.

$\llbracket \text{ASSERT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \approx_0 H \ \& \ H_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w}$ & $w \in H_0$.

We now have almost all the ingredients we need to capture MT. We need to introduce one last thing—negation. The truth-conditional contribution of ‘*not* ϕ ’ is simple—it is true (at a context and relative to a world w) just in case w is a non- ϕ world. We can let the update associated with an utterance of ‘*not* ϕ ’ simply store the complement of the truth-condition expressed by ‘ ϕ ’ in the input context (relative to w), as the value of ‘*comp*’ of the output context. We define this update as follows:

Definition 3.3.10.

$\llbracket \text{NOT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \underset{comp}{\sim} H \ \& \ H_{comp} = \llbracket \neg comp \rrbracket^{G,w}$, where $\llbracket \neg comp \rrbracket^{G,w} = \mathcal{D}_\omega \setminus \llbracket comp \rrbracket^{G,w}$, where \mathcal{D}_ω is a set of possible worlds provided by the model.

Finally, we can now tackle the task of specifying MT pattern of inference. *Prima facie*, we can represent the pattern of the form [*if* ϕ , ψ ’, ‘*not* ψ ’ : ‘*not* ϕ ’], as follows. Let $T_d(\phi)$ and $T_d(\psi)$ stand

⁷⁵Note that this will basically be a new set of candidate worlds for the actual world.

for whatever updates correspond to ϕ and ψ . *Prima facie*, then, the following seems to be the pattern we are after:

$$(30) \quad \text{ASSERT}(\text{IF}(@E, T_d(\phi), T_d(\psi)); \text{ASSERT}(\text{NOT}(T_d(\psi))); \text{ASSERT}(\text{NOT}(T_d(\phi))))^{76}$$

However, the following is not yet the full logical form of (1)–(3) (indeed, it's not a fully specified logical form at all). First, we cannot know whether (30) has the form of MT or not, until we know what updates $T_d(\phi)$ and $T_d(\psi)$ are. Until this has been specified, the form is incomplete; not every instance of the schema (30) is an instance of MT. As we have seen, the updates that constitute a discourse affect which truth-conditions are expressed by it. Whether or not a discourse that is an instance of (30) will be an instance of MT (partly) depends on whether the truth-conditional content expressed by $T_d(\psi)$ in the context obtained after updating with the antecedent, and the truth-conditional content expressed by $\text{NOT}(T_d(\psi))$ in the context obtained after updating with the big premise, indeed do negate each other.⁷⁷ Only if this is the case we'll have an instance of MT. And whether or not this condition obtains will depend on the respective contexts that determine the truth-conditions expressed by $T_d(\psi)$ and $\text{NOT}(T_d(\psi))$, which in turn will depend on the way the updates that result in these contexts proceed. This is partly a matter of what updates $T_d(\phi)$ and $T_d(\psi)$ are, but it is also a matter of which discourse structuring mechanisms organize the premises and the conclusion into a coherent discourse, since we have seen that those mechanisms also update the context in a way that affects the truth-conditions of modal discourse.

We can then state the following generalization: whenever the truth-conditional content expressed by the small premise negates the one expressed by the consequent of the conditional in the big one, the truth of the big and the small premise together will entail the falsity of the truth-conditional content expressed by the antecedent of the conditional in the big premise. (For a proof of the generalization see B.1.6.1.) Provided that the updates in (30) satisfy the constraint that the truth-conditional content expressed by the utterances they represent indeed conform to the pattern of MT, that is, that the proposition expressed by the small premise negates the proposition expressed

⁷⁶I use the standard notation, representing the sequencing of updates with a semicolon. Thus, where K_1 and K_2 are updates, $K_1; K_2$ is also an update, that performs the update K_1 followed by K_2 (Muskens, 1996).

⁷⁷I say 'partly' because, while validity is primarily a matter of logical form, it is also a semantic notion, capturing certain semantic patterns. As we shall see in what follows, we will be able to fully capture validity as a matter of logical form. Moreover, we will be able to prove that all classically valid patterns are associated with a valid logical form.

by the consequent of the big one, this pattern is valid; whenever the premises describe a possible update, the conclusion does so as well.⁷⁸ Only those fully specified instances of (30) that preserve the adequate form to meaning mapping corresponding to MT will be genuine instances of MT; and all those are valid. This is exactly the same point, as the point that we cannot decide what the full logical form of (12)–(14) is until we know how the different occurrences of the pronoun are resolved. This is hardly a threat to MT.

Note that by characterizing MT in this way, we characterize it as a pattern depending partly on the truth-conditional content expressed, not merely on the underlying syntactic form, because only those ways of specifying (30) that ensure that the truth-conditions of the premises and the conclusion conform to the pattern of MT will count as MT. One might instead maintain that MT should be characterized exclusively in terms of a *unique* syntactic form. Yalcin offers this as an additional argument against MT (Yalcin, 2012b). Namely, the standard Kratzerian “restrictor” analysis of conditionals does not recognize the English conditional as a binary operator, and crucially according to this analysis, what is in the scope of the negation in (2) would not even be a constituent of (1) (Kratzer, 1983); thus, Yalcin argues, since there is no single dyadic operator corresponding to the English language conditional (but rather just a multiplicity of different dyadic operators that correspond to different modals), MT, which he takes to be a generalization about this alleged dyadic operator, makes no sense. Since there’s no “stable notion” of “antecedent” and “consequent”, there is no MT.⁷⁹

I find this line of argument unpersuasive; that a certain syntactic/semantic analysis eliminates a unique binary operator corresponding to the conditional should be independent of the fact that a certain *semantic pattern* is valid. Even if something like the Kratzerian analysis of a conditional turns out to be true, that will hardly constitute a demonstration that MT and MP actually do not exist. Perhaps, we should understand the intuitively valid patterns like MT and MP as precisely the patterns that reflect the behavior of modals in modally subordinated environments, but that does not

⁷⁸Here, I’m appealing to the standard dynamic notion of validity: an inference pattern is dynamically valid just in case the sequential updates with the premises followed by the update with the conclusion lead to a non-empty context. (See § B.1.5 for a precise definition.) What we can prove is that dynamic system embeds classical logic: whenever the truth-conditions associated with the premises classically entail the ones expressed by the conclusion, the inference pattern is dynamically valid. See § B.1.

⁷⁹See (Yalcin, 2012b).

mean that such patterns are merely an illusion. (Perhaps, it is useful to reflect on the fact that though we can formalize propositional calculus by means of, say, Sheffer stroke, it would be odd to argue on that basis that in such a system MT is somehow ill conceived.)

What, then, follows for the alleged counterexample we began with? *Prima facie*, the argument's structure is as follows. Where '*p*' stands for "the marble is big", '*q*' for "the marble is red", and '@*E*' as before denotes the top-ranked epistemically live possibility in the context, a first pass representation of (1)–(3) is as follows:

$$(31) \quad \text{ASSERT}(\text{IF}(@E, \langle \text{comp} := p \rangle, \text{LIKELY}(@E, q)); \text{ASSERT}(\text{NOT}(\text{LIKELY}(@E, q))));$$

$$\text{ASSERT}(\text{NOT}(p)))$$

However, (31) is still not a full-blown logical form of (1)–(3), since it leaves out some of the relevant mechanisms that affect the truth-conditions.⁸⁰ An instance of MT requires that the truth-conditional content expressed by the consequent of the big premise is a negation of the truth-conditional content expressed by the small premise, and whether this is the case, depends on the way the relevant updates affect the context which determines the truth-conditions; in particular, here it depends on the way the anaphoric dependency (i.e., the value of '@*E*') of the modal in the big premise and the small one is resolved. We have seen earlier that discourse coherence plays a crucial role in resolving modal anaphora; thus, in order to determine whether (1)–(3) is an instance of MT or not, we need to take into account the contribution of these mechanisms, which are left out of (31). To get a full-blown logical form (1)–(3) we need to describe the way that mechanisms of discourse coherence update the context, as well.

I argued that mechanisms of discourse coherence change the context by updating the prominence ranking of possibilities that are candidates for anaphora resolution. Furthermore, I argued that (1)–(3) is not an instance of MT, because the Elaboration relation between the antecedent and the consequent in the big premise requires that the modal in the consequent further elaborates on the possibility made prominent by the antecedent, while the Contrast relation between the small premise and the big premise requires that the modal in the small premise quantifies over the body

⁸⁰As noted before, once we have the full-blown logical forms we can restore the idea of validity as a matter of logical form: in particular, a sequence of updates expressing classically valid truth-conditional pattern will be dynamically valid, as well.

of information that both premises are about—i.e., the whole set of epistemically accessible worlds discourse initially. So far, I have specified the updates associated with modals and conditionals. Now we need to capture the effects of the mechanisms of discourse coherence on prominence.

We can capture the effect of these mechanisms on the prominence ranking by representing coherence relations as contributing prominence-affecting updates. Let us first characterize Elaboration. I argued that Elaboration promotes the possibility that is elaborated upon. Here's one way of capturing this idea: when an utterance elaborates on a possibility ϕ , a two-fold contribution is made—first, the possibility elaborated upon, ϕ , is promoted to prominence, pushing all other possibilities one position down in the ranking, and second, it is required that the propositional content expressed by the utterance in question, stands in an elaboration relation to ϕ . We can provisionally characterize an elaboration relation between propositions ϕ and ψ , $Elab(\phi, \psi)$, by requiring that it holds just in case ϕ and ψ are centered around the same event or entity, i.e. just in case the event or scenario described by ψ is a part of the event or scenario described by ϕ .⁸¹ Putting all this together, where ϕ is a possibility, and K an update representing the utterance elaborating on ϕ , we can characterize the update associated with Elaboration as follows:

Definition 3.3.11.

$\llbracket ELAB(\phi, K) \rrbracket(w, G, H)$ iff there are G' and G'' such that $G \approx_0 G' \ \& \ G'_0 = \llbracket \phi \rrbracket^{G, w} \ \& \ \llbracket K \rrbracket(w, G', G'') \ \& \ G'' \approx_0 H \ \& \ H_0 = G''_{comp} \ \& \ Elab(\llbracket \phi \rrbracket^{G, w}, H_0)$.

The contrast relation, as we have seen in § 3.2, has a different effect on the context than Elaboration. Its main effect is the following: the two bits of discourse provide contrasting information about some body of information regarding some common topic. In turn, this body of information is made prominent. Thus, to characterize Contrast formally, we need to have a way of specifying the body of information that a given sentence is about—that is, a body of information it contributes information relative to. Here's one way of doing this. Where ϕ is a formula, we say that ϕ is about a set of worlds θ just in case, where G is an input context to ϕ , $\theta = \llbracket @E \rrbracket^{G, w}$. The idea is just, once more, that as a discourse progresses we are trying to narrow down the space of epistemic possibilities—thus, a

⁸¹Cf. Hobbs (1979), Asher and Lascarides (2003). The provisional characterization suffices, because the exact characterization of Elaboration is not crucial for us; the only thing that matters is the way in which the relation affects prominence. Ditto for other coherence relations.

sentence is just about a currently top-ranked epistemic possibility. I will use ‘ θ_ϕ ’ to denote the set of worlds ϕ is about. Then we can characterize Contrast as follows:

Definition 3.3.12.

$\llbracket \text{CONTRAST}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K_1 \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G''_0 = \llbracket \theta_{K_1} \rrbracket^{G,w} \ \& \ \llbracket K_2 \rrbracket(w, G'', H) \ \& \ \llbracket \theta_{K_1} \rrbracket^{G,w} = \llbracket \theta_{K_2} \rrbracket^{G'',w} \ \& \ \text{Contrast}(G''_{\text{comp}}, H_{\text{comp}})$

According to 3.3.12, when two bits of a discourse contrast with each other, a two-fold contribution is made: first, the body of information they are about is made prominent, and second, the propositions expressed by them are required to stand in Contrast relation ($\text{Contrast}(G''_{\text{comp}}, H_{\text{comp}})$), i.e., to provide contrasting information about this body of information, regarding some common topic.⁸²

Finally, now that we have specified the ways in which prominence ranking changes as the discourse evolves, putting all this together, we can return to the original counterexample. Where p stands for “the marble is big” and q for “the marble is red”, we represent (1) and (2) as follows:

$$(32) \quad \text{CONTRAST}(\text{ASSERT}(\text{IF}(@E, \langle \text{comp} := p \rangle, \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle)))), \text{ASSERT}(\text{NOT}(\text{LIKELY}(@E, \langle \text{comp} := q \rangle))))^{83}$$

The following are the key steps in (32).⁸⁴ By 3.3.8, the conditional update introduces the possibility corresponding to the set of top-ranked epistemically accessible worlds in which the proposition expressed by the antecedent holds, i.e. the set of epistemically live worlds in which the marble is big. The consequent provides an elaboration of this possibility; as a result, this possibility is promoted to prominence (as per 3.3.11). Furthermore, it is required that the possibility introduced by the consequent stands in the Elaboration relation to the possibility introduced by the antecedent, which at this point is the possibility ranked at the position 0 (and, so, denoted by ‘ \mathbf{w}_0 ’ in (32)). Since

⁸²I represent CONTRAST as operating on two updates, and ELABORATION as operating on a proposition and an update. This is in line with a more general distinction between two classes of coherence relations, ones that select their arguments structurally (based on syntactic and structural constraints), and ones that select their arguments anaphorically (Webber et al., 2003). I presuppose this distinction here without defending it, due to considerations of space.

⁸³Note that given this formalization, Contrast will not contribute any asserted content on its own. This is a welcome conclusion, but it is inessential. We could in principle make discourse relations a part of asserted content, by imposing additional constraints on the world of evaluation in the specification of the updates associated with coherence relations.

⁸⁴For a detailed derivation, and a proof that (1)–(3) is not an instance of MT, see § B.1.8.

the consequent contains an occurrence of the modal ‘*likely*’, by 3.3.7, the proposition expressed by the consequent of the given utterance of the conditional corresponds to the proposition that the marble is likely red, given the top-ranked possibility, which due to the effect of Elaboration at this point is the set of epistemically accessible worlds in which the marble is big. Thus, to put it simply, the consequent expresses the proposition that, for all that is known, the marble is likely red, given that it is big. By 3.3.8 again, the whole conditional expresses the proposition that for all that is known, if the marble is big, then it is likely red, given that it is big. By 3.3.9, the assertion update requires that the conditional holds of the actual world and promotes the set of epistemically live worlds in which it holds. Due to the effect of Contrast, as specified in 3.3.12, the body of information that the first utterance is about (which, given that the conditional is uttered discourse-initially, as by assumption it is, is just the set of epistemically accessible worlds discourse initially) is promoted to prominence. Then the modal in the small premise will be interpreted with respect to this body of information—given all that is known, the marble is likely red. By 3.3.10, negation then expresses the complement of this possibility; i.e. the small premise expresses the proposition that it is not the case that given all that is known the marble is likely big. The assertion (by 3.3.9, again) makes sure that this propositional content holds of the actual world, and promotes the set of epistemically accessible worlds in which the content holds. Finally, the propositions expressed by the two premises are required to stand in the Contrast relation—i.e. to provide contrasting information regarding a common topic, about the body of information they are about, i.e. the set of epistemically accessible worlds discourse initially.

Crucially, (32) guarantees that the proposition expressed by the small premise, and the one expressed by the consequent of the big premise do *not* contradict each other. The information expressed is the following: *given all that is known, if the marble is big, then it’s likely red, but, given all that is known, the marble is not likely red*. This pattern does not fit the pattern of the premises of MT. So, *a fortiori*, (1)–(3) is not a counterexample to MT.⁸⁵

⁸⁵Recall, the account maintains that (1)–(3) harbors linguistic elements (modals, antecedents of conditionals, and coherence relations) part of the meaning of which is to change the context in a way that affects the truth-conditions (by introducing certain possibilities, making them prominent, and demoting others). These elements are reflected in the logical form of (1)–(3): it is because (1)–(3) harbors these elements that it is not associated with a valid logical form.

3.3.3 Conclusion to Part II

The formal semantics developed in § 3.3 demonstrates that the account argued for in § 3.2.2–§ 3.2.4 can be given a precise formal characterization. The semantics provided in § 3.3 models the two main aspects of the interpretation of modals and conditionals—the way in which utterances containing these expressions change the context, and the way in which the context determines their truth-conditions. We have seen that both aspects of interpretation are crucial, and that they are interrelated. Unless we captured the change in the context prompted by an update associated with an utterance containing a modal or a conditional, we would not be able to tell how the modal can make a possibility available for subsequent anaphora; and unless we calculated in the anaphoric dependency of utterances containing modals or conditionals, we would not be able to correctly predict which proposition a given utterance containing a modal or a conditional expresses, since the anaphorically retrieved restrictor crucially factors into its truth-conditions. Only once both aspects of the interpretation are taken into account can we give a full explanation of the apparent failure of classically valid patterns of inference.

The formal theory developed in § 3.3 thus provides a precise theory of context-change, and captures the way the mechanisms of discourse structure affect the interpretation of modals. As we have seen, these mechanisms are independently motivated, and their impact is systematic and rule governed; most importantly, the underlying logic is classical. Thus, the semantics achieves our two main goals: it preserves the validity of classically valid patterns, while at the same time explaining away the apparent counterexamples.

Chapter 4

Content in a Dynamic Context

4.1 Introduction

A long-standing tradition in theorizing about meaning, going back at least to Frege (1892, 1918) identifies meaning with propositions, where a proposition is, or at least determines, the truth-conditions of a given sentence in a context. A recent trend, however, urges a departure from this tradition. Certain puzzling behavior of modal vocabulary has drawn critics to propose revisionary, non-truth-conditional accounts, according to which, modal sentences do not express propositions, and hence, do not describe the world in the way ordinary non-modal ones do. I shall call this view *non-propositionalism*. It has received different implementations in expressivist semantics (Moss, 2015; Swanson, 2006; Yalcin, 2007) and certain kinds of dynamic update semantics (Gillies, 2010, 2004; von Fintel and Gillies, 2007; Veltman, 1985).¹ Though these semantics differ importantly, they all deny that modal claims express propositional contents. Instead, they hold that the key to the meaning of a modal is its characteristic effect on a conversational context. Let us say that this characteristic effect constitutes the *dynamic aspect of meaning*. Then, non-propositionalists advocate a reduction of the meaning of modal claims to this dynamic aspect of meaning.²

The departure from tradition is radical. Propositions, or truth-conditions, are traditionally taken to play several fundamental roles. For one, they are taken to be what we assert and believe. When

¹ Another type of non-propositionalism we find in relativist semantics (MacFarlane, 2014; Kolodny and MacFarlane, 2010; Egan, 2007). Insofar as similar data have been invoked to motivate relativism, on the one hand, and expressivism and update semantics, on the other, my case against the latter will carry over against the former as well. I do not discuss relativism directly here, however, since such a semantics, though not propositional, is not best described as non-representational, and it is in principle possible to recover propositional content(s) from relativists' semantic contents. Moreover, such accounts tend to focus less on the dynamic aspect of meaning that expressivists and update semanticists tout.

² In dynamic semantics, the dynamic aspect of meaning goes by the name of *context-change potential*. I stick with the neutral term since, as we shall see, how best to understand this dynamic effect is a matter of some debate among proponents of the revisionary accounts. I shall return to this in what follows.

I say ‘It is raining’, in a normal context I will have expressed that it is raining,³ and my utterance will have been true just in case it is raining. Moreover, if you take me to be sincere and reliable, you will likely come to believe that it is raining. This belief in turn can play an action-guiding role; for instance, you might decide to bring an umbrella. This is because my assertion and your belief represent the world to be a certain way. Denying that modal claims express propositions, in turn, implies that for a bulk of natural language claims assertion involves something different than expressing a proposition. When I say ‘It might be raining’, I do not express the proposition that represents the world to be a certain way, e.g., to be such that in it it might be raining. According to non-propositionalists, rather, I’m merely expressing that I’m in a mental state which does not rule out raining. Even if you take me to be sincere and reliable, and you accept what I said, you will not thereby come to believe some proposition that is representing a world to be such that in it, it might be raining. Rather, you will coordinate your mental state with mine, so that, now, you too are leaving open that the possibility of raining.⁴

Moreover, abandoning propositionalism seems to require a revision of the standard understanding of pragmatic inference. Going back to Grice (1975), the content of an assertion—*what is said* by an utterance—is identified with its truth-conditional content; *what is said*, in turn, is the input to further pragmatic reasoning, playing a role in calculating conversational implicatures. But, if modal claims do not traffic in informational content—if they do not express propositions, or truth-conditions—there is nothing that *is said* by them, in Grice’s sense. Thus, those who deny propositionalism face at least a *prima facie* challenge to explain how such utterances can generate implicatures. For instance, under normal circumstances, if I tell you ‘I might eat some of the cookies’, I implicate that it is not the case that I might eat all of them. While a Gricean account would explain this by appealing to pragmatic inference that begins with the asserted content of my utterance—the proposition that I might eat all of the cookies—a denial of propositionalism would require a modification of this account to accommodate for the fact that the modal claim in question does not express any proposition to begin with.

³Strictly speaking, I will have expressed the proposition that it is raining at the time and location of the context. I shall omit this sensitivity to times and locations of the context for the ease of the exposition, since it won’t matter for our purposes.

⁴See e.g. Yalcin (2007).

Thus, a rejection of propositionalism calls for a rethinking of the nature of meaning, communication, assertion and belief. Of course, non-propositionalists are precisely demanding, indeed, offering, a reconceptualization—this is the surprising philosophical import of their view. It might well be that one is required were it to turn out that non-propositionalism best fits our linguistic practice. But, so I shall argue, this is not the case.

While I concur that an important aspect of modal meaning is its dynamic effect on the context in which it is uttered, I shall argue this is but *one aspect* of modal meaning. In addition, modals express propositional contents. Revisionary accounts have misdiagnosed the nature and the complexity of the characteristic effects of modality on a context of utterance. Once we appreciate the nuanced way in which modals interact with context, we shall find that there is nothing particularly non-standard about modal discourse. In short, I shall argue, although the propositionalist is wrong to neglect the dynamic aspect of meaning, the not-propositionalist is also wrong to reduce modal meaning to this dynamic aspect alone.

I shall proceed as follows. First, in § 4.2, I present key challenges for truth-conditional accounts of meaning. In § 4.3, I sketch how non-propositionalists account for these data. I then go on to argue, in § 4.4, that pessimism about the prospects of a truth-conditional account is a result of an overly simplistic understanding of the interaction between context and content, as well as one of how modals are interpreted in a context. In particular, I argue that the account developed in Stojnić (forthcoming.) naturally explains the data, while maintaining that modal vocabulary has ordinary truth-conditional content. In § 4.5, I argue that this account sheds new light on how we should think about both semantic content, and the interplay between context and content.

4.2 Modality and Propositional Content

The traditional semantics for modal expressions treats them as quantifiers over possible worlds Kratzer (1977, 1981). Of course, this quantification is not unrestricted: a modal does not simply quantify over *all* possible worlds. When I say ‘It might be raining’, though ‘might’ is normally understood as an existential quantifier over possible worlds, I am not saying that there is at least one world out of *all possible worlds* in which it is raining. It is normally assumed that the context of utterance restricts the domain of quantification for a modal. So, the modal is a quantifier over a

contextually restricted domain of possible worlds.

The standard account, developed primarily in Kratzer (1977, 1981), treats modal claims, in a context, as expressing propositions. An utterance of ‘It might rain’ is true in a world and a context, just in case there is at least one world in the (contextually restricted) domain of quantification of the modal ‘*might*’ in which it rains. The utterance expresses the proposition corresponding to the set of worlds in which it is true.⁵

A growing number of authors challenge Kratzer’s account maintaining that modal utterances do not express propositions (Yalcin, 2007; Moss, 2015; Gillies, 2010, 2004; von Fintel and Gillies, 2007; Veltman, 1985). Rather, a modal is used exclusively to convey the speaker’s attitude about the proposition expressed by the prejacent. The main motivation for this shift derives from certain puzzling behavior of modals that seemingly evades standard truth-conditional treatments. To see what the problem is, consider the following example, due to Yalcin (2007):

- (1) If it is not raining and it might be raining, then we won’t get wet.
- (2) If it is not raining and the body of information *i* doesn’t rule raining out, then we won’t get wet.⁶

Since, on the standard account, ‘It might be raining’ means that the (contextually determined) body of information *i* does not rule out raining, the contextualist owes us an explanation of the stark contrast between (1) and (2). But whatever body *i* the context delivers, (2) is coherent, but its counterpart, (1), is not.⁷ Call this the problem of *epistemic contradiction*.

A related problem is illustrated with the following example from Yalcin and Knobe (2013):

⁵I shall adopt the standard terminology and call the proposition expressed by ϕ in an utterance of ‘*might* ϕ ’ the *prejacent* of the modal ‘*might*’.

⁶I modified Yalcin’s example slightly, in a way that makes his point fully general, i.e., not presupposing that the contextually determined body of information is the one corresponding to what is known by the speaker, or perhaps a group of conversational participants including the speaker.

⁷Suppose that the relevant body of information is just the set of worlds compatible with what the speaker knows. Then, clearly, both ‘it is not raining and it might be raining’ and ‘it is not raining and for all I know it is’ seem equally bad. One might try to explain this on pragmatic grounds—the content of such assertions undermines the epistemic grounds for those very assertions. But, as Yalcin (2007) points out, the problem is that only the sentence containing a modal continues to be problematic even when it is not asserted, but rather embedded in an antecedent of a conditional, or under a supposition operator. This suggests that ‘it is not raining and it might be raining’ is not merely pragmatically defective.

Fat Tony.

Fat Tony is a mobster who has secretly planted highly compelling evidence of his murder at the docks. The evidence is discovered by the authorities, and word gets out about his apparent death. The next evening, from his safe house, Fat Tony watches a panel of experts on the news discussing the situation. Expert A has had a good look at the evidence found at the scene. ‘Fat Tony is dead,’ he says. Expert B also had a good look at the evidence, but his assessment is more cautious. ‘Fat Tony might be dead,’ B says.⁸

While the Expert A is clearly wrong—Fat Tony is not dead—it is far less clear what to say about Expert B. On the one hand, there is a sense in which we want to say that what he said is true; after all, Fat Tony planted very compelling evidence. On the other, it seems that Fat Tony can react to Expert B’s utterance with ‘That’s false’; after all, Fat Tony knows that he is alive, and so, that it is false that he might be dead. The standard account has to say that the context determines some body of information *i* that the modal ‘might’ in Expert B’s utterance is quantifying over. But which body could that be? It could be the information compatible with what Expert B knows, or perhaps with what Expert B and his interlocutors know. Or it could be the information compatible with the planted evidence. Or it could be the information compatible with what Fat Tony knows. The problem is the first three options would make Fat Tony’s reaction flat out false—after all, according to Expert B’s knowledge, the knowledge of those that have examined the evidence, and the planted evidence, he might be dead. The fourth option would render Expert B’s utterance flat out false—Fat Tony knows he is not dead. Moreover, it would make it mysterious how Expert B could even assert her utterance in the first place—what grounds does she have to assert a claim about Fat Tony’s body of knowledge?

In fact, Yalcin and Knobe (2013) report on the results of an experimental survey that suggests that in this example, the intuitions about the truth or falsity are unclear: some speakers think that Expert B spoke truly, others that he spoke falsely, and others yet are indecisive. And while subjects

⁸Similar cases are discussed in the literature as the so-called ‘eavesdropper’ cases, because they involve at least one agent that is not a party to the conversation, but merely ‘eavesdropping’. See e.g. Egan (2007). The eavesdropper cases are a particularly challenging instance of the more general problem of disagreement. Cf. MacFarlane (2014, 2011); von Fintel and Gillies (2008, 2009).

do tend to disagree with the claim that what Expert B said is false (the disagreement is just worse than flat out disagreement), they do not tend to endorse that what he said is true either (the agreement is just above flat out ambivalence).

So, cases like **Fat Tony** and (1)–(2) seem to pose a problem for the standard truth-conditional account. These problematic cases show, critics maintain, that there is no systematic way of determining plausible truth-conditional content for a modal utterance in a context. Their diagnosis is simple: the reason why we witness this behavior is because modal utterances do not express such content. Against these criticisms, I shall argue that modal discourse expresses ordinary propositional content. But first, in § 4.3, I sketch the key aspects of the non-propositionalist accounts.

4.3 Modals and Context-Change

Motivated by the data in § 4.2, expressivist and dynamic update semanticists alike reject propositionalism and instead adopt a dynamic approach to meaning. Though there are important differences between the two accounts, the key motivating ideas are similar. Both accounts start from a broadly Stalnakerian model of communication, according to which a conversation takes place against a *common ground*, the set of propositions interlocutors mutually accept for the purposes of the conversation. The common ground in turn determines a context—the set of worlds compatible with all the propositions in the common ground.⁹

The context thus characterizes all the open possibilities with respect to what the world is like, given what we accept for the purposes of a conversation. As the discourse evolves, we acquire more information about the world (or what we take the world to be like for the purposes of our conversation), and as a result, the space of open possibilities shrinks. The effect of an assertion, in particular, is to narrow down the space of possibilities with respect to what the world is like. So, when I assert ϕ against a context c , the effect of my assertion is to rule out worlds from c incompatible with the proposition that ϕ expresses in c ; that is, to add the proposition expressed by ϕ to the common ground, eliminating worlds in which this proposition does not hold from the context. The characteristic effect of an assertion is, thus, an addition of the asserted proposition

⁹Formally, thinking of a proposition as a set of worlds, a context is just a set of worlds in the intersection of all the propositions in the common ground.

to the common ground. If an assertion of p is successful, the interlocutors will coordinate on a particular attitude towards p —e.g., they will come to know, or believe, or accept for the purposes of the conversation that p .

Non-propositionalists maintain that (i) the key aspect of meaning of an utterance is the characteristic effect it has on the common ground, and (ii) the effect of an utterance of a modal claim does not amount to asserting a particular propositional content, and thus, adding a proposition to the common ground. Here is a way to motivate such an account. Given the notion of context we began with, let us say that a sentence is true in a context c just in case an utterance of that sentence does not further change c .¹⁰ This idea is intuitive. Suppose I assert that it is raining, against a context c . If successful, my assertion will add the proposition *that it is raining* to the common ground, in turn eliminating all the non-raining worlds from c ; it will thus become established that it is raining. So, learning (or accepting) that it is raining rules out all non-raining possibilities as the candidates for the actual world. But if there are no non-raining worlds in c to begin with—if it is already a part of the common ground that it is raining—then removing all non-raining worlds from c just leaves c as is. So, if ‘It is raining’ is true in a context c , then asserting that it is raining does not change c .

Given this notion of truth in a context, let us turn to modal claims. Consider what it would take for ‘It might rain’ to be true in a context. Having existential force, ‘It might rain’ requires that there is at least one raining world in the context. Suppose c does not rule out that it is raining. Then, there is no proposition you need to add to the common ground—i.e., no worlds you need to remove from c —so as to make ‘It might be raining’ true in c . It is already true in c just in virtue of c not ruling out raining.¹¹ Similarly, starting with a context c that already rules out raining, there is no proposition you can add to the common ground—worlds you can remove from c —to make it the case that ‘It might be raining’ is true in thusly changed c .¹² So, unlike the effect of a non-modal utterance, which adds a proposition asserted by the utterance to the common ground, the effect of

¹⁰This notion of truth comes from dynamic update semantics Veltman (1985); von Stechow (2007); Gillies (2010). It is closely related to Yalcin (2007) notion of *acceptance*, which we will return to below.

¹¹Note also that the context itself is *not* a set of worlds in which it is true that it might be raining—a set of worlds such that there is some raining world accessible from each of them. So, a context which does not rule out raining does not correspond to the proposition that it might rain, either.

¹²Of course, it is possible to *accommodate*—to change the context in the face of infelicity—by expanding it to include some raining worlds. But this does not amount to adding a proposition, let alone the proposition *that it might rain*, to the common ground. Rather, it involves removing some propositions, at the very least the proposition *that it isn't raining*, from the common ground, thus expanding the context to include some raining worlds.

a *might* claim is to *test* it for compatibility with the prejacent. A successful assertion of *might p* is not to coordinate on a particular attitude towards the proposition that *might p*, for there is none. It is rather to coordinate on a particular attitude towards the prejacent proposition, namely, leaving open that *p*. This is reflected in the common ground—to say that *might p* is true in a context *c* is just to deny that *not p* is a part of the common ground—the context does not rule out *p*. Accordingly, non-propositionalists maintain modal claims do not express propositional content. Instead, they propose that the key aspect of meaning is the effect an utterance has on the context. The ways in which dynamic update semanticists and expressivists cash out this proposal somewhat differ, so it will be useful to sketch and compare them. I start with dynamic update semantics.¹³

Dynamic update semantics proposes we take the effect that an utterance has on a context as its *semantic content*: a sentence *s* is interpreted as a function that takes a context *c* and delivers a context *c'* that results from *c* after *s* is uttered. This function captures the characteristic context-change an utterance of *s* brings about. The dynamic notion of truth corresponds to the one defined above—a sentence is true in *c*, if updating with it does not further change *c*. When an utterance does not contain modal expressions it is easy to extract propositional content from these ‘dynamic meanings’, for reasons we have already seen: updating with ϕ just amounts to adding the proposition expressed by ϕ to the common ground.¹⁴ So, dynamic update semantics for a non-modal fragment is obviously not a radical departure from the standard propositionalist account.¹⁵

¹³I shall omit formal details since the general sketch of the framework will suffice for our purposes. For the formal implementation of dynamic update semantics, see e.g. Veltman (1985); von Stechow and Gillies (2007). For the formal details of the expressivist account, see Yalcin (2007).

¹⁴Formally, this means that the dynamic meanings—updates to the context—for a non-modal fragment are intersective: updating the context *c* with ϕ , where ϕ is a non-modal claim, amounts to intersecting *c* with a proposition expressed by ϕ .

¹⁵That such dynamic update semantics for a non-modal fragment of English reduces to adding propositions to the common ground follows from a formal result due to van Benthem (1986), which establishes that any such reduction holds just in case, for all updates, the following two properties hold:

- *Eliminativity*: An update is eliminative just in case it always proceeds by (at most) eliminating worlds from *c*—starting from a context *c*, an update will always deliver a subset of *c*. Taking a propositional fragment of English, closed under negation and conjunction, where atomic sentences are interpreted as propositional constants, let $[\![\phi]\!]$ be the update function associated with a sentence ϕ , and using standard prefix notation, let us write ‘ $c[\![\phi]\!]$ ’, where the context *c* is an argument to the function $[\![\phi]\!]$. Then eliminativity is defined as follows:

$$c[\![\phi]\!] \subseteq c, \text{ for all } c \text{ and } \phi.$$

- *Distributivity*: For any update, *u* the context obtained by updating *c* with *u* could be obtained by updating singleton subsets of *c* with *u*, and then taking a union of the results. Formally:

$$c[\![\phi]\!] = \bigcup_{w \in c} \{w\}[\![\phi]\!]$$

However, things are radically different when it comes to modal claims. ‘It might be raining’, is interpreted, again, as a function that takes a context and returns a context, but one merely tests the context to check if the prejacent is compatible with it: if it is, the context is returned as is (and hence, the modal claim is true in it); otherwise, the context is simply reduced to an absurd state, \emptyset . What’s noteworthy is that we can no longer extract the propositional meaning of a modal claim from its dynamic meaning. Since dynamic meanings of modals are merely testing whether the context as a whole has a certain property or not, they are not reducible to adding some proposition to the common ground, i.e. to narrowing down the space of open possibilities.¹⁶ So, the dynamic meaning of *might p* is not merely superficially dynamic. This is just to say that a *might* claim does not express an ordinary proposition that gets added to the context by intersection. Now, let us turn to expressivism.

According to expressivism, the goal of a conversation is to coordinate interlocutors’ mental states.¹⁷ The Stalnakerian common ground is a reflection of this coordination. The common ground supervenes on the collective attitudes of the speakers; it is a collection of propositions that the interlocutors mutually believe, know or accept (for the purposes of a conversation).¹⁸ While, as before, an assertion of a non-modal claim is a proposal to add a particular proposition to the common ground—to make it common knowledge that the interlocutors all accept it for the purposes of a conversation—an utterance of a modal claim is not a proposal to add a proposition to the common

-
- An update that is *Eliminative* and *Distributive* is also *Intersective*, which just means that it amounts to intersecting the context with a proposition, i.e. to adding a proposition to the common ground:

$$\text{Intersective Update: } c[\![\phi]\!] = c \cap \phi$$

For more on the relation between van Benthem’s results and its relation to truth-conditionality, see Rotschild and Yalcin (2015).

¹⁶Formally, this is because, due to its test behavior, a modal claim is sensitive to the properties of the context as a whole, where these properties need not carry over to each singleton subset of the context. That is, an update with a *might* claim violates distributivity. To see the failure of distributivity, consider the following counterexample from von Stechow and Gillies (2007). Suppose the context c has only two worlds, one in which the proposition p is true (say, w_1), and one in which it is false (say, w_2). Then, since c contains at least one p -world, the update with *might p* returns c as is, i.e., $\{w_1, w_2\}$. But, if we update each singleton subset individually with *might p*, and then take the aggregate of the results, we will not get $\{w_1, w_2\}$, but $\{w_1\}$; since only $\{w_1\}$ will pass the test, $\bigcup_{w \in c} \{w\}[\![p]\!] = \{w_1\}$. So, the update associated with *might p* is not distributive, and hence, it is not intersective.

¹⁷In the exposition I shall follow Yalcin (2007). For related accounts, see e.g. Swanson (2006); Moss (2015); Charlow (2015); Starr (forthcoming).

¹⁸I am simplifying somewhat. Yalcin (2007, 2012a) argues for a version of expressivism that employs a more complex notion of context, one which comprises probability spaces, rather than worlds. This complexity, which has to do largely with probability modals, need not concern us here.

ground. Rather, it is a proposal to strike a particular attitude towards the *prejacent* proposition: for instance, an utterance of *might p* is a proposal to leave open the possibility that *p*. If the interlocutors successfully coordinate on this attitude—if it is common knowledge that they are all leaving open the possibility that *p*—then the common ground will not rule out *p*; this means that the context will have to have a certain property, namely, being compatible with *p*. This means that the effect of the modal claim again is to test whether the context satisfies a particular property.

While dynamic semantics maintains that the compositional semantic content of a modal claim is the update function characterizing context-change, expressivists maintain it is a property of bodies of information (information states, for short): an utterance of ‘It might be the case that *p*’ expresses the property of being compatible with *p*. To make the key ideas behind expressivism more precise, it is instructive to compare it with the standard truth-conditional account. The standard account treats utterances as expressing ordinary propositional content in a context: ϕ is true in a context *c*, at a world *w* just in case the proposition expressed by ϕ at *c* holds in *w*. Expressivists, in turn, further relativize truth to an information state.¹⁹ For a non-modal language, this relativization has no effect: within a non-modal fragment, a sentence ϕ is true in a context *c*, at a world *w* and an information state *s*, just in case the proposition expressed by ϕ in *c* holds in *w*. The relativization to an information state is idle, and so, the propositional content can be easily recovered. But, once we turn to modals, things change, because the relativization to an information state is no longer idle: ‘*might* ϕ ’ is true in a context *c*, at a world *w* and an information state *s*, just in case there is some world in *s* in which the proposition expressed by ϕ in *c* is true. The information state serves as the domain of quantification for a modal operator. Crucially, expressivists deny that the context of utterance determines this information state. Otherwise, we could easily recover the propositional content by letting the context supply the value of the information state parameter. But the puzzling data in § 4.2 precisely show, expressivists maintain, that the context cannot play a role of supplying an information state in a way that would deliver plausible truth-conditions. So, again, in a context, a modal claim does not express a proposition.

The difference between dynamic update semantics and expressivism lies in the division between semantic and pragmatic labor. The dynamics of context-change for an expressivist is not a matter

¹⁹An information state, like a context, is a set of worlds compatible with certain information. For instance, Jane’s information state is the set of worlds compatible with what she knows or accepts.

of the compositional semantic content of an utterance, but is rather a matter of the pragmatics of conversation. Expressivists hold that an utterance of *might p* invites the audience to leave open the possibility *that p*; they successfully coordinate on this attitude when the context does not rule out *p*. This is just to make the context satisfy the property that the compositional semantics assigns to ‘*might p*’. So, as Yalcin (2007) puts it, in uttering *might p*, the speaker is proposing to make the context satisfy (or make explicit that it satisfies) the property of being compatible with *p*. Adopting expressivist terminology (Yalcin, 2007), let us say that an information state *s* *accepts* a sentence ϕ just in case $\forall w \in s: \llbracket \phi \rrbracket^{c,w,s} = 1$.²⁰ Then an utterance of ‘*might p*’ is a proposal to make ‘*might p*’ accepted in the context, in this sense; i.e. it is the proposal to make the context set satisfy the property that ‘*might p*’ expresses.

In sum, both expressivist and dynamic update semantics treat modals as posing conditions on information states, and both agree that the effect of a modal claim is to test whether the context satisfies this condition. They differ insofar as the dynamic approach takes this characteristic effect of a modal claim on context to be the semantic content of a modal utterance, while the expressivist takes it to be part of the pragmatics of conversation. Crucially, on neither approach is the semantic content of a modal reducible to a proposition. And both approaches maintain that the dynamic meaning is the key aspect of the meaning of modal discourse.

Returning to the contrast between (1) and (2) (repeated below), in order to see how dynamic update semantics and expressivism capture it, we need to say how both accounts interpret conditionals:

- (1) If it is not raining and it might be raining, then we won’t get wet.
- (2) If it is not raining and the body of information *i* doesn’t rule raining out, then we won’t get wet.

On the dynamic update semantics approach, a conditional ‘*if p, q*’ is true in *c*, just in case ‘*q*’ is true in *c* hypothetically updated with *p*. The expressivist account endorses a corresponding non-dynamic interpretation: a conditional hypothetically shifts an information state parameter *s* to the

²⁰ $\llbracket \cdot \rrbracket$ is a function that maps a sentence, context, world and an information state to a truth-value.

maximal non-empty subset of s that accepts the antecedent, and checks whether this state accepts the consequent as well. The conditional is true (relative to a context, a world, and the original information state s) just in case it does.

Given this, the dynamic and expressivist accounts provide very similar explanations of the contrast between (1) and (2). On the dynamic account, hypothetically updating a context c with ‘It is not raining and it might be raining’, requires first updating c with ‘it is not raining’ and then updating the resulting context with ‘it might be raining’. ‘It is not raining’ eliminates all the raining worlds from c . ‘It might be raining’ tests thusly updated context c requiring that it is compatible with it raining. But since there are no remaining raining worlds in c , this test fails, and we get an absurd state \emptyset . The conditional requires us to hypothetically entertain a possibility that does not hold in any non-absurd state. By contrast, there are non-absurd information states in which it holds that it is not raining, but that the relevant body of information i is compatible with raining. Hence, the contrast between (1) and (2).

An expressivist offers a similar explanation. The conditional in (1) requires us to hypothetically shift an information state s to a maximal subset of s that accepts the antecedent—‘It is not raining and it might be’. But there is no (non-empty) state that accepts this. ‘It is not raining’ is accepted at an information state s just in case all the worlds in s are non-raining ones. But, an information state accepts ‘It might be raining’, just in case there is a raining world in s . So, there is no non-absurd state that accepts the conjunction, because the two conjuncts pose contradictory conditions on an information state: there is no non-empty s such that all the worlds in it are non-raining ones, yet it contains a raining world. This is why it is infelicitous to hypothetically suppose that it isn’t raining, but it might be. And this explains why (1) is infelicitous. Meanwhile, there are some non-empty states that accept ‘it is not raining, and the relevant body of information i doesn’t rule out raining’; hence, the contrast with (2).

Both accounts can also explain the mixed intuitions about the **Fat Tony** case. Both claim that a modal utterance expresses a proposal to update an information state in a particular way. Then, Fat Tony’s reaction is merely a refusal to go along with the update proposed by Expert B. Intuitively, Fat Tony fails to coordinate his mental state with Expert B and the panel members, because his information state rules out his death. Hence, he doesn’t endorse the attitude that leaves open the

possibility that he is dead, and hence, he rejects Expert B's assertion. Meanwhile, Expert B's utterance can be accepted in the conversational setting involving her and the expert panel (the one that does not include Fat Tony): if it is mutual knowledge that Expert B and the members of the panel (have come to) have the appropriate attitude, it will become common ground that they are all leaving open the possibility that Fat Tony is dead. Hence, they will endorse the modal claim.²¹

So far, so good for the non-propositionalists. However, as I shall argue, these accounts do not adequately capture the interaction between context and modal vocabulary. A host of other data suggests a more nuanced picture of how modals interact with context. I argue that the behavior of modals in context favors the account developed in Stojnić (forthcoming.), which in turn preserves the idea that modals express ordinary propositional content. I start, in §4.4, by presenting the relevant data and arguing that they motivate a more nuanced account of the interaction between modals and context, than what both the non-propositionalists and even the standard contextualists assume.

4.4 Anaphora and Modality

The standard truth-conditional account subscribes to the following claim:

CONTENT IN CONTEXT: The context of an utterance determines its semantic content, which in turn is, or at least determines, the utterance's truth-conditions.²²

But the type of context-sensitivity a modal exhibits calls for a more nuanced view of how context affects interpretation. I want to suggest that modal expressions are sensitive to the prominence of information states in a discourse—their domain of quantification is restricted by the most prominent body of information. This means that modals exhibit a more flexible type of context-sensitivity than what has been assumed, because prominence is a parameter that changes as the discourse progresses. But, as we shall see, this type of context-sensitivity is also more constrained than what the traditional

²¹For a further discussion, see Yalcin and Knobe (2013).

²²CONTENT IN CONTEXT goes back at least to Kaplan (1989b). Context, on this picture, is thought to comprise all the relevant information for determining the semantic content of context-sensitive expressions. Thus, at the very least, a context has to contain information that helps determine the referent of indexicals (such as 'I', 'you', or 'today'); and given other context sensitive expressions, it might have to include more features, relevant for these other context-sensitive expressions.

model assumes, because the changes in prominence are governed by linguistic rules.

To begin, note that how a modal is interpreted in a context depends often on the prior discourse. This is clearly manifested in classic examples of modal subordination, as in (3), where a modal is interpreted relative to another one introduced earlier in a discourse:

- (3) A wolf might walk in. It would eat you first. Roberts (1989)

Such dependency on prior discourse suggests a picture according to which the context evolves as the discourse progresses. Whatever the context is in which the second sentence in (3) is uttered, it has to keep track of the fact that the previous utterance introduced the possibility of a wolf walking in, because the modal ‘would’ in the second sentence is interpreted relative to this possibility—a *wolf would eat you first, if it were to walk in*. Moreover, a dependent interpretation is possible even within a single sentence, as witnessed by the following examples:

- (4) There *may* be other 1961 state committee retirements come April 18, but they *will* be leaving by choice of the Republican voters.²³ (Stone, 1997)
- (5) If a wolf walks in, it would eat you.

Intuitively, in both (4) and (5), the modal that occurs earlier in the sentence introduces the possibility that affects the interpretation of the one later within the same sentence.

This kind of behavior is not restricted to modals. For instance, pronouns, too, are interpreted with respect to some antecedent either introduced earlier in the discourse, or given by the non-linguistic context, as witnessed by the examples of cross-sentential anaphora, and bound anaphora in (6) and (7), respectively:

- (6) A man walked in. He sat down.
- (7) Every woman believes that she’s happy.

In fact, as argued by Stone (1997, 1999), modals and pronouns exhibit the same range of interpretive

²³I cite the original example that Stone provides, but the effect is easily replicated without ‘will’: “If a concert goer arrives late, he or she might not be permitted into the auditorium.” This holds for the example (9) below, as well.

behavior. In addition to the kind of discourse and operator-bound readings we have seen, both modals and pronouns allow for so-called ‘donkey’ anaphora readings, as witnessed by (8) and (9):

(8) If a man owns a donkey, he beats it.

(9) If a submarine cannot self-destruct if an enemy captures it, the enemy will learn its secrets. (Stone, 1997)

Just as ‘it’ in (8) co-varies with the indefinite noun phrase ‘a donkey’ (without being within its syntactic scope), so in (9) the modal in the consequent, ‘will’, co-varies with the sub-constituent of the antecedent clause ‘if the enemy captures it’. Similarly, just as a pronoun can be used to refer to an individual that is salient within a non-linguistic context, a modal can be interpreted with respect to some possibility that is prominent within a non-linguistic context:

(10) (Referring to a certain significant female) She left me. (Partee, 1984)

(11) (Looking at a high-end stereo in an electronics store) My neighbors would kill me. (Stone, 1997)

And just as with the rest of the modal discourse, conditionals also exhibit anaphoric behavior, as witnessed by the following example from Stojnić (forthcoming.):

(12) If a wolf walks in, it might attack. If John has a tranquilizer gun, he will shoot it.

We said before that modals are naturally thought of as quantifiers over possible worlds. But ‘would’ in (3) is not quantifying over *all* possible worlds, or even all the epistemically accessible worlds. Rather, it is quantifying over those epistemically possible worlds in which a wolf walked in. In other words, the restriction on the domain of quantification of the modal ‘would’ is provided by the prior discourse: the modal ‘might’ in the first sentence introduces the possibility of a wolf walking in, that restricts the domain of quantification for the subsequent ‘would’. Since the domain of quantification is just a set of worlds, we can think of the restrictor as a set of worlds—a proposition (in our example the restrictor is the proposition *that a wolf walks in*); the domain of quantification is then restricted by intersecting it with a restrictor proposition. We can think of the context-sensitivity of a modal

as retrieving the restrictor on the domain of quantification. The examples above show that, as is the case with an antecedent of a pronoun, the restrictor on the domain of quantification of a modal is anaphorically retrieved in a context, in the sense that it is either provided by a prior discourse, or by non-linguistic context.

But neither dynamic update, nor expressivist, semantics captures the anaphoric dependencies of the modals described in (3)–(11). In (3), on the dynamic update semantics, its first sentence merely tests the input context, c , checking whether it is compatible with the prejacent proposition; if the context passes this test, it is passed down as an input context for the second sentence. The second sentence tests this context again, requiring that in all the worlds in the context the wolf eats the addressee. But this delivers the intuitively wrong, unrestricted, reading: (3) expresses that a wolf would eat the addressee, were one to come in, not that it would eat the addressee come what may. Similar problems arise for other examples above, all of which require that modals can introduce propositions into the discourse context, and that the context can keep track of those that were introduced earlier in the discourse, and use them to restrict the domain of quantification of subsequent modals. The dynamic update semantics has no built-in resources that allow for this kind of bookkeeping.²⁴

The same problem befalls the expressivist account. On this account, an utterance of the first sentence in (3) makes (it explicit that) the context leaves open the possibility that a wolf walks in. The second sentence, in turn, makes (it explicit that) the context entails that a wolf eats the addressee. This again delivers an unrestricted reading: the wolf would eat the addressee come what may. The problem is the same as before—modals simply test whether the context satisfies a particular property. The account does not capture the idea that they can introduce propositions, which can in turn restrict the domain of subsequent modals.

²⁴To take one other example, note that this account cannot predict the correct interpretation of examples like the following:

- (i) If a wolf might walk in, it would eat you.

Here's the problem: the conditional is true at c just in case the consequent is true at c updated with the antecedent. Since the antecedent contains a modal, updating the input context with it will merely test the input context c for compatibility with the prejacent; the conditional is true in c just in case the consequent is true in the thusly updated c . This gives the reading that a wolf will eat the addressee, given that one might walk in, whereas intuitively, (i) carries the meaning that a wolf will eat the addressee, given that one walks in.

The standard Kratzerian truth-conditional account does not capture the anaphoricity of modal expressions either. In particular, it does not allow for a modal to select a restriction from the prior discourse context.²⁵ However, the standard account need only be modified slightly to do this. One simple way to modify the standard Kratzerian truth-conditions is proposed in Stojnić (forthcoming.). Following Stojnić (forthcoming.), I use ‘ $M(p, q)$ ’ for the truth-condition expressed by an utterance of *might* ϕ , where q is the proposition expressed by the prejacent ϕ of the utterance of *might* ϕ , and p the proposition corresponding to an anaphorically retrieved restrictor, and I use ‘ R ’ to denote a contextually provided accessibility relation, that determines a set of epistemically accessible worlds from a given world w . Then we can define the truth-condition expressed by an utterance of *might* ϕ , ‘ $M(p, q)$ ’, as follows:

Definition 4.4.1.

$$M(p, q) = \{w \mid \exists w' : wRw' \ \& \ w' \in p \ \& \ w' \in q\}$$

4.4.1 says that an utterance of ‘*might* ϕ ’ is true at w relative to a context c if and only if there are worlds R -accessible from w , in which both the anaphorically retrieved restrictor, p , and the prejacent proposition, q , are true. The accessibility relation R provides the domain of quantification, while p is the anaphorically retrieved contextual restriction on this domain.²⁶ Crucially, we factor in the contribution of the restrictor separately, which will allow us to track how it is recovered in a context. (As is standard, ‘*must*’ is the universal dual of ‘*might*’.)

The following modification of the standard Kratzerian (1983) truth-conditions for a conditional, proposed by Stojnić (forthcoming.), captures in turn the anaphoricity of a conditional:²⁷

Definition 4.4.2.

²⁵That the standard Kratzerian semantics for modals does not accommodate this anaphoric behavior has been forcefully argued by Stone (1997). As Stone (1997) points out, the context-sensitivity of modals on Kratzer’s semantics is factored in via complex contextual parameters—a modal base and an ordering source—which are both determined in complex ways, and neither provides a semantic parameter that can be contributed by prior discourse, except in the special case when the antecedent of a conditional restricts the modal in the consequent. See also the discussion in (Kratzer, 2012, ch 4), and Stone (1997); Brasoveanu (2010); Bittner (2014).

²⁶There is still a question of what determines the relevant set of epistemically accessible worlds in a given context, which is here delivered by the R relation. We will turn to this question presently.

²⁷For simplicity, I shall ignore the ordering source—the contextually supplied parameter that ranks worlds relative to their proximity to the world of evaluation, given some standard—but we could easily factor this parameter in, and then state the truth-conditions by making a conditional true (in a world w and at a context c) just in case all the p and q worlds that are closest to w in the domain of quantification are r worlds as well.

$$\text{Cond}(p, q, r) := \{w \mid \forall w' : wRw', \text{ if } w' \in p \ \& \ w' \in q, \text{ then } w' \in r\}$$

where p , as before, is the anaphorically retrieved restrictor, q corresponds to the proposition expressed by the antecedent, and r to the one expressed by the consequent, an utterance of a conditional expresses truth-conditions corresponding to a set of worlds such that for each w in the set, all the worlds R -accessible from w , that are p and q worlds, are r worlds as well; i.e. an utterance of a conditional is true in w if and only if all the p and q worlds in the domain of quantification are r worlds as well. This preserves the standard truth-conditions associated with a conditional, but factors in that a conditional itself is always evaluated against some prominent body of information, that need not correspond to the unrestricted set of epistemically live worlds.

So, while expressivist, and dynamic update, semantics have trouble capturing the fact that modals can select a restrictor from a prior discourse, the truth-conditional account in 4.4.1 and 4.4.2 captures this by design. But we have seen that the truth-conditional account, at least of the standard sort, faces the challenge of explaining the data in § 4.2. I suggest, however, that the truth-conditional account of the form in 4.4.1 and 4.4.2, once complemented by an account of the modal anaphora resolution, accounts for the data straightforwardly. Modals are, just as the standard account would have it, context-sensitive expressions which select a prominent body of information that restricts their domain of quantification. The problem with (1) is not that there is no body of information that the context can select that yields the plausible interpretation; rather, I shall argue, it is that the context determines the body of information that delivers an inconsistent interpretation. As we shall see, the resolution of context-sensitivity of modal expressions is a rule-governed process. The rules that govern the resolution deliver the inconsistent interpretation in (1), but not in (2). This is why we witness the contrast. To illustrate what I have in mind, consider the following example:

(13) John came to the party, but he didn't come.

(14) John came to the party, but the contextually prominent male didn't come.

While (13) is infelicitous, (14) is perfectly fine. But everyone agrees that 'he' is context sensitive, and that in a context it refers to some contextually prominent male.²⁸ I suggest that the source of

²⁸Note that 'John came to the party, but the person I intend to refer to didn't', or 'John came to the party but the contextually determined male didn't', are all likewise perfectly consistent.

contrast in (13) and (14) is exactly analogous to the one (1) and (2). What (13) and (14) show is not that ‘he’ isn’t context-sensitive, or that (13) doesn’t express truth-conditions, but rather, as I shall argue, that there are rules governing the resolution of context-sensitivity, that determine the inconsistent interpretation in (13) and (1), but not in (14) and (2).

As witnessed by (3)–(5), and (9), an utterance containing a modal, or a conditional, does not passively receive its content from the context—rather, it induces an active change in the context, in a way that directly affects the interpretation of subsequent discourse, e.g., by supplying a potential restrictor for a subsequent modal. Moreover, as examples like (4) and (5) show, certain expressions (like modals, or antecedents of conditionals) bring about changes in a context mid-utterance as well. Indeed, it is part of the meaning of an ‘*if*’-clause, or a modal, that it raises a certain possibility to prominence.²⁹ These changes in contextual prominence, as we have seen, affect the resolution of modal anaphora.

Moreover, as I have argued in Stojnić (forthcoming.), not only do modal expressions actively induce changes in context, but our linguistic practices provide a wide range of systematic, robust mechanisms that make propositions prominent in a context, in a way that systematically affects the resolution of modal anaphora. In particular, drawing on resources from the Discourse Coherence Theory (Asher and Lascarides, 2003; Kehler, 2002; Hobbs, 1979), I have argued that the resolution of modal anaphora is sensitive to mechanisms of *discourse coherence*. The key idea of this approach is reflected in a simple but often neglected observation that a discourse is more than a random sequence of sentences. To flesh this out, Coherence Theorists postulate an implicit organization of discourse that establishes inferential connections—*coherence relations*—among utterances (Kehler, 2002; Asher and Lascarides, 2003). This implicit organization arises from the communicative strategies that interlocutors exploit to convey and organize their ideas through an ongoing discourse. For example, consider (15):

(15) Phil tickled Stanley. Liz poked him. Smyth (1994)

²⁹That introducing a hypothetical scenario is a part of the meaning of an ‘*if*’-clause is reflected in the Kratzerian restrictor analysis of the conditional, according to which the sole meaning of an ‘*if*’-clause is to restrict the domain of an (overt or cover) modal expression as well as in the update and expressivist accounts of conditionals, which posit that an ‘*if*’-clause, as a matter of compositional semantics, introduces a hypothetical, local context, in which the consequent is evaluated. However, as we have seen above, these accounts are nevertheless not expressive enough to fully capture the ways in which conditionals and modals make propositions available for subsequent modal discourse.

There are (at least) two ways we can understand (15). The second clause could be taken to describe the result of the event described by the first: Phil tickled Stanley, and *as a result*, Liz poked him (perhaps she was acting in disapproval); or, one could understand the two clauses as comparing and contrasting two parallel events: Phil tickled Stanley, *and similarly* Liz poked him (Liz's action is comparable to Phil's). Coherence theorists capture this difference by positing that, on the one reading, the discourse (15) is organized around the Result relation, which belongs to a broader class of cause-effect (or event-result) relations, while, on the other, it is organized around the Parallel relation, which is an instance of a broader class of resemblance relations. Crucially, the choice of the relation affects the resolution of the pronoun: if the discourse is understood as organized by the Result relation, the pronoun 'him' refers to Phil; if it is organized around Parallel, then 'he' is Stanley. The choice of a coherence relation guides pronoun resolution.

A number of empirical studies corroborate the observation that pronoun resolution co-varies with the choice of coherence relation.³⁰ Moreover, the data suggest that the mutual constraints between these two processes are both systematic and robust: given the choice of a coherence relation, the interlocutors are radically constrained in the possible interpretation of a pronoun. The following example from Kehler (2002) illustrates the point:

- (16) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

(16) is generally judged infelicitous. We expect the pronoun in the second sentence to resolve to Reagan, and intuitively feel the speaker has erred in uttering 'her' instead of 'him'. This infelicity is explained if the choice of coherence relation forces a particular resolution of the pronoun. The discourse is organized by the coherence relation of Parallel, signaled by the scalar relationship between 'admires' and 'absolutely worships'. Parallel requires that the occurrence of a pronoun in the object position be resolved to the object of the first clause, which is Reagan. But since the pronoun is feminine, this results in a gender mismatch, which explains why the utterance is judged infelicitous. Note, it is surprising that out of the blue (16) receives this infelicitous interpretation, given the available referent for 'her' in the first conjunct, which is a well-known subject of Bush's

³⁰See e.g. Wolf, Gibson, and Desmet (2004), Kehler et al. (2008), and Keiser (2009).

admiration. If the correlation were a matter of general pragmatic reasoning, the perceived infelicity of (16) would be mysterious. Given that the infelicitous reading is the one naturally retrieved, this suggests that the effect of a coherence relation on the interpretation of a pronoun is a robust, conventionalized effect.³¹ Coherence relations make certain referents prominent for subsequent anaphora.

In Stojnić (forthcoming.), I argued that the mechanisms of discourse coherence likewise affect the prominence of propositions that are potential restrictors of subsequent modals. We see this already in (3), where the second sentence elaborates on the hypothetical possibility described by the first. I have argued that it is an effect of this Elaboration relation that ‘would’ in the second sentence is understood as restricted by the possibility described by ‘might’ in the first. Elaboration makes the hypothetical scenario in which a wolf walks in, introduced by the modal in the first sentence, prominent, and in turn, ‘would’ in the second sentence is understood as quantifying over the epistemically accessible worlds in which this scenario holds. Similarly, an Elaboration relation between the antecedent and the consequent in (5) promotes the proposition describing the hypothetical scenario introduced by the antecedent, and consequently ‘would’ in the consequent is understood as restricted by this proposition.

Crucially, the mere fact that one modal follows another does not suffice to establish that the second modal is further restricted by the proposition introduced by the first one—the import of coherence is crucial. This is illustrated with examples like (17):³²

- (17) If a wolf walks in, it would eat you. But, one probably won’t walk in.

As before, the Elaboration relation between the antecedent and the consequent of the first sentence in (17) makes the hypothetical scenario introduced by the antecedent the most prominent one, and as a result, ‘would’ in the consequent further describes this scenario. But, the modal ‘won’t’ in the second sentence does not further elaborate upon the scenario described by the two modals in the

³¹For detailed defense of the claim that coherence relations are conventional devices that determine the resolution of a pronoun, see Stojnić, Stone, and Lepore (2013, 2014). One thing Stojnić, Stone, and Lepore (2014) argue is that languages differ with respect to the effects of coherence relations on the interpretation of pronouns, which suggests that, indeed, the effect is a matter of linguistic convention.

³²For similar examples, see Asher and McCready (2007).

first sentence. Rather, the two sentences stand in a relation of Contrast, signaled by the discourse marker ‘but’. They are understood as contrasting two hypothetical scenarios—one in which a wolf walks in, and one in which one does not.

The Contrast relation requires that the two bits of discourse provide contrasting information about some information state regarding a common topic, here, what is possible with respect to a wolf’s entrance.³³ The first sentence sets the stage in determining the body of information the contrast has to be about—(assuming that the conditional is uttered discourse initially) it is interpreted relative to the set of epistemically accessible worlds determined by the context discourse initially, describing what might be the case if a wolf walked in, *given this overall body of knowledge*. The second sentence, then, has to provide a contrasting bit of information about this body of information available discourse initially, regarding the possibility of a wolf’s entrance. The Contrast relation makes this body of information prominent, and consequently, the modal in the second sentence selects it as its restrictor—*given this overall body of knowledge*, a wolf probably won’t come in. This is the intuitively correct interpretation: *given the overall body of knowledge*, *if a wolf walks in*, it would eat the addressee, but *given the same body of knowledge*, one probably won’t walk in.³⁴

These data show, at the very least, that there are systematic mechanisms by which a context changes, not only between utterances, but also within a single utterance. The systematic context-change in turn affects the proposition expressed by subsequent utterances. This suggests that CONTENT IN CONTEXT is false—it is only once we account for the systematic context-change, both within a single utterance, and between utterances, that we can predict the intuitive truth-conditions. Suppose we endorse the idea that modals express truth-conditional content. Then we can capture the idea that the restriction on a domain of quantification of a modal is determined through systematic mechanisms of context-change described above in the following way: as proposed in Stojnić (forthcoming.), suppose that the context represents a ranking of propositions that are candidates for domain restrictors of modals in a discourse according to their relative prominence—the most prominent being the top-ranked one. Then, let the modal retrieve the most prominent proposition representing an epistemically live possibility as the restrictor for its domain of quantification. As the

³³One common way of signaling what the topic is, in English, is by exploiting prosodic accents. (See Rooth (1992), Roberts (1996a)).

³⁴See Stojnić (forthcoming.) for a more detailed exposition, and further examples.

discourse evolves, the prominence ranking changes: modals and conditionals can introduce propositions, and similarly, coherence relations make certain possibilities prominent for subsequent modal anaphora. Importantly, we have seen, the changes in prominence are governed by systematic mechanisms of context-change: by linguistic items that introduce new propositions, and by mechanisms of discourse coherence that promote certain propositions, while demoting others. This allows us to capture the anaphoric behavior of modals manifested in the examples above.

It turns out, given this, we also have resources to explain the puzzling contrast between (1)–(2), repeated below as (18) and (19):

(18) If it is not raining and it might be raining, then we won't get wet.

(19) If it is not raining and the body of information *i* doesn't rule raining out, then we won't get wet.

Let us begin with (18). Naturally, the second conjunct, in the antecedent of the conditional in (18), 'It might be raining,' is understood as elaborating on the hypothetical raining scenario introduced by the first one (compare, e.g., 'If the door is open, and a wolf might walk in, then we are in trouble'). This affects the resolution of modal anaphora in a similar way as before: the Elaboration relation between the two conjuncts makes the proposition comprising the epistemically accessible worlds in which it is not raining prominent. Consequently, '*might*' in the second conjunct in the antecedent selects this proposition as its restrictor, and is thus understood as quantifying over all the relevant epistemically accessible worlds in which it is not raining. So, we get the interpretation according to which the antecedent is expressing the proposition *that it is not raining, and within the epistemically accessible worlds in which it is **not** raining, there is at least one raining world*. This, of course, leads to a contradiction, so it is not a surprise that the conditional winds up being bad. Meanwhile, since the anaphoric dependency of '*might*' in (18) is absent in (19), we have an explanation of the discrepancy between the two examples. The proposition *that it is not raining, and moreover the body *i* doesn't rule out raining* is not a contradiction.³⁵

³⁵Note that my account does not predict that the conditional with the reversed order of conjuncts in the antecedent, i.e. of the form: 'If it might *p*, and not *p*, then...', will automatically be bad. This is a desired result since, as Dorr and Hawthorne (2014) note, switching the order of the conjuncts in some cases (in particular, in Yalcin's original example) makes the conditional felicitous. However, my account does not predict that all such cases will automatically be felicitous

What about the second type of puzzling data that the opponents of truth-conditional accounts point to—examples like **Fat Tony**? It turns out that the anaphoricity of modals easily predicts this kind of behavior as well. Remember, on this account, modals are looking for a prominent proposition that restricts their domain of quantification. We have seen that a proposition can be made prominent by explicitly mentioning it. But, as witnessed by (11), a proposition need not be explicitly mentioned in order for it to be prominent. With this in mind, here is how we can explain the **Fat Tony** example. One can understand Expert B’s comment as elaborating on the proposition corresponding to the body of information of those who have examined the evidence, or perhaps as elaborating on the body of information of those who’ve engaged with the question of whether Fat Tony might be alive. An utterance, indeed, might be ambiguous between these different interpretations. Depending on which proposition is elaborated upon, different propositions are made prominent—on one disambiguation, it would be inappropriate for Fat Tony to react with ‘That’s false’, while on another, this might be a perfectly appropriate reaction.³⁶ This is all perfectly compatible with (and predicted by) an account according to which modals express ordinary propositional content.³⁷

either. This is because the badness of such a conditional will depend on which coherence relations, and other interpretive dependences can be established between the two conjuncts in the antecedent, and between the conditional and the rest of the discourse in which it is embedded, for these factors can all affect the resolution of modal anaphora in a particular case. This is, again, a desired result. And the point holds more generally—embedding any sentence in a broader discourse might give rise to various interpretive dependences that might yield a different interpretation than the one we get when interpreting the same sentence-type out of the blue.

³⁶Dowell (2011) and von Stechow (2009) recognize that the reason why we get puzzling disagreement data of the sorts we see in **Fat Tony** is because there might be more than one body of information that the modal can be interpreted as quantifying over. However, while they recognize the need for a more flexible story about how context impacts the interpretation of an epistemic modal, the accounts they offer are different than the one I advocate here in terms of how this observation is cashed out. Moreover, these accounts do not give us a story about the embedded occurrences of modals.

³⁷A word of clarification is in order. In 4.4.1 and 4.4.2, I assumed that an epistemic accessibility relation R , which provides the domain of quantification for epistemic modals, is provided by the context discourse initially, and further restricted by an anaphorically retrieved restrictor (p), the most prominent proposition in the given context. We could assume that this set of epistemically accessible worlds discourse initially is at the same time the top-ranked proposition discourse initially. But, of course, part of the problem is determining what R is discourse initially and this is precisely the problem stressed by the puzzling disagreement data, and examples like **Fat Tony**. I have suggested that even here, the modal is looking for some prominent body of information, as witnessed by (11), and as I argue is the case with **Fat Tony**. We can easily modify our truth-conditions to accommodate the fact that modals can be restricted by a proposition that hasn’t been explicitly mentioned, but which determines the relevant set of epistemically accessible worlds discourse initially. Let’s define a class of expressions, ‘ P_w ’, interpreted as follows: $P_w = \{w' | w' \in p\}$, where p is some proposition. Now, let p be the proposition comprising the most prominent body of knowledge relative to a world w , and $P_w = \{w' | w' \in p\}$ (we relativize knowledge to worlds in this way, because what, e.g., Tony knows in w might be different from what he knows in w'); moreover, let q be the proposition expressed by a prejacent ϕ of ‘*might* ϕ ’, and $Q_w = \{w' | w' \in q\}$. Then, we define the truth-conditions for ‘*might* ϕ ’ be as follows:

Definition 4.4.3.

$$M(p, q) := \{w \mid \exists w' \in P_w \ \& \ w' \in Q_w\}$$

Note, that this behavior is something we see with other context-sensitive expressions as well. Suppose you are at a dinner party in a restaurant, and everyone invited seems to be having fun. You utter (20):

(20) Everyone is having a great time.

The waiter, who is working his third shift that day, overhears this from a distance, and reacts with (21):

(21) Well, that's false. I'm having a terrible day over here.

It seems that these examples give rise to the same intuitions as **Fat Tony**. There is a sense in which the waiter's utterance could be interpreted as true, and then one in which it is understood as false. Just as in the case of modals, this has to do with how we resolve the restriction on the domain of the quantifier. I suggest that the two case are exactly alike in this respect.

So, pessimism towards truth-conditional accounts rests on an overly simplistic conception of how context interacts with content, one that presupposes **CONTENT IN CONTEXT**. We have seen, however, that **CONTENT IN CONTEXT** is false: context changes with evolving discourse, both between utterances and within a single utterance. Moreover, the mechanisms by which the context changes are both systematic and robust. This allows us to meet the key challenges from non-truth-conditional accounts. But, as I shall argue in the next section, there is still something right about the non-propositionalist accounts.

This captures the idea that a modal can be restricted by some proposition prominent in the non-linguistic context. *Mutatis mutandis*, we make the same change for the conditional. Where p and P_w are defined as in 4.4.3, q is the antecedent proposition, $Q = \{w' | w' \in q\}$, r is the consequent proposition, and $R = \{w' | w' \in q\}$, the truth-condition for a conditional is defined as follows:

Definition 4.4.4.

$Cond(p, q, r) := \{w \mid \forall w' \text{ if } w' \in P_w \ \& \ w' \in Q_w \text{ then } w' \in R_w\}$

For the most part, we can disregard this modification, and work with the definitions in 4.4.1 and 4.4.2.

4.5 Content in a Dynamic Context

My account maintains that modals express truth-conditional content. We have seen that the mechanisms affecting the resolution of context-sensitivity are systematic and robust; certain expressions—e.g. modals, and antecedents of conditionals—systematically affect context in a way that affects the truth-conditions expressed. Moreover, we have seen that mechanisms structuring the information flow in a discourse—in particular, mechanisms of discourse coherence—systematically affect context in way that also affect the interpretation of modal discourse. I want to suggest that these mechanisms contribute another layer of semantic content—over and above truth-conditions. In particular, I shall argue that these mechanisms contribute the dynamic aspect of meaning, capturing the change in context, brought about by an utterance. Thus, in this, and only in this, regard I agree with the non-propositionalist—a dynamic effect on context is an important aspect of meaning. But it is only one aspect of meaning. Once we correctly characterize the dynamics of context-change, the truth-conditions naturally fall out of this characterization.

The account I sketched above suggests that the resolution of context-sensitivity is sensitive to the changes in context induced by certain linguistic expressions (e.g., modals, antecedents of conditionals), and other linguistic items (e.g. coherence relations). The following picture of context-content interaction emerges. Suppose we think of a context as a conversational record, in the sense of Lewis (1979), an abstract ‘scoreboard’ that tracks the moves and contributions interlocutors make in the flow of a discourse, and that comprises information relevant for interpretation, such as who’s speaking, what the conversation is about, etc. Among other things, a conversational record tracks propositions put into play in the course of a conversation, as well as their relative prominence. As the discourse progresses, the state of the record changes. Interlocutors make different conversational moves, reflected on the scoreboard. New questions are raised, others answered; new propositions are made prominent, others demoted.

Utterances can change the record in a way that affects the truth-conditions expressed. For instance, we have seen that they can introduce propositions that can serve as restrictors for subsequent modals or conditionals. How do we capture the fact that utterances induce these kinds of systematic changes in context? For instance, how do we capture that modal expressions introduce possibilities

that can serve as restrictors for subsequent anaphora? In addition to expressing propositional content, I suggest that we interpret utterances as updates to context, i.e. as instructions on how to change the state of the scoreboard. To make this more concrete, let us focus again on modal anaphora. The key aspect of a conversational context that matters for the resolution of modal anaphora is the relative prominence of propositions that are candidate restrictors for a modal. Since we are only dealing with modal anaphora, we can abstract away from all other aspects of the conversational record needed for other purposes.

Suppose then that a context is a ranking of propositions relative to their prominence—those higher in the ranking being more prominent than those lower.³⁸ We can see this as modeling one aspect of a Lewisian scoreboard—the relative prominence of propositions within a discourse. At the beginning of a conversation, the top-ranked proposition is the set of epistemically accessible worlds determined by the initial context—it might be the body of knowledge compatible with the common ground, or some other body of knowledge that is relevant given our present practical purposes.³⁹ We interpret utterances as instructions to update the context, promoting certain propositions, and demoting others.⁴⁰ So, an utterance of ‘might p ’ is not merely associated with truth-conditional content, defined in 4.4.1; it is also interpreted as an update to a conversational context the effect of which is to promote a new proposition—the proposition corresponding to the top-ranked set of epistemically accessible worlds in which the prejacent holds.

To illustrate, take the first sentence of (3). Suppose that (3) is uttered out of the blue, and that the relevant body of information that it is about is some set of epistemically accessible worlds, which comprises a proposition p . Given that p is the top ranked proposition in the context, the modal selects it as its restrictor: so, the utterance expresses the proposition that given p , the possibility of a wolf walking in is open. But the utterance also makes the proposition comprising all the epistemically accessible p -worlds in which a wolf walks in prominent, demoting all others one position in the ranking. Thus, apart from merely expressing a propositional content, a modal claim also affects the

³⁸Here, I borrow resources from centering theories (see Sidner (1983), Grosz, Joshi, and Weinstein (1995), and Bittner (2014)).

³⁹Recall the **Fat Tony** example: depending on which body of information the speaker is understood as commenting on, different bodies of information might be made prominent.

⁴⁰I undertook a formal development of this account in Stojnić (forthcoming.). Here we need not concern ourselves with formal details—the general sketch of the framework will suffice for our purposes.

conversational record, by updating the prominence ranking. It introduces a proposition that a conversational record can track across the discourse, so that it can serve as a restrictor for subsequent modals.⁴¹

Apart from modals and conditionals, we have seen, the mechanisms of discourse coherence also affect context, by changing the prominence ranking. In (3), since the second sentence is elaborating on the hypothetical scenario introduced by the first one—the epistemic possibility of a wolf walking in—Elaboration makes this proposition the new top-ranked one (demoting all others one position down). Thus, we interpret coherence relations as making a two-fold contribution: first, they contribute updates to the conversational context (e.g., Elaboration promotes the possibility elaborated upon) and second, they signal a particular relation between the relevant bits of discourse (e.g., it is signaled that the proposition expressed by the utterance of the second sentence in (3) stands in an Elaboration relation with the epistemic possibility of a wolf walking in).⁴² In other words, apart from signaling that a particular relation holds between different bits of discourse, coherence relations also contribute updates to the context that affect prominence ranking.

My account thus treats utterances as having two levels of content. One is characterized by the updates, the role of which is to structure and rearrange the context and contextual parameters as the discourse evolves. This kind of content encodes the characteristic effect an utterance has on the conversational record. The other is ordinary propositional content. The account is designed to capture both aspects of modal meaning at the same time. Both are necessary and are related. Unless we factor in the context-change that an utterance induces, we wouldn't be able to predict how the utterance sets the context for subsequent discourse, e.g., by making a proposition prominent for subsequent modal anaphora. This context-change, in turn, affects the truth-conditions—in particular,

⁴¹Plausibly, an utterance of a declarative sentence also makes the proposition expressed by the utterance prominent. For example, an utterance of 'It is raining,' promotes the proposition that it is raining, which is witnessed by the fact that we can easily pick it up as a propositional anaphor, e.g. by following up with 'That is true'. If that is so, then the modal claim in (3) first makes the proposition comprising all the p -worlds in which a wolf walks in prominent, demoting all others one position down, and then it makes prominent the proposition that given p , a possibility of a wolf walking in is open. Note that, as a consequence, both propositions are stored on the conversational record, and potentially available as restrictors for subsequent modals. Since the second sentence in (3) elaborates on the epistemic possibility of a wolf walking in, introduced by the first, as a result of the Elaboration relation, this proposition is again promoted, and thus, the modal in the second sentence selects it as a restrictor.

⁴²We characterize an Elaboration relation between propositions p and q , by requiring that it holds just in case p and q are centered around the same event or entity, i.e. just in case the event or scenario described by q is a part of the one described by p (cf. Hobbs (1979), Asher and Lascarides (2003)). This provisional characterization suffices; the only thing that matters is the way in which the relation affects prominence.

given the anaphoric behavior of modals, the proposition expressed by a modal utterance depends on which propositions were made prominent by the prior discourse. The account is thus modeling a two-way interaction between a context and an utterance—on the one hand, an utterance changes the context, and on the other, it depends on the context to determine the proposition it expresses. The changes to the context that an utterance induces, in turn, affect the propositional content of the subsequent discourse. Far from being incompatible with the propositional content, the dynamic meaning is precisely what allows us to systematically compute it.⁴³

Note that my account treats the dynamics of prominence as governed by linguistic mechanisms, building into the semantics proper the mechanisms affecting the prominence ranking. Couldn't we, instead, keep the anaphoric account of modals and conditionals, as specified in 4.4.1 and 4.4.2, but leave the mechanisms structuring the context (and thus governing anaphora resolution) out of the semantics? In particular, couldn't the dynamics of prominence be governed by speaker's intentions and general pragmatic cues? Nothing in principle rules out such an account. But the account wouldn't satisfactorily explain our key datum, namely, the contrast between (1) and (2).

Recall, the Elaboration relation between the conjuncts in the antecedent of (1) makes the possibility described by the first conjunct prominent, and that is why the modal in the second conjunct selects it as its restrictor. This in turn delivers an inconsistent interpretation. If we thought that the effect on prominence associated with the Elaboration were merely a byproduct of pragmatic reasoning, we would expect the interlocutors to re-interpret, selecting some other body of information as the restrictor for the modal (perhaps the body of information of the speaker, delivering the unrestricted reading: 'if it's not raining, and for all I know it is...'). Instead, we recover the inconsistent interpretation. Other data, too, suggest that the effect of discourse coherence on the interpretation of modal anaphora is a matter of linguistic rules, rather than a byproduct of pragmatic reasoning. It seems that hearers follow a narrow set of linguistic cues that signal how the discourse is organized when resolving anaphoric dependencies. We have seen this already with pronominal anaphora in (16) and with modal anaphora in (1). To take another example, consider the contrast between (22) and (23):

⁴³As shown in Stojnić (forthcoming.), this approach allows us to characterize the truth-conditions in a way that preserves the ordinary classical notion of validity and entailment.

(22) John might come to the party, and he probably won't.

(23) John might come to the party, but he probably won't.

While (22) seems odd, (23) is perfectly fine. The only difference between the two is that (23) is organized around the Contrast relation, signaled by 'but', while in (22) we naturally expect Elaboration. We have seen before, in § 4.4, that the effect of Contrast is to promote the body of information that the two bits of discourse in (23) are providing contrasting information about, which yields the intuitive interpretation of (23)—given the information available discourse initially, John might come to the party, but given this same information, he probably won't. On the other hand, the expected effect of Elaboration would deliver an inconsistent interpretation for (22)—it is compatible with the information available discourse initially that John will come to the party, but he will probably not come, *given that he comes*. If this effect of Elaboration were a mere byproduct of pragmatic reasoning, there would be no reason, in principle, why the hearers would not be selecting the body of information available discourse initially as the restrictor for the modal 'won't' in (22), which would essentially yield the interpretation we get in (23). Moreover, note that if one were to re-interpret (22) in this way when faced with inconsistency, one could no longer understand (22) as organized by the Elaboration relation—one cannot freely change the resolution of modal anaphora while holding the relation fixed. This again suggests that the effect of the relation is linguistically encoded.

Further support for my linguistic account comes from the examples like the following:

(24) A wolf might walk in. But, then again one will probably not walk in. It would eat Harvey.

In (24), we cannot understand 'would' in the final sentence as restricted by the proposition describing the epistemic possibility of a wolf walking in, introduced by the modal in the first. While this is expected on a linguistic account, which maintains that the resolution of modal anaphora is governed by the underlying discourse structure, it is less clear why there would be such a constraint on a purely pragmatic account—after all, there is a prominent proposition in the context that could serve as the restrictor, and which would yield a plausible interpretation. In other words, the possible resolutions of a modal anaphora are much more constrained than what one would expect if they

were determined by pragmatic principles.

Just as the robustness of the effect of coherence on prominence seems to be linguistically constrained, so too does the effect of modal expressions. It is a part of the meaning of a modal (or an antecedent of a conditional) that it makes the proposition introduced by the preadjacent prominent for subsequent modal anaphora. Note that a proposition is made prominent regardless of whether the modal is embedded, say under a negation, as in (26) and (27), or in a question, as in (25), or an antecedent of a conditional, as it is in (1):

(25) Might a wolf walk in? It would probably eat Harvey.

(26) No wolf will come in. It would have to get through the security gate.

(27) John will probably not come to the party. That's a pity. He would have fun.

In all these examples, the dynamics of prominence is in play. If the effect on prominence were a matter of a pragmatic effect of an utterance of a modal, it would be surprising that they persist even under embeddings, and even when the possibility that is made prominent is questioned (as in (25)), or explicitly said not to obtain (as in (26) or (27)).⁴⁴

In addition, just as languages differ with respect to the effects of coherence relations on the interpretation of pronouns, which suggests that, indeed, the effect is a matter of linguistic convention (Stojnić, Stone, and Lepore, 2014), there is likewise cross-linguistic variation in how languages exploit coherence in signaling shifts in prominence that affect modal anaphora. For example, the data presented by Asher and McCready (2006) show that the direct translation of (3) in Japanese is infelicitous, unless there is an overt discourse marker signaling an Elaboration relation.⁴⁵ This, again, would be surprising if the interpretation according to which the second modal is understood as elaborating on the scenario described by the first one, were merely a byproduct of general reasoning.

These data suggest, then, that not only are the mechanisms of context-change that affect the

⁴⁴Moreover, the effect is present even when the possibilities made prominent are extremely far-fetched, which, too, would be somewhat surprising on a pragmatic view, since general reasoning often allows us to disregard far-fetched scenarios.

⁴⁵Asher and McCready (2006) are not themselves committed to the view that coherence relations contribute prominence-affecting updates, nor do they use these data to argue that coherence relations have conventionalized effects on the resolution of modal anaphora. Still their data illustrates this point nicely.

resolution of modal anaphora systematic and robust, but they are also part of grammar. This, in turn, suggests that not only do we need to take into account mechanisms of context-change in order to derive the plausible truth-conditions for modal utterances, but moreover, language encodes the effects of these mechanisms along (and beyond) ordinary truth-conditional content. This is precisely what the account which interprets utterances as carrying both dynamic and propositional content maintains.⁴⁶

4.6 Conclusion

Taking stock, I have argued against the non-propositionalists. In particular, I have argued that their argument from the puzzling behavior of modal vocabulary in context does not warrant the conclusion that modal claims fail to express propositional content. The data is better explained once we recognize that modals are anaphoric expressions, and that the resolution of modal anaphora is sensitive to systematic and robust mechanisms of discourse structure and discourse coherence. Once we take these mechanisms into account, we can easily explain the puzzling behavior of modals that fueled non-propositionalist accounts, while maintaining that modals express ordinary propositional content.

The view I am defending maintains that context systematically changes with the evolving discourse, both between utterances, and also within a single utterance. Thus, it rejects the traditional picture in *CONTENT IN CONTEXT*. Hence, though modal claims express ordinary truth-conditions, there is an important sense in which my view is dynamic: much like the non-propositionalist accounts, it too crucially relies on the idea that an important aspect of meaning is the characteristic effect an utterance has on context, and this aspect of meaning goes beyond the truth-conditional content. However, unlike these other accounts, I have argued that once we adequately capture the dynamic aspect of meaning, the ordinary propositional meaning naturally falls out.

More precisely, I have defended a theory according to which utterances are interpreted as carrying two types of content—they are interpreted as instructions to update existing dependencies in a context, and they express propositional content. Both types are needed, and they are interrelated: the

⁴⁶See the discussion in Lepore and Stone (2015) and Stone and Stojnić (2015) for arguments that discourse relations themselves are learned conventions for signaling how utterances are organized into a coherent discourse.

kind of context-change that an utterance induces on a context directly affects the truth-conditional content of the subsequent discourse. Thus, it turns out, though dynamic semantics is often understood as competing with traditional truth-conditional semantics, the two are in sync.

Appendices

Appendix A

Formal Definitions for the Attention-Coherence Approach

A.1 Appendix: Formal definitions

We conclude with complete definitions for a logical language that formalizes our Attention–Coherence approach to pronouns. We assume a set of individual constants \mathcal{C} ; a set of predicate symbols \mathcal{P} , each taking a specified number of arguments; and variables x_i for each natural number i . The interpretation of constants is set up in terms of frames and models in the usual way.

- **A Frame** is a tuple $F = \langle D_w, R, D_e, D_t \rangle$ where D_w is a domain of possible worlds, R is a (transitive and reflexive) accessibility relation on D_w , D_e is a domain of individuals, and D_t is a domain of truth values ($D_t = \{0, 1\}$). We require that the domains be disjoint: $D_t \cap D_w = D_w \cap D_e = D_t \cap D_e = \emptyset$.
- **A Model** is a pair $\mathcal{M} = \langle \mathcal{F}, I \rangle$, where \mathcal{F} is a frame and I is an interpretation function, which assigns to each individual constant an element of \mathcal{D}_e and each n -place predicate constant a set of pairs $\langle w, \sigma \rangle$ with $w \in \mathcal{D}_w$, and σ an n -tuple of elements of \mathcal{D}_e .

The Attention–Coherence approach captures interpretive dependencies across formulas using dynamic semantics. Operations on sequences of individuals play a key role in the semantics. Given the individuals defined by a Frame F , we use the following notation to specify these operations.

- i_m
If m is an integer, then i_m is the m th element of i .
- $i_{m..n}$
If m and n are integers, then $i_{m..n}$ is a sequence containing the subsequence of elements of i in order from element number m up through the element that precedes n (if any).

- $i_{m...}$

If m is an integer, then $i_{m...}$ is the sequence containing the complete subsequence of elements of i in order beginning from element number m .

- $i + j$

If i is a sequence and j is a sequence, then $i + j$ is the sequence containing the elements of i in order followed by the elements of j in order.

Note then that $i = i_{0,k} + i_{k...}$

- $u.i$

If $u \in D_t \cup D_w \cup D_e$ is an individual and i is a sequence, then $u.i$ is the sequence that begins with u and continues with the elements of i in order.

- $w(i)$

If i is a sequence, then $w(i)$ is the first element u of i such that $u \in D_w$.

Now we can define the expressions of the language and their interpretations in a model \mathcal{M} .

- Individual expressions

- if t is an individual constant, then t is an individual expression
(represents the name of an individual)
- the variable x_m is an individual expression
(represents a discourse reference contributed by argument structure)
- if p is a unary predicate, then $@p$ is an individual expression
(represents a syntactically unconstrained anaphor)
- if p is a unary predicate and o is an individual expression, then $@p^o$ is an individual expression
(represents a syntactically constrained anaphor)

The interpretation of individual expressions at a sequence i and world w :

- $\llbracket t \rrbracket i, w = I(t)$ for interpretation function I .
(Access constants from model.)

- $\llbracket x_m \rrbracket i, w = i_m$.

(Look up values of variables. We need variables to manage argument structure; otherwise, it will be very cumbersome to deal with the syntax–semantics interface for transitive and ditransitive verbs; we need to potentially distinguish the order in which arguments are introduced, how salient they are after the utterance, and what role they play in the described event. Having variables clears this all up. Basically, x_0 will correspond to the subject, x_1 to the direct object, x_2 to the indirect object, and so forth.)

- $\llbracket @p \rrbracket i, w = i_0$ if $\langle w, i_0 \rangle \in I(p)$.

$\llbracket @p \rrbracket i, w = \llbracket @p \rrbracket i_{1\dots}, w$ otherwise.

(Find most prominent referent that agrees with anaphor.)

- $\llbracket @p' \rrbracket i, w = i_0$ if $\langle w, i_0 \rangle \in I(p)$ and $i_0 \neq \llbracket t \rrbracket i, w$.

$\llbracket @p' \rrbracket i, w = \llbracket @p' \rrbracket i_{1\dots}, w$ otherwise.

(Find most prominent free referent that agrees with anaphor.)

Conditions:

- If r is an n -place predicate symbol and t_1 through t_n are individual expressions, then $r(t_1, \dots, t_n)$ is a condition.

The interpretation of conditions:

- $\llbracket r(t_1, \dots, t_n) \rrbracket i, w$ is true if and only if $\langle w, \langle \llbracket t_1 \rrbracket i, w, \dots, \llbracket t_n \rrbracket i, w \rangle \rangle \in I(r)$

(interpret atomic conditions by making sure the specified entities are in the specified relation at the world of evaluation)

Dynamic updates:

- $\langle \alpha k \rangle$ is an update, for $k \in \mathbb{N}$.

(push new indefinite assignment for variable x_k)

- $\langle \pi k t \rangle$ is an update, where t is an individual expression, and $k \in \mathbb{N}$.

(push new assignment of t as a value of a variable x_k)

- $[\varphi]$ is an update if φ is a condition.
(restrict the values of variables)
- $H;K$ is an update if H and K are updates
(composition—conjunction)
- $\Box K$ is an update if K is an update
(metaphysical necessity)

The interpretation of dynamic updates: At each possible world w , the interpretation of a dynamic update is a relation on sequences:

- $\llbracket \langle \alpha k \rangle \rrbracket(w, i, j)$ if and only if $j = i_{0,k} + o.i_{k...}$ for some individual $o \in D_e$.
- $\llbracket \langle \pi kt \rangle \rrbracket(w, i, j)$ if and only if $o = \llbracket t \rrbracket i, w$ and $j = i_{0,k} + o.i_{k...}$ for some individual o of type e .
- $\llbracket [\varphi] \rrbracket(w, i, j)$ if and only if $j = i$ and $\llbracket \varphi \rrbracket i, w$ is true.
- $\llbracket H;K \rrbracket(w, i, j)$ if and only if there is some sequence h such that $\llbracket H \rrbracket(w, i, h)$ and $\llbracket K \rrbracket(w, h, j)$.
- $\llbracket \Box K \rrbracket(w, i, j)$ if and only if $j = i$ and for all worlds v accessible from w , there is some k such that $\llbracket K \rrbracket(v, i, k)$.

Following Kaplan, we can define an initial context of a model M as any sequence (a, x, y) where $a \in D_w$ (representing the actual world of the context), $x \in D_e$ (representing the speaker of the context) and $y \in D_e$ (representing the addressee of the context). (In a more general language, this could be extended by whatever parameters are appropriate for the interpretation of relevant indexical elements.) Then we can define truth in a context and a model as abstractions over the basic dynamic updates:

- H is true in a model M and initial context J for M , if and only if there is some sequence i such that $\llbracket H \rrbracket(a, J, i)$.
- H is valid if and only if H is true in every model M for every initial context for M .
- H entails_s K if and only if for any model M and initial context J for M , if H is true at M and J , then K is true at M and J .

This is the “static” sense of entailment: K is a summary of H .

- H entails_D K if and only if for any model M and initial context J for M , and any assignment i such that $\llbracket H \rrbracket(a, J, i)$, there is an assignment k such that $\llbracket K \rrbracket(a, i, k)$.

This is the “dynamic” sense of entailment: K doesn’t add information to H . This version of entailment is the one that’s usually given in treatments of dynamic semantics designed to accommodate anaphora, because it allows anaphoric links not only between the premises but from the premises to the conclusion.

Worked out examples:

- A man met Sam. He greeted him.

| formula | gloss | output |
|--|--------------------------------------|---|
| $\langle \alpha 0 \rangle; [\mathbf{man}(x_0)];$ | “A man (is the subject)” | (m, \dots) where m is a man |
| $\langle \pi 1 s \rangle;$ | “Sam (is the object)” | (m, s, \dots) |
| $[\mathbf{met}(x_0, x_1)];$ | “(the subject) met (the object)” | (m, s, \dots) where m met s |
| $\langle \pi 0 @ \mathbf{he} \rangle;$ | “He (is the subject)” | (m, m, s, \dots) since m is a he |
| $\langle \pi 1 @ \mathbf{he}^{x_0} \rangle;$ | “him (is the object)” | (m, s, m, s, \dots) since m, s are he but $m = x_0$ |
| $[\mathbf{greeted}(x_0, x_1)]$ | “(the subject) greeted (the object)” | (m, s, m, s, \dots) where m greeted s |

Recall that you should be reading “0” as “the subject” and “1” as “the direct object,” which is their meaning in the formalism.

- John was disappointed with Tim. He fired him.

| formula | gloss | output |
|---|---|--|
| $\langle \pi 0 \mathbf{j} \rangle$ | “John (is the subject)” | (j, \dots) |
| $\langle \pi 1 \mathbf{t} \rangle$; | “Tim (is the object)” | (j, t, \dots) |
| $[\mathbf{disapp.with}(x_0, x_1)]$; | “(the subject) was disappointed with (the object)” | (j, t, \dots) where j was disappointed with t |
| $[\mathbf{Narration}(x_0)]$; $\langle \pi 0 x_0 \rangle$; | “continue a narrative about the subject” | (j, j, t, \dots) where j was disappointed by t |
| $\langle \pi 0 @ \mathbf{he} \rangle$; | “He (is the subject)” | (j, j, j, t, \dots) since j is a he |
| $\langle \pi 1 @ \mathbf{he}^{x_0} \rangle$; | “him (is the object)” | (j, t, j, j, t, \dots) since j, t are he , but $j = x_0$ |
| $[\mathbf{fired}(x_0, x_1)]$ | “(the subject) fired (the object)” | (j, t, j, j, t, \dots) where j fired t |

– John was disappointed with Tim. He did sloppy work.

| formula | gloss | output |
|--|---|---|
| $\langle \pi 0 \mathbf{j} \rangle$ | “John (is the subject)” | (j, \dots) |
| $\langle \pi 1 \mathbf{t} \rangle$; | “Tim (is the object)” | (j, t, \dots) |
| $[\mathbf{disapp.with}(x_0, x_1)]$; | “(the subject) was disappointed with (the object)” | (j, t, \dots) where j was disappointed with t |
| $[\mathbf{Explanation}(x_0, x_1)]$; $\langle \pi 0 x_1 \rangle$; | “explain appealing to the object” | (t, j, t, \dots) where j was disappointed by t |
| $\langle \pi 0 @ \mathbf{he} \rangle$; | “He (is the subject)” | (t, t, j, t, \dots) since t is a he |
| $\langle \alpha 1 \rangle$; $[\mathbf{work}(x_1)]$; | “A work (is the object)” | (t, w, t, j, t, \dots) where w is some work |
| $[\mathbf{sloppy}(x_1)]$; | (the object) is sloppy | (t, w, t, j, t, \dots) where w is sloppy |
| $[\mathbf{did}(x_0, x_1)]$ | “(the subject) did (the object)” | (t, w, t, j, t, \dots) where j did t |

We can reduce the difference in the minimal pair to an attentional shift associated with two different coherence relations, Narration and Explanation, represented explicitly in logical form.

Appendix B

Language for Modality with Coherence

B.1 Appendix: Formal definitions

In this appendix, I first provide a dynamic language that formalizes my approach to modality with anaphora (§ B.1.1–§ B.1.5), and then prove that the system preserves the validity of MP and MT (§ B.1.6). Then I go on to provide a translation from a fragment of English into this language and into a simple classical language and prove that our dynamic interpretation preserves classical interpretation (§ B.1.7–B.1.7.4). Finally, I prove that Yalcin’s counterexample is not an instance of MT (§ B.1.8).

B.1.1 Syntax:

In this section I specify the expressions of the language. We first start by listing the basic vocabulary:

- Propositional expressions: the elements of the set \mathcal{C} of constants (p, q, r, \dots), and the elements of the set \mathcal{V} of variables ($comp$ and \mathbf{w}_n for $n \in \mathbb{N}$).
- Unary predicates: P, Q, R
- Unary operator: $@$
- Update expressions: K, H
- Connectives: \wedge, \neg
- Identity: $=$

The following are atomic formulae (atoms) in our language:

- Propositional expressions are atoms.

- $@P$ is an atom, where P is a unary predicate.
- Nothing else is an atom in our language.

These are the conditions in our language:

- All atoms are conditions.
- $\phi = \psi$ is a condition, where ϕ, ψ are conditions. (Stands for identity.)
- $\neg\phi$, where ϕ is a condition. (Stands for negation.)
- $\phi \wedge \psi$ is a condition where ϕ, ψ are conditions. (Stands for conjunction.)

These are the update expressions:

- $\langle comp := \phi \rangle$ is an update expression, where ϕ is an atom.
- If ϕ is a condition, then $[\phi]$ is an update expression.
- $K; K'$ is an update expression, if K is an update expression and K' is an update expression.
- $MIGHT(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $MUST(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $LIKELY(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $IF(\phi, K_1, K_2)$ is an update expression, if ϕ is a condition and K_1 and K_2 are update expressions.
- $AND(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.
- $NOT(K)$ is an update expression, if K an update expression.
- $ASSERT(K)$ is an update expression, if K is an update expression.
- $ELAB(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $CONTRAST(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.

B.1.2 Models:

I define frames and models in the usual way:

- **A Frame** is a tuple $\mathcal{F} = \langle \mathcal{D}_w, \mathcal{D}_t = \{0, 1\}, R, \mathcal{P} \rangle$ such that \mathcal{D}_t is a domain of truth values ($\mathcal{D}_t = \{0, 1\}$), \mathcal{D}_w is a finite domain of possible worlds, $\mathcal{D}_t \cap \mathcal{D}_w = \emptyset$, with R , a (transitive and reflexive) accessibility relation defined over \mathcal{D}_w , and \mathcal{P} , a probability measure over \mathcal{D}_w , that maps each subset of \mathcal{D}_w to $[0, 1]$, satisfying the following constraints:
 - i $\mathcal{P}(\mathcal{D}_w) = 1$
 - ii $\mathcal{P}(p \cup q) = \mathcal{P}(p) + \mathcal{P}(q)$, when p and q are disjoint subsets of \mathcal{D}_w .
 - iii \mathcal{P} is a regular probability measure: if $p \neq \emptyset$ then $\mathcal{P}(p) > 0$.
- **A Model** is a pair $\mathcal{M} = \langle \mathcal{F}, I \rangle$, where \mathcal{F} is a frame and I an interpretation function, which assigns to each propositional constant p a subset of \mathcal{D}_w and each predicate constant P a set of subsets of \mathcal{D}_w .

B.1.2.1 Truth-conditional contributions of modals and conditionals:

Let us define some meta-language abbreviations that will help us state the truth-conditions associated with updates associated with modals and conditionals. These correspond to propositions expressed by modals and conditionals.

- Where R is the accessibility relation, and \mathcal{P} the probability measure over sets of possible worlds provided by the model:

Definition B.1.1. (Definition 3.3.2 in the text)

$M(p, q) := \{w \mid \exists w' : wRw' \ \& \ w' \in p \ \& \ w' \in q\}$ —*might* q , relative to some possibility, p .

Definition B.1.2.

$N(p, q) := \{w \mid \forall w' : wRw', \text{ if } w' \in p \text{ then } w' \in q\}$ —*must* q , relative to some possibility p .

Definition B.1.3. (Definition 3.3.3 in the text)

$P(p, q) := \{w \mid \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p \ \& \ w' \in q\}) / \mathcal{P}(\{w' \mid wRw' \ \& \ w' \in p\}) > .5\}$ —*probably* q , given p .

Definition B.1.4. (Definition 3.3.4 in the text)

$Cond(p, q, r) := M(p \& q, r) = \{w \mid \forall w' : wRw', \text{ if } w' \in p \& w' \in q, \text{ then } w' \in r\}$ —if q, r , relative to p .

B.1.3 Describing operations on stacks (sequences of worlds) and sets of stacks.

Here I define operations on stacks and sets of stacks, which I will use to define the semantics for our language later on. Formally, a stack is just a function from a finite convex subset of \mathbb{N} plus *comp* to a set of worlds plus \perp , where ' \perp ' denotes an undefined value.¹ (I'll assume that '*comp*' is a designated position on the stack. Where the stack is intended to model prominence ranking, '*comp*' is not affecting the prominence ranking, as described in § 3.3.)

- Where $m \in \mathbb{N}$, and i is a stack, i_m is the m^{th} member of the stack if m is within the domain of i , and $i_m = \perp$ otherwise. (i_{comp} is the member of the stack stored at the designated position *comp*.)—Selecting a member of the stack.
 - Where G is a set of stacks (i.e. a 'context'), g a stack, and u a world, $G_m = \bigcup_{g \in G} \{u \mid g_m \neq \perp \& g_m = u\}$, for $m \in \mathbb{N}$ or $m = comp$.—Getting the m^{th} element in the set of stacks G .
- For $m, n \in \mathbb{N}$, and a stack i , $i_{m,n}$ is a stack j defined on the set $\{0, \dots, n-m\} \cup \{comp\}$ such that for $k \in \mathbb{N}$, $j_k = i_{(m+k)}$ if j is defined on k , and $j_{comp} = i_{comp}$.
 - Where G is a context, and g and j are stacks, $G_{m,n} = \bigcup_{g \in G} \{j \mid j = g_{m,n}\}$ and for $H = G_{m,n}, H_{comp} = G_{comp}$.
- For $m \in \mathbb{N}$, and a stack i , $i_{m...}$ is the stack j defined on the set $\{k \in \mathbb{N} \mid i \text{ is defined at } (m+k)\} \cup \{comp\}$ such that, for $k \in \mathbb{N}$, $j_k = i_{(m+k)}$ and $j_{comp} = i_{comp}$.
 - Where G is a context, and g, j are stacks, $G_{m...} = \bigcup_{g \in G} \{j \mid j = g_{m...}\}$ and for $H = G_{m...}, H_{comp} = G_{comp}$.
- If i is a stack with a finite domain with maximal element $k-1$ then for a stack j , $i+j$ is a stack h where, for $x \in \mathbb{N}$, $h_x = i_x$ if i is defined at x , and $h_x = j_{(x-k)}$ otherwise (and $h_{comp} = i_{comp}$).

¹A set of numbers S is convex just in case if $x \in S, y \in S$ and $x < m < y$ then m is in S .

- Where u is a world and i is a stack, u, i is a stack j , such that $j_0 = u$, and for all $n \in \mathbb{N}$, such that $n > 0$, $j_n = i_{(n-1)}$ if i is defined on n , and $j_n = \perp$ otherwise and $j_{comp} = i_{comp}$.—Appending to a stack.
 - Where G is a context, u is a world, and g, j are stacks, $G_{u...} = \bigcup_{g \in G} \{j | j = u, g\}$ and for $H = G_{u...}, H_{comp} = G_{comp}$.
- $g[n]g'$ iff $g_m = g'_m$ for $m \neq n$ (where $m, n \in \mathbb{N} \cup \{comp\}$).
- $G \sim_n G'$ iff $\{g' | g[n]g', g \in G\} = \{g' | g[n]g', g \in G'\}$ (where $n \in \mathbb{N} \cup \{comp\}$).
- $G \approx_n G'$ iff $\{g_{0,n} + g_{n+1...} | g \in G'\} = G$ and $G_{comp} = G'_{comp}$.

B.1.4 Semantics:

The Interpretation of Atoms: The interpretation of an expression e , relative to the interpretation function I a context G , and a world w :

- $\llbracket p \rrbracket^{G,w} = I(p)$, if $p \in \mathcal{C}$.
 - Constants.
- $\llbracket \mathbf{w}_m \rrbracket^{G,w} = G_m$, if $\mathbf{w}_m \in \mathcal{V}$ and $m \in \mathbb{N}$
 - Variables.
- $\llbracket comp \rrbracket^{G,w} = G_{comp}$
 - A designated position on the stack.
- $\llbracket @P \rrbracket^{G,w} = \emptyset$ if $G_0 = \perp$, $\llbracket @P \rrbracket^{G,w} = G_0$, if $G_0 \in I(P)$, and $\llbracket @P \rrbracket^{G,w} = \llbracket @P \rrbracket^{G_{1...},w}$ otherwise.
 - Find the top ranked entity in G , satisfying P .

The Interpretation of Conditions:

- $\llbracket \phi = \psi \rrbracket^{G,w} = \mathcal{D}_\omega$, if $\llbracket \phi \rrbracket^{G,w} = \llbracket \psi \rrbracket^{G,w}$; $\llbracket \phi = \psi \rrbracket^{G,w} = \emptyset$, otherwise.
 - Identity.

- $\llbracket \neg\phi \rrbracket^{G,w} = \mathcal{D}_\omega \setminus \llbracket \phi \rrbracket^{G,w}$.
 - Negation.
- $\llbracket \phi \wedge \psi \rrbracket^{G,w} = \llbracket \phi \rrbracket^{G,w} \cap \llbracket \psi \rrbracket^{G,w}$.
 - Conjunction.

The Interpretation of Update Expressions

- $\llbracket \langle comp := p \rangle \rrbracket(w, G, H)$ iff $G \underset{comp}{\sim} H$ & $H_{comp} = \llbracket p \rrbracket^{G,w}$
- $\llbracket [\phi] \rrbracket(w, G, H)$ if and only if $H = G$ and $w \in \llbracket \phi \rrbracket^{G,w}$
- $\llbracket [K; K'] \rrbracket(w, G, H)$ iff $\exists G' : \llbracket [K] \rrbracket(w, G, G')$ and $\llbracket [K'] \rrbracket(w, G', H)$
- The following are updates that describe how propositional content (B.1.2.1) in context is determined. Where p is a proposition (an anaphorically retrieved restrictor) and ‘@E’ denotes the top-ranked proposition that is the subset of the epistemically accessible worlds:²
- $\llbracket \text{MIGHT}(\phi, K) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket [K] \rrbracket(w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \text{ \& } G'' \underset{comp}{\sim} H$ & $H_{comp} = M(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{MUST}(\phi, K) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket [K] \rrbracket(w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \text{ \& } G'' \underset{comp}{\sim} H$ & $H_{comp} = N(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{LIKELY}(\phi, K) \rrbracket(w, G, H)$ if and only if there is a G' and G'' such that $\llbracket [K] \rrbracket(w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \text{ \& } G'' \underset{comp}{\sim} H$ & $H_{comp} = P(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{IF}(\phi, K_1, K_2) \rrbracket(w, G, H)$ iff there is a G', G'', G''' and G'''' such that $\llbracket [K_1] \rrbracket(w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \text{ \& } \llbracket [K_2] \rrbracket(w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \text{ \& } G'''' \underset{comp}{\sim} H$ & $H_{comp} = \text{Cond}(\llbracket \phi \rrbracket^{G,w}, G'_{comp}, G'''_{comp})$
- $\llbracket \text{AND}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G', G'', G''' and G'''' such that $\llbracket [K_1] \rrbracket(w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \text{ \& } \llbracket [K_2] \rrbracket(w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \text{ \& } G'''' \underset{arg0}{\sim} H$ & $H_{comp} = G'_{comp} \cap G''_{comp}$

²For generality, I let the restrictor in the definition be any proposition p . However, as argued above, epistemic modals and conditionals select the top-ranked possibility in a given context (‘@E’) as their restrictor.

- $\llbracket \text{NOT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \underset{comp}{\sim} H \ \& \ H_{comp} = \llbracket \neg comp \rrbracket^{G', w}$
- $\llbracket \text{ASSERT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \ \& \ G' \underset{0}{\approx} H \ \& \ H_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G, w} \ \& \ w \in H_0$

In order to define the truth-conditions for updates associated with coherence relations, we assume the following abbreviations:

Definition B.1.5. $\text{Elab}(\phi, \psi)$ iff ϕ and ψ are centered around the same event or entity, i.e. iff the event or scenario described by ψ is a part of the event or scenario described by ϕ .

Definition B.1.6. A formula, ϕ , is about of body of information θ iff, where G is the input context to ϕ , $\theta = \llbracket @E \rrbracket^{G, w}$, where ‘ E ’ is a predicate denoting the property of being an epistemically accessible proposition, and thus, ‘ $@E$ ’ denotes the top-ranked epistemically accessible proposition. I use ‘ θ_ϕ ’ to denote the body of information that ϕ is about.

Definition B.1.7. $\text{Contrast}(\phi, \psi)$ iff ϕ and ψ describe contrasting information about some body of information regarding a common topic.

- $\llbracket \text{ELAB}(\phi, K) \rrbracket(w, G, H)$ iff there are G' and G'' such that $G \underset{0}{\approx} G' \ \& \ G'_0 = \llbracket \phi \rrbracket^{G, w} \ \& \ \llbracket K \rrbracket(w, G', G'') \ \& \ G'' \underset{0}{\approx} H \ \& \ H_0 = G''_{comp} \ \& \ \text{Elab}(\llbracket \phi \rrbracket^{G, w}, H_0)$.
- $\llbracket \text{CONTRAST}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K_1 \rrbracket(w, G, G') \ \& \ G' \underset{0}{\approx} G'' \ \& \ G''_0 = \llbracket \theta_{K_1} \rrbracket^{G, w} \ \& \ \llbracket K_2 \rrbracket(w, G'', H) \ \& \ \llbracket \theta_{K_1} \rrbracket^{G, w} = \llbracket \theta_{K_2} \rrbracket^{G'', w} \ \& \ \text{Contrast}(G''_{comp}, H_{comp})$

B.1.5 Truth, validity, entailment.

- K is true, relative to a context G , a world w , and a model \mathcal{M} , if there is some H , s.t. $H \neq \emptyset$ and $\llbracket K \rrbracket(w, G, H)$. K is false (relative to a context G , a world w , and a model \mathcal{M}), otherwise.
- K is valid iff it's true in all models.
- K_1 entails K_2 iff for any model \mathcal{M} , any context G , and any world w if there is a G' such that $G' \neq \emptyset$ and $\llbracket K_1 \rrbracket(w, G, G')$, then there is a G'' such that $G'' \neq \emptyset$ and $\llbracket K_2 \rrbracket(w, G', G'')$.

B.1.6 Modus Ponens (MP) and Modus Tollens (MT):

We want to prove that MP and MT are preserved within our system. In order to do that, we need to have a way of characterizing MT and MP correctly.

B.1.6.1 MT

Let us start with MT. We need to have a way of individuating instances of MT, first. Let us begin with the following discourse:

1. $\text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{Assert}(\text{NOT}(K_3))$

- where ‘@E’ denotes the top-ranked epistemically accessible proposition.

What we want to show is that when the propositions expressed by K_2 and $\text{NOT}(K_3)$ in their respective contexts contradict each other, then if the proposition expressed by $\text{IF}(@E, K_1, K_2)$ and $\text{NOT}(K_3)$ both hold (relative to the world of evaluation w), then the proposition corresponding to the intersection of the truth-conditional contribution of K_1 and $\llbracket @E \rrbracket^{G,w}$, where G is the input context, will be false (relative to the world of evaluation w). (Note that for our proof it does not really matter what ‘@E’ is; the proof goes through regardless of what we take the anaphoric restrictor to be. I use ‘@E’ because, as I argue in the paper, epistemic modals and conditionals select top-ranked epistemic possibility in the context as a restrictor; but this is orthogonal to our proof.)

Proof.

2. $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{ASSERT}(\text{NOT}(K_3)) \rrbracket(w, G, H)$ iff there is a G' , such that $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)) \rrbracket(w, G, G')$ and $\llbracket \text{ASSERT}(\text{NOT}(K_3)) \rrbracket(w, G', H)$.
3. Take such a G, G' and H .
4. $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)) \rrbracket(w, G, G')$ iff there is a G'' , such that $\llbracket \text{IF}(@E, K_1, K_2) \rrbracket(w, G, G'') \ \& \ G'' \approx_0 G' \ \& \ G'_0 = G''_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ w \in G'_0$.
5. Take such a G'' .

6. $\llbracket \text{IF}(@E, K_1, K_2) \rrbracket(w, G, G'')$ iff H', H'', H''' and H'''' , such that $\llbracket K_1 \rrbracket(w, G, H') \& H' \approx_0 H'' \& H''_0 = H'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& \llbracket K_2 \rrbracket(w, H'', H''') \& H''' \approx_0 H'''' \& H''''_0 = H'''_{comp} \cap \llbracket @E \rrbracket^{H'',w} \& H'''' \sim_{comp} G'' \& G''_{comp} = \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp})$.
7. Take such H', H'', H''' , and H'''' .
8. $\llbracket \text{ASSERT}(\text{NOT}(K_3)) \rrbracket(w, G', H)$ iff there is a G''' such that $\llbracket \text{NOT}(K_3) \rrbracket(w, G', G''') \& G''' \approx_0 H \& H_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G,w} \& w \in H_0$.
9. Take such G''' .
10. $\llbracket \text{NOT}(K_3) \rrbracket(w, G', G''')$ iff there is a G'''' such that $\llbracket K_3 \rrbracket(w, G', G''') \& G'''' \sim_{comp} G''' \& G'''_{comp} = \llbracket \neg comp \rrbracket^{G''',w}$.
11. Now, what is left to prove is that when $H'''_{comp} \cap G'''_{comp} = \emptyset$ & $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp}) \& w \in G'''_{comp}$, then $w \notin H'_{comp} \cap \llbracket @E \rrbracket^{G,w}$ (i.e. when the truth-conditional contribution of the small premise negates the truth-conditional contribution of the consequent, and both premises hold at a world w , then the antecedent does not hold at w).
12. *Per reductio*, suppose that $H'''_{comp} \cap G'''_{comp} = \emptyset$ & $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp}) \& w \in G'''_{comp}$, and also $w \in H'_{comp} \cap \llbracket @E \rrbracket^{G,w}$.
13. Then, w is such that $\forall w' : wRw'$, if $w' \in \llbracket @E \rrbracket^{G,w} \& w' \in H'_{comp}$ then $w' \in H'''_{comp}$, and also $w \in G'''_{comp}$.
14. Since, $w \in G'''_{comp}$, and $H'''_{comp} \cap G'''_{comp} = \emptyset$, then $w \notin H'''_{comp}$.
15. Suppose that $w \in \llbracket @E \rrbracket^{G,w} \& w \in H'_{comp}$. (Since R is reflexive and transitive, we know that wRw .)
16. Then, $w \in H'''_{comp}$.
17. Contradiction!
18. So, if $H'''_{comp} \cap G'''_{comp} = \emptyset$ & $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp}) \& w \in G'''_{comp}$, then $w \notin H'_{comp} \cap \llbracket @E \rrbracket^{G,w}$. □

- Note that whenever the proposition expressed by $\text{NOT}(K_3)$ negates the one expressed by K_2 , and the one expressed by $\text{NOT}(K_4)$ negates the one expressed by the proposition expressed by the intersection of the truth-conditional contribution of K_1 and $\llbracket @E \rrbracket^G, w$, where G is the input context for the conditional, given (§ B.1.4) and (§ B.1.5), $\text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{Assert}(\text{NOT}(K_3))$ will entail $\text{Assert}(\text{NOT}(K_4))$.

B.1.6.2 MP

As with MT, we need to have a way of individuating instances of MP, first. Let us begin with the following discourse:

1. $\text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{ASSERT}(K_3)$

What we want to show is that when the propositions expressed by K_3 and the one corresponding to the intersection of the propositions expressed by K_1 and ‘@E’ are identical, if the proposition expressed by $\text{IF}(@E, K_1, K_2)$ and K_3 hold (relative to a world of evaluation w), the proposition corresponding to the truth-conditional contribution of K_2 will be true (relative to a world of evaluation w , as well).

Proof.

2. $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{ASSERT}(K_3) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)) \rrbracket(w, G, G')$ and $\llbracket \text{ASSERT}(K_3) \rrbracket(w, G', H)$.
3. Take such a G, G' and H .
4. $\llbracket \text{ASSERT}(\text{IF}(@E, K_1, K_2)); \rrbracket(w, G, G')$ iff there is a G'' such that $\llbracket \text{IF}(@E, K_1, K_2) \rrbracket(w, G, G'') \ \& \ G'' \approx_0 G' \ \& \ G'_0 = G''_{comp} \cap \llbracket @E \rrbracket^{G, w} \ \& \ w \in G'_0$.
5. Take such a G'' .
6. $\llbracket \text{IF}(@E, K_1, K_2) \rrbracket(w, G, G'')$ iff H', H'', H''' and H'''' such that $\llbracket K_1 \rrbracket(w, G, H') \ \& \ H' \approx_0 H'' \ \& \ H''_0 = H'_{comp} \cap \llbracket @E \rrbracket^{G, w} \ \& \ \llbracket K_2 \rrbracket(w, H'', H''') \ \& \ H''' \approx_0 H'''' \ \& \ H''''_0 = H'''_{comp} \cap \llbracket @E \rrbracket^{H'', w} \ \& \ H'''' \sim_{comp} G'' \ \& \ G''_{comp} = \text{Cond}(\llbracket @E \rrbracket^{G, w}, H'_{comp}, H'''_{comp})$.
7. Take such H', H'', H''' , and H'''' .

8. $\llbracket \text{ASSERT}(K_3) \rrbracket(w, G', H)$ iff there is a G''' such that $\llbracket K_3 \rrbracket(w, G', G''') \ \& \ G''' \approx_0 H \ \& \ H_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G,w} \ \& \ w \in H_0$.
 9. It is easy to show given § B.1.4 that if $\llbracket (K_3) \rrbracket(w, G', G''')$, whatever form K_3 has, its truth-conditional bit will be stored as a value of G'''_{comp} .
 10. Suppose $\llbracket (K_3) \rrbracket(w, G', G''')$.
 11. So, now we just have to prove that when $G'''_{comp} = \llbracket @E \rrbracket^{G,w} \cap H'_{comp}$, and $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp})$, and $w \in G'_{comp}$, then $w \in H'''_{comp}$ (i.e. when the truth-conditional contribution of the small premise and the antecedent are identical, and both the antecedent and the small premise hold at w , then the consequent holds at w , as well).
 12. *Per reductio*, suppose that $G'''_{comp} = \llbracket @E \rrbracket^{G,w} \cap H'_{comp}$, and $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp})$ and $w \in G'_{comp}$, and also $w \notin H'''_{comp}$.
 13. Then w is such that $\forall w' : wRw'$, if $w' \in \llbracket @E \rrbracket^{G,w} \ \& \ w' \in H'_{comp}$ then $w' \in H'''_{comp}$, and also $w \in G'_{comp}$.
 14. By (12), $w \notin H'''_{comp}$.
 15. Then, by (14) and reflexivity of R , $\exists w' : wRw'$, and $w' \in \llbracket @E \rrbracket^{G,w} \ \& \ w' \in H'_{comp}$ and $w' \notin H'''_{comp}$, namely, w .
 16. Contradiction!
 17. So, if $G'''_{comp} = \llbracket @E \rrbracket^{G,w} \cap H'_{comp}$ and $w \in \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{comp}, H'''_{comp})$ and $w \in G'_{comp}$, then $w \in H'''_{comp}$. \square
- Note that whenever the proposition expressed by K_3 is identical to the one corresponding to the intersection of the propositions expressed by K_1 and $@E$, and the one expressed by K_2 identical to the one expressed by K_4 , given (§ B.1.4) and (§ B.1.5), $\text{ASSERT}(\text{IF}(@E, K_1, K_2)); \text{Assert}(K_3)$ will entail $\text{Assert}(\text{NOT}(K_4))$.

B.1.7 Relation between the dynamic (B.1.7.2) and classical (B.1.7.3) interpretations

I shall now prove that my dynamic interpretation preserves a classical one. To this end, I shall first give a dynamic translation for a fragment of English, specifying the updates associated with utterances containing modals and conditionals. Then I will give a classical translation for the same fragment, and prove that the dynamic interpretation preserves the truth-conditions assigned by the classical interpretation. For ease of comparison between the two translations, I shall avail myself of abstract level of logical forms (ALFs) for the relevant fragment of English. The reader should bear in mind that we do not have to take a stand on the existence of a level of representation corresponding to ALFs. This level of representation is merely a dispensable convenience that helps us compare the two interpretations in a streamlined way.

B.1.7.1 Abstract Logical Forms (ALFs) for a Fragment of English:

Terms:

- Propositional constants from our base language in § B.1.1 (set C).

Atoms:

- All terms are atoms, and nothing else is an atom.

ALFs:

- Atoms are ALFs.
- If ϕ and ψ are ALFs, then $might(\phi, \psi)$ is an ALF. (Stands for “it might be the case that ψ , given the restrictor ϕ ”.)
- If ϕ and ψ are ALFs, then $must(\phi, \psi)$ is an ALF. (Stands for “it must be the case that ψ , given the restrictor ϕ ”.)
- If ϕ and ψ are ALFs, then $likely(\phi, \psi)$ is an ALF. (Stands for “it’s likely the case that ψ , given the restrictor ϕ ”.)
- If ϕ, ψ and γ are ALFs, then $if(\phi, \psi, \gamma)$ is an ALF. (Stands for “given the restrictor ϕ if ψ , then γ ”.)

- If ϕ and ψ are ALFs, then $and(\phi, \psi)$ is an ALF. (Stands for “ ϕ and ψ .”)
- If ϕ is an ALF, then $not(\phi)$ is an ALFs. (Stands for “Not ϕ .”)
- If ϕ is an ALF, then $Assert(\phi)$ is an ALF. (Assertion operator—makes sure that the proposition ϕ is asserted.)

B.1.7.2 Dynamic Interpretation:

I first give a translation of the relevant fragment of English into our dynamically interpreted language (B.1.7.2.1), and then a translation of the same fragment into a classically interpreted language (B.1.7.3)

B.1.7.2.1 Dynamic Translations

In this section, I provide a translation of the relevant fragment of English, into our dynamically interpreted language defined in § B.1.1–§ B.1.4. I’ll assume the ALFs for the relevant fragment of English defined in B.1.7.1, (e.g. $might(\phi, \psi)$ for “it might be the case that ψ ”, where the modal is anaphorically dependent on $\llbracket \phi \rrbracket^{G,w}$, for an input context G .)

(Base case, where $T_d(\phi)$ is a translation of a formula ϕ into our dynamic system.)

- If ϕ is an atom, then $T_d(\phi) = \langle comp := \phi \rangle$.

(Recursive case)

- $T_d(might(\phi, \psi)) = \text{MIGHT}(\phi, T_d(\psi))$
- $T_d(must(\phi, \psi)) = \text{MUST}(\phi, T_d(\psi))$
- $T_d(likely(\phi, \psi)) = \text{LIKELY}(\phi, T_d(\psi))$
- $T_d(if(\phi, \psi, \gamma)) = \text{IF}(\phi, T_d(\psi), T_d(\gamma))$
- $T_d(and(\phi, \psi)) = \text{AND}(T_d(\phi), T_d(\psi))$
- $T_d(not(\phi)) = \text{NOT}(T_d(\phi))$

- $T_d(\text{Assert}(\phi)) = \text{ASSERT}(T_d(\phi))$

Next, I shall introduce the classical interpretation of modals and conditionals in B.1.7.3, and then go on to prove that our dynamic interpretation of modals and conditionals (B.1.7.2.1) preserves the classical interpretation in (B.1.7.4).

B.1.7.3 Classical Interpretation

I will now introduce a classical toy modal language, and provide a translation of the relevant fragment of English into this modal language, so I can compare our dynamic interpretation and classical interpretation. Assume the following modal language:

Terms:

- Propositional terms: propositional constants (p, q, r) .

Atomic formulae:

- All terms are atoms, and nothing else is an atom.

Now we introduce well-formed formulae:

- All atoms are well-formed formulae.
- If ϕ and ψ are formulae, then $\Diamond(\phi, \psi)$ is a well-formed formula.
- If ϕ and ψ are formulae, then $\Box(\phi, \psi)$ is a well-formed formula.
- If ϕ and ψ are formulae, then $\phi \rightarrow \psi$ is a well-formed formula.
- If ϕ and ψ are formulae, then $\phi \wedge \psi$ is a well-formed formula.
- Nothing else is a well-formed formula.

B.1.7.3.1 Classical Semantics:

I now define classical semantics for the simple modal language. I assume models are defined as in § B.1.2. (Assuming the definition of models in § B.1.2, and sets of sequences in § B.1.3), where R is

an accessibility relation provided by the model, and ϕ a restriction on the domain of quantification of a modal:

- $\llbracket p \rrbracket^{G,w} = w \in I(p)$
- $\llbracket \Diamond(\phi, \psi) \rrbracket^{G,w} = \{w \mid \exists w': wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w}, w' \in \llbracket \psi \rrbracket^{G,w}\}$
- $\llbracket \Box(\phi, \psi) \rrbracket^{G,w} = \{w \mid \forall w': \text{if } wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w} \text{ then } w' \in \llbracket \psi \rrbracket^{G,w}\}$
- $\llbracket \phi \wedge \psi \rrbracket^{G,w} = \llbracket \phi \rrbracket^{G,w} \cap \llbracket \psi \rrbracket^{G,w}$
- $\llbracket (\phi \wedge \psi) \rightarrow \gamma \rrbracket^{G,w} = \{w \mid \forall w': \text{if } wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w} \cap \llbracket \psi \rrbracket^{G,w} \text{ then } w' \in \llbracket \gamma \rrbracket^{G,w}\}$

B.1.7.3.2 Classical Translation:

Assuming the same ALFs for the relevant fragment of English as in B.1.7.1, I now define the following translations of the relevant bits of the fragment into our classical language.

Where $T_c(p)$ is a translation of a formula p into classical system:

- Where ϕ is an atom $T_c(\phi) = \llbracket \phi \rrbracket^{G,w}$
- $T_c(\text{might}(\phi, \psi)) = \llbracket \Diamond(\phi, \psi) \rrbracket^{G,w}$
- $T_c(\text{must}(\phi, \psi)) = \llbracket \Box(\phi, \psi) \rrbracket^{G,w}$
- $T_c(\text{if}(\phi, \psi, \gamma)) = \llbracket (\phi \wedge \psi) \rightarrow \gamma \rrbracket^{G,w}$

This translation does not capture the systematic effects of context on the interpretation of modals: it doesn't control for where the restrictor comes from. But, note that we don't have to care about what the restrictor ϕ is. We can simply assume that it is the domain of all possible worlds (\mathcal{D}_w), the set of epistemically live worlds, or any other proposition. So long as all the anaphoric restrictors are denoting the same proposition, what the restrictor is won't matter, because by resolving the restrictor always to the same proposition we are making sure that the domain of quantification for modals is held constant throughout, as in a simple classical modal logic.

B.1.7.4 Proof.

- We want to prove that our dynamic interpretation (B.1.7.2) preserves the classical interpretation (in B.1.7.3). In particular, we prove that for any $T_d(p)$, if $\llbracket T_d(p) \rrbracket(w, G, H)$, then $H_{comp} = \llbracket T_c(p) \rrbracket^{G,w}$, where $\llbracket T_c(p) \rrbracket^{G,w}$ is the corresponding translation of the formula p in classical system; $Assert(T_d(p))$ guarantees that $\llbracket T_c(p) \rrbracket^{G,w}$ is true at the actual world. We do a proof by induction.

- * Base case. First prove that for an atom p , and translation $T_d(p)$, if $\llbracket T_d(p) \rrbracket(w, G, H)$, then $H_{comp} = \llbracket p \rrbracket^w$.

Proof.

2. By (B.1.7.2.1), $T_d(p) = \langle comp := p \rangle$.
3. By (§ B.1.4), we have $\llbracket \langle comp := p \rangle \rrbracket(w, G, H)$ iff $G \underset{comp}{\sim} H$ & $H_{comp} = \llbracket p \rrbracket^{G,w}$.
4. Suppose $\llbracket \langle comp := p \rangle \rrbracket(w, G, H)$.
5. By (2)–(4), $H_{comp} = \llbracket p \rrbracket^{G,w}$, and $\llbracket p \rrbracket^{G,w}$ iff $w \in I(p)$. Thus, $H_{comp} = \llbracket T_c(p) \rrbracket^{G,w}$, which we were set to prove. \square

- * Recursive case.

Proof.

- **IH:** Assume that for a formula ϕ of the depth k or less, if $\llbracket T_d(\phi) \rrbracket(w, G, H)$, then $H_{comp} = \llbracket \phi \rrbracket^w$, where $\llbracket \phi \rrbracket^w$ is the corresponding classical interpretation of the formula ϕ .
- Consider a formula ϕ of the depth $k + 1$. We prove that the **IH** holds for the possible ways of constructing ϕ :

- i Let $\phi = \text{might}(\chi, \psi)$. Then, by B.1.7.2.1, $T_d(\phi) = \text{MIGHT}(\chi, T_d(\psi))$. Suppose that $\llbracket \text{MIGHT}(\chi, T_d(\psi)) \rrbracket(w, G, H)$. We know by (§ B.1.4) that $\llbracket \text{MIGHT}(\chi, T_d(\psi)) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket T_d(\psi) \rrbracket(w, G, G') \& G' \underset{0}{\approx} G'' \& G''_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& G'' \underset{comp}{\sim} H \& H_{comp} = M(\llbracket \chi \rrbracket^{G,w}, G'_{comp})$. Take such a G' and G'' . We have that $\llbracket T_d(\psi) \rrbracket(w, G, G')$; thus, by **IH**, $G'_{comp} = \llbracket \psi \rrbracket^{G,w}$. Then, since $H_{comp} = M(\llbracket \chi \rrbracket^{G,w}, G'_{comp})$, given Definition (B.1.1) and (B.1.7.3.2), H_{comp} is equivalent to $\llbracket T_c(\text{might}(\chi, \psi)) \rrbracket^{G,w}$, by simple math.

- ii Let $\phi = \text{must}(\chi, \psi)$. Then, by B.1.7.2.1, $T_d(\phi) = \text{MUST}(\chi, T_d(\psi))$. Suppose that $\llbracket \text{MUST}(\chi, T_d(\psi)) \rrbracket(w, G, H)$. So, by (§ B.1.4), there is a G' and G'' such that $\llbracket T_d(\psi) \rrbracket(w, G, G')$ & $G' \approx_0 G''$ & $G'_0 = G'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w}$ & $G'' \sim_{\text{comp}} H$ & $H_{\text{comp}} = N(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}})$. Take such a G' and G'' . Since $\llbracket T_d(\psi) \rrbracket(w, G, G')$, by **IH**, $G'_{\text{comp}} = \llbracket \psi \rrbracket^{G,w}$. Then, since $H_{\text{comp}} = M(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}})$, given Definition (B.1.2) and (B.1.7.3.2), H_{comp} is equivalent to $\llbracket T_c(\text{must}(\chi, \psi)) \rrbracket^{G,w}$, by simple math.
- iv Let $\chi = \text{if}(\chi, \psi, \gamma)$. Then, by B.1.7.2.1, $T_d(\chi) = \text{IF}(\chi, T_d(\psi), T_d(\gamma))$. Suppose that $\llbracket \text{IF}(\chi, T_d(\psi), T_d(\gamma)) \rrbracket(w, G, H)$. So, by (§ B.1.4), we know that there is a G', G'', G''' and G'''' such that $\llbracket T_d(\psi) \rrbracket(w, G, G')$ & $G' \approx_0 G''$ & $G'_0 = G'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w}$ & $\llbracket T_d(\gamma) \rrbracket(w, G'', G''')$ & $G''' \approx_0 G''''$ & $G'''_0 = G'''_{\text{comp}} \cap \llbracket @E \rrbracket^{G'',w}$ & $G'''' \sim_{\text{comp}} H$ & $H_{\text{comp}} = \text{Cond}(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}}, G'''_{\text{comp}})$. Take such a G', G'', G''' and G'''' . Since $\llbracket T_d(\psi) \rrbracket(w, G, G')$, by **IH** we know that $G'_{\text{comp}} = \llbracket \psi \rrbracket^{G,w}$ and since $\llbracket T_d(\gamma) \rrbracket(w, G'', G''')$, we know by **IH** that $G'''_{\text{comp}} = \llbracket \gamma \rrbracket^{G,w}$. Then, since $H_{\text{comp}} = \text{Cond}(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}}, G'''_{\text{comp}})$ given Definition (B.1.4) and (B.1.7.3.2), H_{comp} is equivalent to $\llbracket T_c(\text{cond}(\chi, \psi, \gamma)) \rrbracket^{G,w}$, by simple math. \square

B.1.8 Yalcin's Counterexample:

I shall now formally represent Yalcin's counterexample in our system and prove that the counterexample is not an instance of MT (and, thus, is not a counterexample). Recall the two premises of the counterexample:

- (1) If the marble is big, then it's likely red.
- (2) The marble is not likely red.

Let $I(p)$ be the proposition *that the marble is big*, and $I(q)$ be the proposition *that the marble is red*. Then we represent the two premises of the counterexample as in (32), repeated below:

$$(32) \quad \text{CONTRAST}(\text{ASSERT}(\text{IF}(@E, \langle \text{comp} := p \rangle, \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle)))), \text{ASSERT}(\text{NOT}(\text{LIKELY}(@E, \langle \text{comp} := q \rangle))))$$

Let ' K_1 ' stand for ' $\text{IF}(@E, \langle \text{comp} := p \rangle, \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle)))$ ' and ' K_2 ' for

‘NOT(LIKELY(@E, <comp := q>))’. Then, we show that (32) is not an instance of MT.

Proof. First, we shall calculate the truth-conditions expressed by the consequent of the big premise, then, we shall calculate the truth-conditions expressed by the small premise, and then we show that the two do not contradict each other.

1. $\llbracket \text{CONTRAST}(\text{ASSERT}(K_1); \text{ASSERT}(K_2)) \rrbracket(w, G, H)$ iff there is a G' and G'' such that
 $\llbracket \text{ASSERT}(K_1) \rrbracket(w, G, G') \ \& \ G' \approx_0 G'' \ \& \ G''_0 = \llbracket \theta_{\text{ASSERT}(K_1)} \rrbracket^{G,w} \ \& \ \llbracket \text{ASSERT}(K_2) \rrbracket(w, G'', H)$
 $\ \& \ \llbracket \theta_{\text{ASSERT}(K_1)} \rrbracket^{G,w} = \llbracket \theta_{\text{ASSERT}(K_2)} \rrbracket^{G'',w} \ \& \ \text{Contrast}(G''_{\text{comp}}, H_{\text{comp}})$. (By (§ B.1.4).)
2. Take such a G, G', G'' and H . Then $\llbracket \text{ASSERT}(K_1) \rrbracket(w, G, G')$ and
 $\llbracket \text{ASSERT}(K_2) \rrbracket(w, G'', H)$. (By (1).)
3. First we calculate the truth-condition of the consequent of the big premise.
 $\llbracket \text{ASSERT}(K_1) \rrbracket(w, G, G')$ iff there is a G''' such that $\llbracket K_1 \rrbracket(w, G, G''') \ \& \ G''' \approx_0 G' \ \& \ G'_0 =$
 $G'''_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w} \ \& \ w \in G'_0$. (By (§ B.1.4).)
4. Take such G''' . $\llbracket K_1 \rrbracket(w, G, G''')$ iff there is a H', H'', H''' and H'''' such that $\llbracket \langle \text{comp} :=$
 $p \rangle \rrbracket(w, G, H') \ \& \ H' \approx_0 H'' \ \& \ H'_0 = H'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w} \ \&$
 $\llbracket \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle)) \rrbracket(w, H'', H''') \ \& \ H''' \approx_0 H'''' \ \& \ H'''_0 = H'''_{\text{comp}} \cap$
 $\llbracket @E \rrbracket^{H'',w} \ \& \ H'''' \sim_{\text{comp}} G''' \ \& \ G'''_{\text{comp}} = \text{Cond}(\llbracket @E \rrbracket^{G,w}, H'_{\text{comp}}, H'''_{\text{comp}})$. (By (§ B.1.4)
and the definition of K_1 .) Note that by § B.1.4 and B.1.7.2.1, H'''_{comp} stores the truth-
conditions of the consequent of the big premise.
5. Take such H', H'', H''' and H'''' .
 $\llbracket \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle)) \rrbracket(w, H'', H''')$ iff there are some J and J' such
that $H'' \approx_0 J \ \& \ J_0 = \llbracket \mathbf{w}_0 \rrbracket^{H'',w} \ \& \ \llbracket \text{LIKELY}(@E, \langle \text{comp} := q \rangle) \rrbracket(w, J, J') \ \& \ J' \approx_0 H''' \ \&$
 $H'''_0 = J'_{\text{comp}} \ \& \ \text{Elab}(\llbracket \mathbf{w}_0 \rrbracket^{H'',w}, H'''_0)$. (By (§ B.1.4).)
6. Take such J and J' . $\llbracket \text{LIKELY}(@E, \langle \text{comp} := q \rangle) \rrbracket(w, J, J')$ iff there is a J'' and J'''
such that $\llbracket \langle \text{comp} := q \rangle \rrbracket(w, J, J'') \ \& \ J'' \approx_0 J''' \ \& \ J'''_0 = J''_{\text{comp}} \cap \llbracket @E \rrbracket^{J,w} \ \& \ J''' \sim_{\text{comp}} J' \ \&$
 $J'_{\text{comp}} = P(\llbracket @E \rrbracket^{J,w}, J'''_{\text{comp}})$. (By (§ B.1.4).)

7. Thus, by (B.1.3), $J'_{comp} = \{w | \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{J,w} \ \& \ w' \in J'_{comp}\}) / \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{J,w}\}) > .5\}$.
8. By 5 and the definition of \approx_n , $J'_{comp} = H'''_{comp}$.
9. By (4)–(6), and (B.1.7.2.1), $H'''_{comp} = \{w | \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket p \rrbracket^{G,w} \cap \llbracket @E \rrbracket^{G,w} \ \& \ w' \in I(q)\}) / \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket p \rrbracket^{G,w} \cap \llbracket @E \rrbracket^{G,w}\}) > .5\}$. These are the truth-conditions expressed by the consequent of the big premise. Now we calculate the truth-conditions of the small premise.
10. From (2) we have: $\llbracket \text{ASSERT}(K_2) \rrbracket(w, G'', H)$.
By (§ B.1.4), $\llbracket \text{ASSERT}(K_2) \rrbracket(w, G'', H)$ iff there is a I such that $\llbracket K_2 \rrbracket(w, G'', I) \ \& \ I \approx_0 H \ \& \ H_0 = I_{comp} \cap \llbracket @E \rrbracket^{G'',w} \ \& \ w \in H_0$.
11. Take such I . By definition of K_2 and (§ B.1.4), we have: $\llbracket K_2 \rrbracket(w, G'', I)$ iff there is a I' such that $\llbracket \text{LIKLEY}(@E, \langle comp := q \rangle) \rrbracket(w, G'', I') \ \& \ I' \sim_{comp} I \ \& \ I_{comp} = \llbracket \neg comp \rrbracket^{I',w}$.
12. Take such I' . Then, by (§ B.1.4), $\llbracket \text{LIKLEY}(@E, \langle comp := q \rangle) \rrbracket(w, G'', I')$ iff there is a I'' and I''' such that $\llbracket \langle comp := q \rangle \rrbracket(w, G'', I'') \ \& \ I'' \approx_0 I''' \ \& \ I'''_0 = I'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \ \& \ I''' \sim_{comp} I' \ \& \ I'_{comp} = P(\llbracket @E \rrbracket^{G'',w}, I'''_{comp})$.
13. By (11), (6), and (B.1.3), the truth-condition expressed by the small premise is as follows: $\mathcal{D}_w \setminus \{w | \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{G'',w} \ \& \ w' \in I'''_{comp}\}) / \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{G'',w}\}) > .5\}$, where \mathcal{D}_w is the domain of possible worlds from the model.
14. By (1), (B.1.6), (§ B.1.4) and (B.1.7.2.1), we have: $\mathcal{D}_w \setminus \{w | \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{G,w} \ \& \ w' \in I(q)\}) / \mathcal{P}(\{w' | wRw' \ \& \ w' \in \llbracket @E \rrbracket^{G,w}\}) > .5\}$, where \mathcal{D}_w is the domain of possible worlds from the model. This is the truth-condition expressed by the small premise.
15. From (9) and (13), we see that the truth-conditions expressed by the big premise and the consequent of the small one do not contradict each other.³ Hence, (32) does not correspond to the premises of an instance of MT. \square

³Yalcin's scenario with an urn with 100 marbles can be used to generate a model in which both propositions are true. In particular, suppose the domain of worlds \mathcal{D}_w is partitioned according to a color-size distribution: into big and blue, small and blue, big and red and small and red worlds. Where $I(p)$ is the proposition *that the marble is big*, $I(q)$

the proposition *that the marble is red*, $I(r)$ the proposition *that the marble is blue* and $I(s)$ the proposition *that the marble is small*, let the probability measure assign the following probabilities: $\mathcal{P}(I(p) \cap I(q)) = .3$, $\mathcal{P}(I(p) \cap I(r)) = .1$, $\mathcal{P}(I(s) \cap I(q)) = .1$, and $\mathcal{P}(I(s) \cap I(r)) = .5$.

Bibliography

- Alahverdzhieva, Katya and Lascarides, Alex;. 2011. “An HPSG Approach to Synchronous Deixis and Speech.” In Stefan Muller (ed.), *Proceedings of the 18th International Conference on Head-Driven Phase Structure Grammar (HPSG)*, 6–24. Seattle, USA.
- Altshuler, Daniel and Stojnić, Una. 2016. “The Meaning of ‘Now’.” Manuscript.
- Asher, Nicholas and Lascarides, Alex. 2003. *Logics of Conversation*. Cambridge University Press.
- Asher, Nicholas and McCready, Eric. 2006. “Modal Subordination in Japanese: Dynamics and Evidentiality.” *University of Pennsylvania Working Papers in Linguistics* 12.
- . 2007. “Were, Would, Might and a Compositional Account of Counterfactuals.” *Journal of Semantics* 24:93–129.
- Barker, Christopher. 2002. “Dynamics of Vagueness.” *Linguistics and Philosophy* 25:1–36.
- Barwise, Jon and Etchemendy, John. 1987. *The Liar*. Oxford University Press.
- Beaver, David and Clark, Brady. 2008. *Sense and Sensitivity: How focus determines meaning*. Oxford: Wiley- Blackwell.
- van Bentham, Johan. 1986. *Essays in Logical Semantics*, volume Studies in Linguistics and Philosophy. Dordrecht Reidel.
- van den Berg, Martin. 1996. *Some Aspects of the Internal Structure of Discourse: The Dynamics of Nominal Anaphora*. Ph.D. thesis, University of Amsterdam.
- Bittner, Maria. 2001. “Topical Referents for Individuals and Possibilities.” In Rachel Hastings (et al.) (ed.), *Proceedings from SALT XI*, 36 – 55. CLC Publications, Cornell University, Ithaca.
- . 2014. *Temporality: Universals and Variation*. Wiley-Blackwell.
- Bledin, Justin. 2014. “Logic Informed.” *Mind* 123:277–316.
- . forthcoming. “*Modus Ponens* Defended.” *Journal of Philosophy* .
- Brasoveanu, Adrian. 2010. “Decomposing Modal Quantification.” *Journal of Semantics* 27:437–527.
- Cantwell, John. 2008. “Changing the Modal Context.” *Theoria* 74:331–351.
- Cappelen, Herman and Hawthorne, John. 2009. *Relativism and Monadic Truth*. Oxford University Press.
- Charlow, Nate. 2015. “Prospects for an Expressivist Theory of Meaning.” *Philosophers’ Imprint* 15:1–43.

- Chomsky, Noam. 1981. *Lectures on Government and Binding: The Pisa Lectures*. Holland: Foris Publications. (Reprint: 7th Edition. Berlin and New York: Mouton de Gruyter, (1993).).
- Cohen, Stewart. 1998. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy* 76:289–306.
- Dekker, Paul. 2011. "Dynamic Semantics." In Klaus von Heusinger, Claudia Maienborn, and Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning, Volume 1*, 923–945. Berlin: Mouton De Gruyter.
- DeRose, Keith. 1995. "Solving the skeptical problem." *Philosophical Review* 104:1–52.
- . 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*, volume 1. Oxford University Press.
- Dorr, Cian and Hawthorne, John. 2014. "Embedding Epistemic Modals." *Mind* 122:867–913.
- Dowell, Janice L. 2011. "Flexible Contextualist Account of Epistemic Modals." *Philosophers' Imprint* 11:1–25.
- Dretske, Fred. 1981. "The Pragmatic Dimensions of Knowledge." *Philosophical Studies* 40:363–378.
- . 1991. "Knowledge: Sanford and Cohen." In Brian McLaughlin (ed.), *Dretske and His Critics*, 185–196. New York: Oxford University Press.
- Egan, Andy. 2007. "Epistemic Modals, Relativism and Assertion." *Philosophical Studies* 133:1–22.
- . 2010. "Disputing About Taste." In Richard Feldman and Ted A. Warfield (eds.), *Disagreement*, 247–292. Oxford: Oxford University Press.
- Egan, Andy, Hawthorne, John, and Weatherson, Brian. 2005. "Epistemic Modals in Context." In Gerhard Preyer and Peter Georg (eds.), *Contextualism in Philosophy*, 131 – 169. Oxford University Press.
- Elbourne, Paul. 2008. "Implicit Content and the Argument from Binding." In *Proceedings of Semantics and Linguistic Theory XVIII (SALT 18)*, 284–301. Ithaca: CLC Publications.
- Evans, Gareth. 1980. "Pronouns." *Linguistics Inquiry* 11:337–362.
- Fiengo, Robert and May, Robert. 1994. *Indices and Identity*. Cambridge, MA: MIT Press.
- . 2006. *De Lingua Belief*. Cambridge, MA: MIT Press.
- von Fintel, Kai. 1994. *Restrictions on Quantifier Domains*. Ph.D. thesis, University of Massachusetts-Amherst.
- von Fintel, Kai and Gillies, Anthony S. 2007. "An Opinionated Guide to Epistemic Modality." In Tamar Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, 32–62. Oxford University Press.
- . 2008. "CIA Leaks." *Philosophical Review* 117:77–98.
- . 2009. "'Might' Made Right." In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 108–130. NY: Oxford University Press.

- Frege, Gotlob. 1892. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik* 100:25–50.
- . 1918. "Der Gedanke." *Beiträge zur Philosophie des deutschen Idealismus* 1:58–77.
- Gillies, Anthony S. 2004. "Epistemic Conditionals and Conditional Epistemics." *Noûs* 38:585–616.
- . 2010. "Iffiness." *Semantics and Pragmatics* 3:1–42.
- Ginzburg, Jonathan. 2012. *The Interactive Stance: Meaning for Conversation*. Oxford University Press.
- Glanzberg, Michael. 2007a. "Context and Unrestricted Quantification." In Agustín Rayo and Gábor Uzquiano (eds.), *Absolute Generality*, 45–74. Oxford University Press.
- . 2007b. "Context, Content, and Relativism." *Philosophical Studies* 136:1–29.
- Grice, H.P. 1975. "Logic and Conversation." In Peter Cole and Jerry L. Morgan (eds.), *Syntax and semantics 3: Speech acts*, 41–58. New York: Academic Press.
- . 1989. "Indicative Conditionals." In *Studies in the Way of Words*, 58–85. Cambridge, MA: Harvard University Press.
- Grosz, Barbara J., Joshi, Aravind K., and Weinstein, Scott. 1995. "Centering: A framework for modelling the local coherence of Discourse." *Computational Linguistics* 21:203–225.
- Halliday, M.A.K. 1967. "Notes on Transitivity and Theme in English (Part 2). :." *Journal of Linguistics* 3:199–244.
- Hamann, Cornelia. 1991. "Adjectival Semantics." In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 657–673. de Gruyter, New York.
- Hawthorne, John. 2004s. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts.
- Hobbs, Jerry R. 1979. "Coherence and Coreference." *Cognitive Science* 3:67 – 90.
- . 1990. *Literature and Cognition*, volume (CSLI Lecture Notes 21). University of Chicago Press.
- Holliday, Wesley H. and Icard, Thomas F. III. 2013. "Measure Semantics and Qualitative Semantics for Epistemic Modals." *Proceedings of SALT* 23:514–534.
- Kameyama, Megumi. 1996. "Indefeasible Semantics and Defeasible Pragmatics." In M Kanazawa, C Pinon, and H de Swart (eds.), *Quantifiers, Deduction and Context*, 111–138. CSLI Publications.
- Kamp, Hans. 1981. "A Theory of Truth and Semantic Representation." In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof (eds.), *Formal Methods in the Study of Language*. Mathematisch Centrum, University of Amsterdam, Amsterdam.
- Kaplan, David. 1989a. "Afterthoughts." In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 565–614. Oxford University Press.

- . 1989b. “Demonstratives.” In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 481–563. Oxford University Press.
- Kehler, Andrew. 2002. *Coherence, Reference and the Theory of Grammar*. CSLI Publications.
- Kehler, Andrew, Kertz, L, Rodhe, Hannah, and Elman, J. 2008. “Coherence and Coreference Revisited.” *Journal of Semantics* 25:1–44.
- Keiser, E. 2009. “Effects of Anaphoric Dependencies and Semantic Representations on Pronoun Interpretation.” In *Anaphora Processing and Applications*, 121 – 130. Heidelberg: Springer.
- Kendon, Adam. 2004. *Gesture: Visible Action as Utterance*. Cambridge University Press.
- Kennedy, Christopher. 2007. “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives.” *Linguistics and Philosophy* 30:1–45.
- King, Jeffrey C. 2014a. “The Metasemantics of Contextual Sensitivity.” In Alexis Burgess and Bret Sherman (eds.), *New Essays on Metasemantics*, 97–118. Oxford University Press.
- . 2014b. “Speaker Intentions in Context.” *Noûs* 48:219–237.
- King, Jeffrey C. and Stanley, Jason. 2005. “Semantics, Pragmatics and the Role of Semantic Content.” In Zoltán Gendler Szabó (ed.), *Semantics vs. Pragmatics*, 111–164. Oxford University Press.
- Klein, Ewan. 1991. “Comparatives.” In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 673–691. de Gruyter, New York.
- Knott, Alistair. 1996. *A Data Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, Department of Artificial Intelligence, University of Edinburgh.
- Kolodny, Niko and MacFarlane, John. 2010. “Ifs and Oughts.” *Journal of Philosophy* 107:115–143.
- Kratzer, Angelika. 1977. “What *Must* and *Can* Must and Can Mean.” *Linguistics and Philosophy* 1:337 – 355.
- . 1981. “Notional Category of Modality.” In Hans-Jürgen Einkmeyer and Hannes Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, 38 – 74. Berlin: de Gruyter.
- . 1983. “Conditionals.” In Anne M. Farley, Peter Farley, and Karl-Erik McCullough (eds.), *Papers From the Parasession on Pragmatics and Grammatical Theory*, 1–15. Chicago: Chicago Linguistic Society.
- . 1991. “Modality.” In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, 639–50. Berlin: de Gruyter.
- . 1998. “An Investigation of the Lumps of Thought.” *Linguistics and Philosophy* 12:607–653.
- . 2012. *Modals and Conditionals*. Oxford University Press.
- . 2014. “Situations in Natural Language Semantics.” <http://plato.stanford.edu/entries/situations-semantics/> .

- Kripke, Saul. 1977. "Speaker's Reference and Semantic Reference." In Peter A. French, Theodore E. Uehlig, and Howard Wettstein (eds.), *Midwest Studies in Philosophy*, volume II, 255–276. University of Minnesota Press.
- . 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- van Kuppevelt, Jan. 1995. "Discourse Structure, Topicality and Questioning." *Journal of Linguistics* 31:109–147.
- Lappin, Shalom. 1993a. "Ellipsis Resolution at S-Structure." In *Proceedings of NELS*, volume 23, 255–269. University of Massachusetts, Amherst.
- . 1993b. "The Syntactic Basis of Ellipsis Resolution." In Steve Berman Arild Hestvik (ed.), *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich*, volume 340, Bericht Nr. 29–1992, SFB 340. University of Stuttgart, University of Tuebingen and IBM Germany.
- . 1996. "The Interpretation of Ellipsis." In Shalom Lappin (ed.), *Handbook of Contemporary Syntactic Theory*, 145–175. Blackwell.
- Lepore, Ernest and Cappelen, Herman. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell Publishing.
- Lepore, Ernest and Stone, Matthew. 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.
- Lewis, David K. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.
- . 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8:339 – 359.
- Lewis, Karen. ms. "The Speaker Authority Problem for Context-sensitivity." Manuscript. Barnard College, Columbia University Department of Philosophy.
- Locke, John. 1689. *An Essay Concerning Human Understanding*. Peter H. Nidditch (ed.) Oxford: Calderon Press (1975).
- MacFarlane, John. 2011. "Epistemic Modals are Assessment-Sensitive." In Brian Weatherson and Andy Egan (eds.), *Epistemic Modality*, 144–178. NY: Oxford University Press.
- . 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press.
- Mann, William C and Thompson, Sandra A. 1988. "Rhetorical Structure Theory: Toward a functional theory of text organization." *Text* 8:243–281.
- May, Robert. 1977. *The Grammar of Quantification*. Ph.D. thesis, MIT.
- McGee, Van. 1985. "A Counterexample to Modus Ponens." *Journal of Philosophy* 92:462 – 471.
- Moens, Marc and Steedman, Mark. 1988. "Temporal Ontology and Temporal Reference." *Computational Linguistics* 14:15–28.
- Moss, Sarah. 2015. "On the Semantics and Pragmatics of Epistemic Modals." *Semantics and Pragmatics* 8:1–81.

- Muskens, Reinhard. 1996. "Combining Montague Semantics and Discourse Representation." *Linguistics and Philosophy* 19:143–186.
- Neale, Stephen. 1990. *Descriptions*. Cambridge, MA: MIT Press.
- . 2004. "This, That and the Other." In Marga Reimer and Anna Bezuidenhout (eds.), *Descriptions and Beyond*, 68–182. NY: Oxford University Press.
- Partee, Barbara. 1973. "Some Structural Analogies between Tenses and Pronouns in English." *The Journal of Philosophy* 70:601–609.
- . 1984. "Nominal and Temporal Anaphora." *Linguistics and Philosophy* 7:243–286.
- Perry, John. 1979. "The Problem of the Essential Indexical." *Noûs* 13:3–21.
- Potts, Christopher. 2005. *The Logic Of Conventional Implicatures*. Oxford University Press.
- Recanati, François. 1996. "Domains of Discourse." *Linguistics and Philosophy* 19:445–475.
- Reimer, Marga. 1992. "Three Views of Demonstrative Reference." *Synthese* 93:373–402.
- Roberts, Craige. 1989. "Modal Subordination and Pronominal Anaphora in Discourse." *Linguistics and Philosophy* 12:683 – 721.
- . 1996a. "Information Structure: Towards an integrated formal theory of pragmatics." In Jae Hak Yoon and Andreas Kathol (eds.), *OSUWPL Volume 49: Papers in Semantics*, 91–136. The Ohio State University Department of Linguistics.
- . 1996b. "Information Structure: Towards an integrated formal theory of pragmatics." *OSUWPL: Papers in Semantics* 49.
- . 2002. "Demonstratives as Definites." In Kees van Deemter and Rodger Kibble (eds.), *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*, 89–196. Stanford: CSLI Press.
- . 2004. "Context in Dynamic Interpretation." In Lawrence R. Horn and Gregory Ward (eds.), *The Handbook of Pragmatics*, 197–220. Blackwell.
- . 2011. "Solving for Interpretation." Manuscript. Ohio State University.
- Rooth, Mats. 1992. "A Theory of Focus Interpretation." *Natural Language Semantics* 1:75–116.
- Rotschild, Daniel and Yalcin, Seth. 2015. "On the Dynamics of Conversation." *Noûs* 49.
- Russell, Bertrand. 1905. "On Denoting." *Mind* 14:479–493.
- Schaffer, Jonathan and Szabó, Zoltán Gendler. 2013. "Epistemic comparativism: a contextualist semantics for knowledge ascriptions." *Philosophical Studies* 2:1–53.
- Schlenker, Philippe. 2013. "Temporal and Modal Anaphora in Sign Language (ASL)." *Natural Language and Linguistic Theory* 31:207–234.
- Schoubye, Anders J. and Stokke, Andreas. 2015. "What is Said?" *Nous* 49:1–35.
- Sidner, Candice. 1983. "Focusing in the Comprehension of Definite Anaphora." In Michael Brady and Robert C. Berwick (eds.), *Computational Models of Discourse*, 267–330. MIT Press.

- Smyth, Ron. 1994. "Grammatical Determinants of Ambiguous Pronoun Resolution." *Journal of Psycholinguistic Research* 23:197–229.
- Stalnaker, Robert. 1978. "Assertion." *Journal of Semantics and Syntax* 9:315–332.
- Stanley, Jason. 2000. "Context and Logical Form." *Linguistics and Philosophy* 23:391–434.
- . 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, Jason and Szabó, Zoltán Gendler. 2000. "On Quantifier Domain Restriction." *Mind and Language* 15:219–261.
- Starr, William. forthcoming. "Dynamic Expressivism about Deontic Modality." In Nate Charlow and Matthew Chrisman (eds.), *Deontic Modality*. Oxford University Press.
- Stojnić, Una. 2016. "The Metasemantics of Domain Restriction." Manuscript. Rutgers University.
- . forthcoming. "One Man's *Modus Ponens*...: Modality, Coherence and Logic." *Philosophy and Phenomenological Research*.
- Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2013. "Deixis (Even Without Pointing)." *Philosophical Perspectives* 27:502 – 525.
- . 2014. "Discourse and Logical Form." Ms.
- . forthcoming. "Distinguishing Ambiguity from Underspecificity." In Ken Turner and Lawrence Horn (eds.), *An Atlas of Meaning*, volume Current Research in the Semantics/Pragmatics Interface. Leiden: Brill.
- Stone, Matthew. 1997. "The Anaphoric Parallel Between Modality and Tense." *IRCS Report* 97–06, *University of Pennsylvania*.
- . 1999. "Reference to Possible Worlds." Technical Report 49, Rutgers University, New Brunswick, RuCCS.
- Stone, Matthew and Stojnić, Una. 2015. "Meaning and Demonstration." *Review of Philosophy and Psychology* 6:69–97.
- Swanson, Eric. 2006. *Interactions with context*. Ph.D. thesis, MIT.
- Tonhauser, Judith, Beaver, David, Roberts, Craige, and Simons, Mandy. 2013. "Toward a Taxonomy of Projective Content." *Language* 89:66–109.
- Travis, Charles. 1989. *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford: Oxford University Press.
- Vallduví, Enric. 1992. *The Informational Component*. Garland Press, New York.
- . 1993. "Information packaging: A Survey." Report of the word order, prosody, and information structure initiative, University of Edinburgh.
- Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.
- Walker, Marilny A., Masayo, Iida, and Cote, Sharon. 1994. "Japanese Discourse and the Process of Centering." *Computational Linguistics* 20:193–232.

- Webber, Bonnie L. 1988. "Tense and Discourse Anaphor." *Computational Linguistics* 14:61–73.
- Webber, Bonnie L., Stone, Matthew, Joshi, Aravind, and Knott, Alistair. 2003. "Anaphora and Discourse Structure." *Computational Linguistics* 29:545–587.
- Wettstein, Howard K. 1984. "How to Bridge the Gap Between Meaning and Reference." *Synthese* 58:63–84.
- Wilkins, David. 2003. "Why Pointing With the Index Finger is Not a Universal (in Sociocultural and Semiotic Terms)." In S Kita (ed.), *Pointing: Where Language, Culture and Cognition Meet*. Lawrence Erlbaum Associates.
- Willer, Malte. 2014. "Dynamic Thoughts on Ifs and Oughts." *Philosophers' Imprint* 14:1–30.
- Winograd, Terry. 1972. *Understanding Natural Language*. Academic Press.
- Wolf, Florian, Gibson, Edward, and Desmet, Timothy. 2004. "Discourse Coherence and Pronoun Resolution." *Language and Cognitive Processes* 19:665–675.
- Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116:983 – 1026.
- . 2010. "Probability Operators." *Philosophy Compass* 5:916–937.
- . 2011. "Nonfactualism about Epistemic Modality." In Brian Weatherson and Andy Egan (eds.), *Epistemic Modality*, 295–332. NY: Oxford University Press.
- . 2012a. "Bayesian Expressivism." *Proceedings of the Aristotelian Society* 112:123–160.
- . 2012b. "A Counterexample to *Modus Tollens*." *Journal of Philosophical Logic* 41:1001 – 1024.
- . 2015. "Semantics and Metasemantics in the Context of Generative Grammar." In Alexis Burgess and Bret Sherman (eds.), *Metasemantics: New Essays on the Foundations of Meaning*, 17–54. Oxford University Press.
- Yalcin, Seth and Knobe, Joshua. 2013. "Epistemic Modals and Context: Experimental Data." *unpublished* 1–21.